

Volume XI ♦ Issue 1

2020

Logos & Episteme

an international journal
of epistemology

**Romanian Academy
Iasi Branch**



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

Founding Editor

Teodor Dima (1939-2019)

Editorial Board

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

<http://logos-and-episteme.acadiasi.ro/>

https://www.pdcnet.org/pdc/bvdb.nsf/journal?openform&journal=pdc_logos-episteme

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, România

Panayot Butchvarov

University of Iowa, USA

Mircea Dumitru

Universitatea din București, România

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 Mittelstrab

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

Nicolas Rescher

University of Pittsburgh, USA

Rahman Shahid

Université Lille 3, France

Ernest Sosa

Rutgers University, USA

Alexandru Surdu

Academia Română, România

John F. Symons

University of Texas at El Paso, USA

TABLE OF CONTENTS

RESEARCH ARTICLES

Frederik J. ANDERSEN, Uniqueness and Logical Disagreement.....	7
Daniele BERTINI, The Conflict of Rigidity and Precision in Designation.....	19
Nicole DULAR, Harder, Better, Faster, Stronger: Epistemic Standards and Moral Beliefs.....	29
Kok Yong LEE, Stakes-Shifting Cases Reconsidered—What Shifts? Epistemic Standards or Position?.....	53
Timothy PERRINE, On Some Arguments for Epistemic Value Pluralism.....	77
Gregory STOUTENBURG, Luck, Knowledge, and Epistemic Probability.....	97

DISCUSSION NOTES/ DEBATE

James SIMPSON, Gettier Beliefs and Serious Beliefs: A Reply to Biro and Forrai.....	113
Notes on the Contributors.....	119
<i>Logos and Episteme</i> . Aims and Scope.....	121
Notes to Contributors.....	123

RESEARCH ARTICLES

UNIQUENESS AND LOGICAL DISAGREEMENT

Frederik J. ANDERSEN

ABSTRACT: This paper discusses the *uniqueness thesis*, a core thesis in the epistemology of disagreement. After presenting uniqueness and clarifying relevant terms, a novel counterexample to the thesis will be introduced. This counterexample involves *logical disagreement*. Several objections to the counterexample are then considered, and it is argued that the best responses to the counterexample all undermine the initial motivation for uniqueness.

KEYWORDS: the uniqueness thesis, rational uniqueness, logical disagreement, propositional justification

1. Introduction

The uniqueness thesis (henceforth denoted 'UT') concerns a relation between a body of evidence, a doxastic attitude and a proposition. Jonathan Matheson, a proponent of the thesis, defines UT as follows:

(UT) For any body of evidence E and proposition P, E justifies at most one doxastic attitude toward P.¹

UT features frequently in the epistemology literature² and is motivated by arguments concerning peer disagreement—if two epistemic peers³ disagree about a

¹ Quote from Jonathan Matheson, "The Case for Rational Uniqueness," *Logos & Episteme* II, 3 (2011): 360.

² See for example Thomas Kelly, "Evidence Can Be Permissive," in *Contemporary Debates in Epistemology*, eds. Matthias Steup, John Turri, and Ernest Sosa (Wiley-Blackwell, 2013), 298–312, Roger White, "Evidence Cannot Be Permissive," in *Contemporary Debates in Epistemology*, eds. Steup, Turri, and Sosa, 312–323, Luis Rosa, "Justification and the Uniqueness Thesis," *Logos & Episteme* III, 4 (2012): 571–577, Jonathan Matheson, "The Case for Rational Uniqueness," *Logos & Episteme* II, 3 (2011): 359–373, Earl Conee, "Rational Disagreement Defended," in *Disagreement*, eds. Richard Feldman and Ted A. Warfield (Oxford University Press, 2010), 69–90.

³ Roughly put, two agents in disagreement are epistemic peers when neither side is epistemically superior with respect to the proposition at hand, i.e., when the two are similar enough in all relevant factors such as evidence, track record, time constraints etc.

proposition P, is it then possible that they are both justified in their doxastic attitudes toward P? If UT is true, then the answer is negative.

Importantly, there are in fact several non-equivalent definitions of UT in the literature.⁴ Thomas Kelly, for example, favors a formulation of UT saying that there is *exactly one* justified doxastic attitude given a body of evidence,⁵ while Matheson prefers *at most one*, as we have just seen. Matheson notes that in most cases there will be exactly one justified doxastic attitude given a body of evidence, but in some situations, there may be no justified doxastic attitude toward P whatsoever. This can arguably happen when one is not able to, or when it is simply not possible to, comprehend the proposition at hand.⁶ If one takes comprehension of P to be a necessary condition for the existence of a justified doxastic attitude toward P, then it seems most reasonable to use Matheson's weaker definition of UT. Thus, this is what we will assume here. Further, we will adopt Matheson's assumption that the term 'doxastic attitude' can only refer to the following three possibilities: *belief that P*, *disbelief that P* and *suspension of judgement with respect to P*, i.e., the possibility space of attitudes that one can take toward a proposition P is exhausted by these three attitudes.⁷ Now, UT puts a constraint on the total number of doxastic attitudes that a body of evidence can justify toward a proposition. According to UT any body of evidence E justifies at most one doxastic attitude toward P. In other words, according to UT, there exists no body of evidence E such that E justifies both belief and disbelief toward P. Similarly, of course, the thesis implies that there exists no E such that E justifies both a (dis)belief in P and suspension of judgement with respect to P. In the paper "The Case for Rational Uniqueness," Matheson makes two further clarifying remarks about UT:

(UT) [...] makes no reference to individuals or times since (UT) claims (in part) that who possesses the body of evidence, as well as when it is possessed, makes no difference regarding which doxastic attitude is justified (if any) toward any particular proposition by that body of evidence.⁸

(UT) concerns propositional justification, rather than doxastic justification. That is, the kind of justification relevant to (UT) is solely a relation between a body of

⁴ This is noted by Matheson, "The Case for Rational Uniqueness," 360-361.

⁵ Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, eds. Feldman and Warfield, 119.

⁶ See Richard Feldman, "Epistemological Puzzles About Disagreement," in *Epistemology Futures*, ed. S. Hetherington (Oxford University Press, 2006) for a motivation of this view.

⁷ This assumption is common in the contemporary literature, see for example Rosa, "Justification and the Uniqueness Thesis," Matheson, "The Case for Rational Uniqueness," Kelly, "Peer Disagreement and Higher-Order Evidence."

⁸ Quote from Matheson, "The Case for Rational Uniqueness," 360.

evidence, a doxastic attitude, and a proposition. How individuals have come to have the doxastic attitudes they have toward the proposition in question will not be relevant to our discussion. Further, individuals can be propositionally justified in adopting attitudes toward propositions which they psychologically cannot adopt [...] Importantly, it is not a necessary condition for being justified in believing p that one be able to demonstrate that one is justified in believing.⁹

The first of these quotes states that according to UT a given body of evidence E justifies exactly the same doxastic attitude (if any) towards P , no matter the subject that assesses E and at what time this is done. In the second quote, Matheson distinguishes between *propositional* and *doxastic* justification, where the former is a relation between a body of evidence, a doxastic attitude and a proposition, the latter concerns *how* a given individual came to adopt a specific doxastic attitude towards a proposition, i.e., doxastic justification is concerned with one's reasons for actually adopting a certain attitude toward P . Doxastic justification presumes that a given individual has a certain attitude toward P , and the question is then whether or not this individual has sufficient reason to be justified in having that attitude. When it comes to propositional justification, on the other hand, it is irrelevant whether any individual is ever concerned with P ; the crux of propositional justification is that a justification-relation between a body of evidence, a doxastic attitude and a proposition holds, not whether any individual realizes this. Understood in this way propositional justification refers to an external relation, and an individual can accordingly be propositionally justified in a doxastic attitude towards P even though this individual has not adopted the relevant attitude psychologically. And hence, it is not necessary for a subject to be able to demonstrate or defend this given attitude towards P in order for it to be propositionally justified. Matheson tells us that UT is a thesis concerning propositional justification rather than doxastic justification.

2. Clarifications

Before we move on to consider the announced counterexample to UT, let us pause to further specify what is meant by 'justification' and 'evidence' in the rest of the text. We will deliberately stay on a high level of generality in order not to exclude too many accounts of justification and evidence from the later discussions in sections 3 and 4.

When using the term 'justification,' this use is naturally restricted to the epistemic domain, we are not concerned with any practical issues whatsoever. So, in other words, our concern is with the justification of doxastic attitudes towards

⁹ Quote from Matheson, "The Case for Rational Uniqueness," 360-361.

propositions. This kind of justification is regulated by epistemic norms, i.e., truth-conducive norms, and as indicated in section 1, we are concerned with *propositional* justification, rather than doxastic justification.¹⁰

Our use of the term ‘evidence’ assumes that we can all agree that evidence can stem from many different sources like direct visual perception, testimony from individuals or media, scientific experiments etc. The only constraints we will force on our understanding of evidence from the outset are: (1) evidence must be propositional (and thus truth-apt), (2) any piece of evidence must be true, (3) any piece of evidence must (at least in principle) be accessible to human beings, and (4) evidence should be supportive of doxastic attitudes, where *support* may be interpreted probabilistically, but does not have to be.

(2) is arguably the most controversial among these four constraints. However, for our purposes there is a very good reason for including this factivity condition. To see this, suppose that one could have false (misleading) pieces of evidence in one’s body of evidence E. Then, given the further assumption that false evidence can support anything, we could easily have a situation where a true bit of evidence e_1 supports the belief that P, while a false bit of evidence e_2 supports the belief that not-P. This would in effect trivialize the debate about UT; on this account of evidence UT is obviously false.¹¹ Hence, we should either accept that evidence is factive or we should deny that false evidence can support anything. For the rest of the paper we will take the first option.

3. The Argument from Logical Disagreement

Consider now the following case against UT:

Logical Disagreement. Two logicians, S1 and S2, are walking into an empty auditorium where they find a deduction written on a blackboard. S1 and S2 are simultaneously looking at the board. As it happens, S1 is a classical logician, while S2 is an intuitionist. Now, the deduction consists in a finite number of steps, so all steps of the deduction except for the conclusion C will serve as a common body of

¹⁰ The literature on epistemic justification is vast, but prominent examples of theories of justification can be found in: Richard Feldman and Earl Conee, *Evidentialism* (Oxford University Press, 2004), Timothy Williamson, *Knowledge and Its Limits* (Oxford University Press, 2000), Ernest Sosa, *Knowledge in Perspective* (Cambridge University Press, 1991), William Alston, *Epistemic Justification* (Cornell University Press, 1989), Alvin Goldman, *Epistemology and Cognition* (Harvard University Press, 1986), Laurence Bonjour, *The Structure of Empirical Knowledge* (Harvard University Press, 1985), Richard Feldman and Earl Conee, “Evidentialism,” *Philosophical Studies* 48 (1985): 15-34, Alvin Goldman, “What Is Justified Belief?,” in *Justification and Knowledge*, ed. G. Pappas (Springer, 1979), 1-23.

¹¹ Thanks to Francesco Berto for pressing this point about false (misleading) evidence.

evidence E, i.e., a set of propositions that are represented in a language that both logicians fully comprehend. The central question is then whether E entails C. Suppose that C on line n is the result of applying DNE (double negation elimination) to not-not-C on line $n - 1$.¹² As S₁ accepts classical logic, she also accepts the inference from not-not-C to C, while S₂, given her intuitionist convictions, denies DNE as a rule of inference and thus denies that C comes out supported by E.

In this case we have a situation in which two agents possess exactly (!) the same evidence (the propositions represented by lines $n - 1$ on the blackboard), but they are justified in diverging doxastic attitudes towards the relevant proposition in question, namely C. We see that E justifies S₁ in her belief that P, while E justifies (at least) suspension of judgement regarding P for S₂ (P is not supported by E). Thus, the case is a clear counterexample to UT as the number of attitudes that E justifies exceeds one. Of course, as the reader will have noticed by now, the case is concerned with a special type of evidence, i.e., evidence of the completely formal type that we find in pure logic and mathematics. This means that the counterexample is narrow in the sense that it does not indicate the existence of counterexamples to UT among other types of evidence.¹³ However, this will be completely irrelevant as long as we regard UT as a general epistemic principle. If

¹² Using standard notation DNE is an inference from $\Gamma \vdash \neg\neg\phi$ to $\Gamma \vdash \phi$, where ‘ Γ ’ denotes a set of sentences in a given language, ‘ \vdash ’ denotes deducibility from left to right and ‘ ϕ ’ picks out a single sentence of the language. Some readers may point out that it is underspecified in the case above whether S₁ and S₂ disagree over an *instance* or a *schema* of DNE. This is true, but it will not make a significant difference to the main argument of the paper.

¹³ However, some epistemologists have suggested that there are counterexamples to UT among other types of evidence. Consider, for example, a case where S₁ and S₂ discuss which football team will win the national league this season. Suppose that their discussion takes place the day before the final match day, and at this point of the season only two teams can win; either team A or team B. Suppose further that the only evidence available to the subjects is a certain newspaper statistic, which shows the scores of the season so far. According to this statistic, team A is in front of team B by the smallest possible margin. Now, S₁ is convinced that team A will take the championship due to the statistical support for this (they are ahead at this point). However, S₂ suspends judgement about who will be the champions as team A leads with the smallest possible margin and it is still possible for team B to make it. In such a case the proponent of UT should say that at most one of the subjects’ doxastic attitudes is justified, but one might reasonably argue that this is wrong. In such borderline cases it seems that at least two out of three doxastic attitudes could be justified. If this is right, we have a counterexample to UT using another type of evidence, i.e., empirical data. Find similar borderline cases in Kelly, “Evidence Can Be Permissive,” 299-300.

the case holds, we will have the necessary and sufficient counterexample needed to reject UT.

4. Objections and Responses

As the case presented above will be very hard to accept for many readers (for various reasons), the rest of the paper aims to motivate the argument from logical disagreement. The strategy here is simple. While discussing various objections to **Logical Disagreement**, it will become clear that the UT-proponent can only avoid the counterexample by undermining the initial motivation behind UT, i.e., explaining away the counterexample to UT will lead to an indirect defeat of the thesis. In the following, five objections to **Logical Disagreement** will be scrutinized (subsections 4.1-4.5). The first two will simply be rejected, the third will be found underdeveloped, and while the remaining two can actually explain away the counterexample to UT, this can only be done by undermining the motivation behind the principle.

4.1 Evidence Is Contingent

Objection 1. Even though the evidence E present in **Logical Disagreement** satisfies our four rudimentary constraints on evidence (cf. section 2) as E is propositional, factive, accessible and supportive, E is still not a genuine body of evidence. This is because only contingent propositions can be evidence. Thus, UT is not even applicable in **Logical Disagreement**.

First of all, there is no principle reason why necessary propositions such as the ones found in pure mathematics and logic cannot be counted as evidence. Propositions of logic and mathematics can clearly serve the supportive role of evidence very well, i.e., such propositions speak in favor of certain hypotheses in the strongest possible way (by entailment). Hence, if any proposition is able to justify a belief, it seems that pure logical or mathematical propositions are ideal candidates. Habit may dictate, perhaps leading back to acceptance of Hume's Fork, that some of us cannot see the point in taking purely formal premises of deductive arguments as evidence, but without further qualification this is obviously not a good argument for accepting such an exclusion in philosophical or scientific work. Moreover, accepting **Objection 1** leads to absurd consequences when we hold other plausible epistemic principles to be true. Take for example Timothy Williamson's principle $E = K$, i.e., evidence equals knowledge.¹⁴ If we accept that our evidence is coextensive with our knowledge, and that **Objection 1** holds, it directly follows

¹⁴ Williamson, *Knowledge and Its Limits*, chapter 9.

that we cannot have pure mathematical or logical knowledge. To deny that we can and do have such knowledge would not only be absurd, it would be intellectual suicide.

4.2 Communication Breakdown

Objection 2. The case **Logical Disagreement** misrepresents the interaction between classical logicians and intuitionists. Where the classical logician works with a philosophical presupposition of a world of mathematical objects independent of the thinking subject (objects that obey the laws of classical logic and can stand in set-theoretic relations), this is radically different from the intuitionists who advocate for constructive methods and take mathematics to be about mental constructions. As a result of this schism, the two logicians in the proposed case would run into an insurmountable communication breakdown, i.e., the DNE-inference acceptable to the classical logician would not even be understandable to the intuitionist – it would be nonsense. To quote Brouwer: *“Let us now consider the concept: ‘denumerably infinite ordinal number.’ From the fact that this concept has a clear and well-defined meaning for both formalist and intuitionist, the former infers the right to create the ‘set of all denumerably infinite ordinal numbers,’ the power of which he calls aleph-one, a right not recognized by the intuitionist.”*¹⁵ Something similar to what Brouwer describes in the interaction between diverse logical traditions in this quote occurs in **Logical Disagreement** with respect to DNE, i.e., the intuitionist does simply not comprehend the final step of the deduction on the blackboard. Thus, suspension of judgement is not a justified doxastic attitude for the intuitionist in this case; the supposed logical connection between E and C is gibberish to her. Rather, **Logical Disagreement** represents the kind of case where there is no justified doxastic attitude for the intuitionist to have. Hence, UT would be saved (at least the *at most one doxastic attitude*-version of the thesis). The case allows only one justified attitude, namely the attitude of the classical logician.

This objection overstates the divide between the classical and intuitionist traditions. Comprehension of classical logic is often presupposed in discussions of non-classical logical systems, e.g., as a meta-theory. Indeed, it is stipulated in **Logical Disagreement** that the deduction found on the blackboard is written in a language that both logicians fully comprehend. We do not need more than noticing and appreciating this very stipulation in order to slide off the objection. Further, we can strengthen this reply by noticing that it is not the case that when there is logical disagreement, one party has automatically misunderstood (or lacks) some concept. The disagreement may just be the result of one side having false

¹⁵ Quote from Luitzen Egbertus Jan Brouwer, “Intuitionism and Formalism,” *Bulletin of the American Mathematical Society* 20 (1913): 91.

beliefs. So, in **Logical Disagreement**, it need not be the case that the intuitionist (supposing that she got it wrong) lacks some concept about how negation works, or has misunderstood or changed its meaning. Negation means whatever it means, also in the intuitionist's mouth, she just has false beliefs about that meaning.¹⁶

4.3 Logical Monism

Now, let us turn to the more challenging objections.

Objection 3. The evidence does in fact justify exactly one doxastic attitude in **Logical Disagreement**, it is just that we do not know which attitude it is. For we do not know which logic is the "correct" model of logical consequence, but surely there is only one correct logic in the end. Thus, UT survives the case even though the underlying logical disagreement leaves us in the dark with respect to what doxastic attitude is justified.

This objection begs the question against logical pluralists (something like Beall & Restall-style pluralists), i.e., the view that there is more than one true logic; there is not always a single answer to the question whether a proposition *P* logically follows from a set of propositions (premises), in some cases there are more than one correct answer. A rough motivation for this kind of pluralism is that classical logic(s), relevance logic(s), intuitionistic logic(s) etc., all have a rightful place in formalizing and restraining logical inference as various important aspects of our pre-theoretic notion of logical consequence can be explicated by each of these approaches to logic. Clearly, begging the question against the pluralist in this way merely relocates the tension from an infight between UT-supporters and -deniers to a clash between logical monism and pluralism, so it seems like a dissatisfying option. Of course, some UT-supporters might be happy to say that logical pluralism is false, and thus they will have a way to save their principle, but this strategy should be supported by strong independent reasons. It will not be enough for the UT-supporter to accept logical monism because it seems like the default position amongst epistemologists. Hence, **Objection 3** is underdeveloped as it stands, and UT-supporters opting for this way out have further work to do. Developing the back and forth between logical monists and pluralists any further here would take us beyond the scope of this paper, but find a few useful references in the footnote below.¹⁷

¹⁶ A similar point is made by Williamson; see Timothy Williamson, *The Philosophy of Philosophy* (Oxford University Press, 2007), chapter 4.

¹⁷ For more on logical pluralism in the Beall & Restall-style, see e.g., JC Beall and Greg Restall, *Logical Pluralism* (Oxford University Press, 2005), JC Beall and Greg Restall, "Logical Pluralism," *Australasian Journal of Philosophy* 78, 4 (2000): 475–493. Other kinds of logical pluralism can be

4.4 Splitting the Evidence

Objection 4. As S1 and S2 belong to two opposing traditions in logic and thus do not accept the same rules of inference, it is actually not the case that they possess the same evidence in the situation described. Surely, considered as a set of propositions, the evidence is the same for both subjects, but due to the subject's diverse logical backgrounds the evidence splits in two. The case really presents both E and E*, where the acceptable inference rules of classical logic are tacitly accepted to induce E and the rules of intuitionist logic are tacitly accepted to induce E*. No purely formal body of evidence (or set of propositions) supports anything pre-theoretically. Pre-inquiry acceptance of a logical system (or another kind of systematic method) is necessary to even generate evidence. Pre-theoretically, the question of which doxastic attitude is supported by a formal body of evidence is empty. Hence, **Logical Disagreement** is not a counterexample to UT since each body of evidence only justifies one doxastic attitude.

Prima facie, this objection seems to have something going for it. Indeed, it might save UT seen as a general epistemic principle since at most one doxastic attitude can be justified per body of evidence. However, at the same time it undermines the initial appeal of UT. For if we need a prior systematic method in order to even generate formal evidence, we get a kind of evidential relativism. To illustrate, take an arbitrary set of purely formal propositions. This set does not constitute a unique body of evidence, as would be natural to suppose, instead it constitutes as many different bodies of evidence as there are acceptable systematic methods of inquiry. This moves our discussion away from evidence to a discussion of acceptable methods, but this discussion should not be relevant to UT. UT should not be true only relative to preferred methodology. For let us remind ourselves of how strong a thesis UT really is: it concerns all bodies of evidence, no matter what subject possesses it and no matter the time and circumstances. The crucial point is that UT is supposed to motivate a certain response to peer disagreement, i.e., at most one peer can be justified in such disagreements. But if formal evidence is relativized to method, the scope of UT is reduced drastically. You can now only share formal evidence with those from your own methodological equivalence class, and there can be as many of those classes as there are acceptable methods. This kind of relativism is clearly not desirable for a UT-proponent, and thus saving UT using

found in: Steward Shapiro, *Vagueness in Context* (Oxford University Press, 2006), Rudolf Carnap, *The Logical Syntax of Language* (Open Court, 1937/2002). For an overview, see Gillian Russell, "Logical Pluralism," *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), ed. Edward N. Zalta, URL = <https://plato.stanford.edu/archives/sum2019/entries/logical-pluralism/>.

Objection 4 turns out to be a Pyrrhic victory.¹⁸ However, some might hesitate to admit that **Objection 4** leads to evidential relativism regarding formal evidence, for it may be objected that E and E* do not have the same epistemic status. There could be good and purely epistemic reasons for favoring E over E* (or *vice versa*) the reply goes. As noted above, E is the body of evidence induced by the tacit acceptance of classical logic, while E* is the result of tacitly accepting intuitionist logic, but surely logicians do not just accept any old system of logic, they have epistemic reasons for accepting whatever system they favor. Thus, S₁'s *total* evidence pool may very well include evidence for accepting DNE, law of the excluded middle etc., which the intuitionist lacks. Similarly, S₂'s *total* evidence pool may well include evidence for denying DNE, law of the excluded middle etc., which the classical logician does not have in her possession. Further, S₁'s reasons may be better than S₂'s ditto (or *vice versa*).

Although this worry is legitimate, it will not save UT. First, it is underspecified in the literature whether UT is meant to apply to the *total* bodies of evidence in this sense, i.e., including pieces of evidence supporting one's methods used to generate evidence. There are hints about the importance of evidence for evidence-generating methods in the literature on *deep disagreement*,¹⁹ but usually such evidence is taken as background information, and thus not as included in whatever body of evidence is under consideration in standard (deep) disagreement cases. Thus, it is not clear what UT-proponents would say about cases involving such *total* bodies of evidence. Further, one could easily rewrite **Logical Disagreement** stipulating that the two logicians were (known) epistemic peers.

¹⁸ Other epistemologists have suggested that one way in which uniqueness might fail is if there is a plurality of methods (in a broad sense) which one could reasonably use to generate evidence. Accordingly, the counterexample **Logical Disagreement** presented here and my discussion about formal evidence being relativized to acceptable methods might reasonably be subsumed under a broader style of argument against uniqueness, namely that UT fails because evidence (of various types) is relative to acceptable methods. For further discussion of this general style of argument see Greta Turnbull, "Why dinosaur paleobiology shows us that reasonable disagreement is possible," unpublished manuscript, Steven Hales, "Motivations for Relativism as a Solution to Disagreements," *Philosophy* 89, 1 (2014): 63-82, Alvin Goldman, "Epistemic Relativism and Reasonable Disagreement," in *Disagreement*, eds. Feldman and Warfield, 187-215.

¹⁹ For discussions of deep disagreement, see Klemens Kappel, "Higher Order Evidence and Deep Disagreement," *Topoi* (2018): 1-12, Michael Lynch, "After the Spade Turns: Disagreement, First Principles and Epistemic Contractarianism," *International Journal for the Study of Skepticism* (6) (2016): 248-259, Klemens Kappel, "The Problem of Deep Disagreement," *Discipline Filosofiche* 22 (2) (2012): 7-25, Michael Lynch, "Epistemic Circularity and Epistemic Incommensurability," in *Social Epistemology*, eds. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford University Press 2010), 262-277.

Then, insofar as evidential symmetry is necessary for peerhood, this would exclude any evidence from the case besides the common evidence. Of course, one could then say that if S_1 is a classical logician and S_2 an intuitionist, they cannot be epistemic peers, but in that case, we are back to square one; formal evidence becomes relativized to your own methodological equivalence class and relativism looms.

4.5 Individualistic Versus Social Epistemology

Objection 5. UT is most plausibly defended as an intra-personal thesis, but **Logical Disagreement** is an inter-personal case. Thomas Kelly distinguishes between intra-personal and inter-personal versions of UT.

UT_{Intra}: Given that my evidence is E, there is some doxastic attitude D that is the only fully rational doxastic attitude for me to take towards proposition p[...].²⁰

UT_{Inter}: Given evidence E, there is some doxastic attitude D that is the only fully rational doxastic attitude for anyone to take towards proposition p[...].²¹

UT_{Intra} holds as a general epistemic principle.

This objection saves UT as a general epistemic principle in the intra-personal domain, but as should be clear, it completely undermines the core motivation for the thesis, which is social. Instead of relativizing to methods as in **Objection 4**, E is now relativized to subjects, and an even worse kind of relativism is unavoidable.

I agree that **UT_{Intra}** is true. Take a perceptual case. If S clearly sees that there is a computer in front of her on the table and this visual perception constitutes her evidence, then under normal circumstances there will be at most one justified doxastic attitude for her to adopt towards the proposition *<there is a computer on the table>*, i.e., S is justified in believing the proposition to be true and nothing besides this. Likewise, **UT_{Intra}** is true in logical cases in so far as we assume that the subject in play has accepted a certain logical system prior to inquiry. This blocks cases where **Logical Disagreement** is reformulated as a single person-case with an eclectic logician who is neither dogmatic regarding the classical nor the intuitionist tradition in logic, but is fully competent in both traditions anyway. Given our assumption, this logician cannot be intra-personally justified in more than one doxastic attitude towards P, e.g., the eclectic logician cannot be justified in a belief that P as well as a suspension of judgement with

²⁰ Quote from Kelly, "Evidence Can Be Permissive," 307. Note that even though Kelly uses the term 'rational' instead of 'justified' in this quote, it will not make any substantial difference for our purposes.

²¹ See footnote 20.

respect to P based on the same body of evidence. However, as mentioned above, admitting that only **UT_{Intra}** is true comes with an unbearable cost for the UT-proponent. For with the embrace of this view, UT is no longer relevant to the peer disagreement debate which it was supposed to be central to. As **UT_{Intra}** is compatible with multiple doxastic attitudes being justified in cases of peer disagreement, the initial motivation behind UT is now completely lost. Thus, UT-proponents should not accept **Objection 5** as it indirectly undermines UT.

5. Concluding Remarks

This paper has introduced a new counterexample to UT which involves logical disagreement. To legitimize this example and strengthen the case for it, I have shown that five different objections trying to save UT from **Logical Disagreement** fails. Two of the five objections were simply fended off, one needed further development to pose any real threat, while explaining away the counterexample with either one of the remaining two options resulted in an unbearable indirect defeat of the thesis. Hence, in the absence of successful objections to **Logical Disagreement**, I recommend that we hesitate in accepting UT as a general epistemic principle.²²

²² Thanks to Francesco Berto, Jessica Brown and Klemens Kappel for helpful comments on earlier versions of the paper.

THE CONFLICT OF RIGIDITY AND PRECISION IN DESIGNATION

Daniele BERTINI

ABSTRACT. My paper provides reasons in support of the view that vague identity claims originate from a conflict between rigidity and precision in designation. To put this strictly, let x be the referent of the referential terms P and Q . Then, that the proposition “that any x being both a P and a Q ” is vague involves that the semantic intuitions at work in P and Q reveal a conflict between P and Q being simultaneously rigid and precise designators. After having shortly commented on an example of vague identity claim, I make the case for my proposal, by discussing how reference by baptism conflicts with descriptive attitudes towards understanding conceptual contents.

KEYWORDS: vagueness, rigid designators, imprecise designators, identity claims

Vague identity claims are very ordinary linguistic items. Consider the following statement:

1. Florentine Neoplatonists are intellectuals working at the court of the Medici during the Renaissance.

Such proposition is an identity claim because it can be analysed as follows:

2. For any x , x is a Florentine Neoplatonist if and only if x is an intellectual working at the court of the Medici during the Renaissance.

It is plain that someone may doubt that (1) is vague. An historian may claim that at least one individual which worked as intellectual at the court of the Medici during the Renaissance was not a Neoplatonist. This being the case, (2) would result definitely false.¹ As a consequence, the interpretation of (1) as an identity claim would exclude that (1) is vague.

However, such a reply is problematic because the referential expressions involved in (1), namely, *being an intellectual working at the court of the Medici during the Renaissance* (MI) and *being a Florentine Neoplatonist* (FN) are both vague, where vagueness is commonly understood to refer to the existence of fuzzy

¹ I'm not assuming that vague propositions are not epistemic; rather, they are epistemic in a non definitive way, namely, they need a precisification of their meaning in order to have a truth-value. If a proposition is definitely false does not need such a move to acquire a truth-value. As a consequence, it cannot be vague (although, it can contain vague terms as constituent).

boundaries separating groups of objects. (Let P be a referential term. Suppose that x is clearly a P and that y is clearly not a P. P is vague if borderline cases exist between x and y, and there are therefore fuzzy boundaries between being a P and not being a P. The issue at stake is that a vague term needs a stipulation of its meaning: there are indeed no fixed insights for establishing those objects to which it applies and those to which it does not).

For example, both *intellectual* and *Neoplatonism* are open to different construals. A jurist or a financial expert is an intellectual worker, but not necessarily an intellectual in the traditional sense, unless the individual has an interest in intellectual concerns which do not relate to their work. As a consequence, *being an intellectual* is a property which those who serve as intellectual workers may or may not have. Consider a list of all intellectual workers at the court of the Medici during the Renaissance and order them by their depth of interest in intellectual concerns. At one end of the list is an individual with no interest in intellectual concerns, while at the other is somebody interested in nothing but intellectual concerns. All other individuals stand between the two opposites: the smaller the intellectual concerns, the closer to the non-intellectual boundary. Where is the dividing line between being an intellectual and not being an intellectual?

Similar considerations hold for FN. For each theoretician working for the Medici consider whether he or she endorsed a qualifying feature of Neoplatonist philosophy, say P (think of a qualifying feature for x-ism as a feature you should endorse if you intend to be counted among x-ists). Now order them by the strength of their endorsement of P. At one end place a theoretician who did not endorse P (and is thus certainly not a Neoplatonist); at the other end place a theoretician who endorsed it at the maximum extent (and thus certainly is a Neoplatonist). Place all other individuals from lower endorsement of P to higher. Where is the dividing line between Neoplatonists and non-Neoplatonists?

Now, if MI and FN are vague referential expressions, a claim which concerns the identity of MI and FN turns out to be vague too. Actually, the semantic indeterminacy of such terms infects any proposition which stipulates their identity: if I do not have a non vague individuation criteria for being a MI and a FN, the identity between MI and FN is fluctuating over different construals of their meaning. These lines of reasoning lead then to the conclusion that (1) is a vague identity claim and highlight two points of interest. First, the informational content of (1) is indisputable: the identity claim between MI and FN provides facts about them, because their identities, taken together, convey a representation of reality that can turn out to be true or false. For example, Renaissance scholars debate over

claims like this one to capture relevant facts about the cultural policies of the Medici dynasty, the history of patronage of the arts, the relationship between political power and exhibition of wealthiness, and so on. Second, (1) is vague because MI and FN are imprecise designators.² By consensus view, a designator is precise if and only if there is something determinately denoted thereby and so it is not vague what the designator picks out; on the contrary, a designator is imprecise if it is not precise.³

My view is that vague identity claims originate from a conflict of rigidity and precision in designation. Conflict here means something along the following lines. In ordinary predication, if a designator is rigid (e.g., a qualified name), the designator picks out precisely a set of objects. Consider a referential expression as *red car*. According to the ordinary use of the term (which implies a non vague construal of the referential expression *red car*), when you sign an agreement for buying a red car, you precisely know what you are buying. Or, if your friend standing at the window says to you: hey, *there is a wonderful red car out there!* you precisely know which kind of object you could see if you walked out there. There are different cars and different degrees of red, but the designator *red car* individuates precisely a set of objects. On the contrary, if a designator is not precise (e.g., an ambiguous term), the designator cannot individuate rigidly a set of objects. Consider the ambiguous term *religion*. Whoever has dealt with the difficult task to define what a religion is, perfectly knows how resistant is the term to a strict definition. Actually, some scholars use *religion* inclusively, others do not. As a consequence, some count as religions what others refute to classify as such. The moral of the story is that, grossly speaking, rigidity and precision in designation stand side by side. If you are a friend of degree approaches in philosophy, you can say that the more a term is rigidly employed, the more the term is precise in designation.

Such a conclusion is not true for vague identity claims. In a substantive sense, vagueness consists in that the proportionality of rigidity and precision is broken: the incapability to access commonly agreed methods for establishing what counts as something has for consequence that rigidity and imprecision in designation are simultaneous features of one and the same proposition. This characterization means that a proposition is a vague identity claim if the

² Garrett Evans, "Can There Be Vague Object?" *Analysis* 38, 4 (1978): 208; Richmond H. Thomason, "Identity and Vagueness," *Philosophical Studies* 42, 3 (1982): 329-332; David Lewis, "Vague Identity: Evans Misunderstood," *Analysis* 48, 3 (1988): 128-130.

³ Dominic Hyde, *Vagueness, Logic, and Ontology* (London & New York: Routledge, 2016), 116-17.

designating behaviour of the involved referential terms fluctuates between rigidity and imprecision.

What I mean by the notion of fluctuation can be spelt out as follows. Let P be a vague referential term. The referentiality of P determines that P is used in a rigid way. That is to say, speakers of a language wherein P occurs understand *prima facie* P as if it individuates a set of objects. Nonetheless, since P is vague, P cannot pick out precisely a set of objects. As a consequence, P is used to refer imprecisely to a set with fuzzy boundaries. Now, while in ordinary predication the rigidity and precision of a referential expression are proportional, in vague predication they are conflicting. Saying that relevantly vague terms in vague identity claims fluctuate between rigidity and imprecision in designation intends to capture that once a term is employed, such a term is employed rigidly, although since it is vague, it cannot be used precisely.

To put this more precisely, let x be the referent of the referential term P (and Q). Then, the proposition “that any x being a P (and a Q)” is vague involves that the semantic intuitions at work in P (and Q) reveal the predicative impossibility to establish a proportionality relation between referential rigidity and precision of P (and Q).

I will set forth a case as evidence for my thesis. Michelangelo is possibly the most important intellectual among Florentine Neoplatonists. Suppose dividing Michelangelo’s lifeline into different segments. Let M_x stands for “Michelangelo at the age of x”, so that:

M_0 stands for Michelangelo’s lifeline segment at the age of 0;

M_1 stands for Michelangelo’s lifeline segment at the age of 1;

...

...

...

M_x stands for Michelangelo’s lifeline segment at the age of x;

...

...

...

M_n stands for Michelangelo’s lifeline segment at the age of n.

According to Kripke's seminal analysis for proper names, which established the definition of the notion, the name *Michelangelo* is a rigid designator.⁴ This means that the use of the name *Michelangelo* to refer to the individual universally known by that name is rigidly determined. I can conceive counterfactual worlds wherein Michelangelo did not become an artist, was not a Florentine Neoplatonist, or even had never been born. Even so, each of these conceptions refers to the very same individual, Michelangelo.⁵

Now, each segment of Michelangelo's lifeline is represented by $M_0, M_1, \dots, M_x, \dots, M_n$, abbreviated expression that refer to Michelangelo at a certain age. They are therefore referential terms for Michelangelo and can be used in identity statements where the proper name *Michelangelo* occurs:

M_0 is Michelangelo;

M_1 is Michelangelo;

...

...

...

M_x is Michelangelo;

...

...

...

M_n is Michelangelo.

Since any of these statements has for content the individual universally known as Michelangelo, the proper name *Michelangelo* works as a rigid designator in each of them, accordingly to Kripke's definition. So far, so good.

However, each of $M_0, M_1, \dots, M_x, \dots, M_n$ is used imprecisely in referring to Michelangelo, because: (a) any M_x and M_{x+1} are continuous over slight temporal changes, and, as a consequence, it cannot be detected what criteria strictly individuate them; (b) they cannot be substituted one for another in statements about Michelangelo within any predicative context. For example, Michelangelo completed his statue *David* in 1504, when he was 29. So, while it is true that M_i carved *David* (where i is greater than 29), it is not true that M_i carved it (where i is

⁴ Saul A. Kripke, *Naming and Necessity* (Cambridge: Harvard University Press, 1980), 48.

⁵ Kripke, *Naming and Necessity*, 77.

less than 29). As a consequence, if x is less than 29, and y is greater than 29, M_x and M_y are both Michelangelo, but what is true of M_x is not true of M_y .

It seems evident then that, while it is not vague whether Michelangelo carved *David* and each of $M_0, M_1, \dots, M_x, \dots, M_n$ is Michelangelo (according to the intuition that proper names are rigid designators), $M_0, M_1, \dots, M_x, \dots, M_n$ are imprecise designators for the individual known as Michelangelo and constitute a set of vague descriptions for him.

Why do they constitute a set of vague descriptions? After all, each member of the set describes Michelangelo: it seems there are no borderline cases in being Michelangelo. To see why they do, consider what mereological constitution is. In a contribution to the debate on the metaphysics of time,⁶ Ted Sider argues that individuating compounded objects in a given instant of time always raises issues of vagueness because the diachronic composition which determines how an object is numerically distinct from others can be captured in terms which may generate a sorites paradox.⁷ The core problem consists in that the temporal changes of an object seems to be continuous over a range of slightly indiscernible differences.

For example, it is notorious that while in his early years Michelangelo worked mainly as sculptor and painter, in his later years he accepted exclusively jobs in architecture. Naturally, he began to work as architect from his early years, and continued to paint and carve privately in his later years too. As a consequence, although it is true that Michelangelo was mainly a sculptor and a painter in his youth and an architect in his old age, it is not easy to see when the change in his artistic inclinations occurs.

According to the logic of Sider's argument, a tri-dimensionalist reading of Michelangelo's life is committed to the acceptance of ontic vagueness. If being a painter and a sculptor is essential to Michelangelo (as it seems reasonable to assume), and being an architect is essential too, Michelangelo fluctuates from being a painter and a sculptor to being an architect; there are fuzzy boundaries between his early and late years. This means that the object Michelangelo had a fluctuating nature, and it is not determinate for which value of i the proposition *M_i was mainly a sculptor and a painter* is true. On the contrary, a four-dimensionalist reading of Michelangelo's life does not raise such a problem, because if Michelangelo had extended in time, he had temporal parts for which he was a painter and a sculptor, temporal parts for which he was a painter, a sculptor, and to

⁶ Ted Sider, *Four-Dimensionalism: An Ontology of Persistence and Time* (Oxford: Clarendon Press, 2001).

⁷ Achille Varzi, "Change, Temporal Parts, and the Argument from Vagueness," *Dialectica* 59, 4 (2005): 488-89.

a lesser extent an architect, temporal parts for which he was an architect and to a lesser extent a painter and a sculptor, and, finally, temporal parts for which he was an architect.

Now, since Sider holds that ontic vagueness is not an option, in order to block the assumption of this kind of vagueness from the problems related to diachronic composition, it is necessary to endorse a four-dimensionalist theory of time. Varzi and others doubt that such move actually works, since there seems not to be an implicature relation between diachronic composition and four-dimensionalism. Nonetheless, if one inclines to evaluating vagueness as a semantic fact, Sider's argument has to be blocked somewhere. Achille Varzi provides an analytical overview of what costs rejecting one or the other premise of the argument involves.⁸ However, my view is that one can pursue a strategy which is not set forth by Varzi's conclusions, by denying that Sider's argument should be answered.

The point of the matter is actually that, independently of how a theorist approaches the nature of vagueness (whether it be ontic or semantic), vagueness generates from linguistic uses. That is to say, a purely linguistic story about designation, namely, about how vagueness is structurally related to the use of predicates which work rigidly and imprecisely at once, may suffice to provide an account for how vague statements work. Such a story can be compatible with a number of different theories; it might be the case that vagueness is exclusively a semantic fact, as well as that vagueness generates from a linguistic use because of the fluctuating nature of things out there.

The key to my approach consists in distinguishing between the individual to which a proper name refers and the conceptions thereof. Although any conception (or counterfactual proposition) of M_0 , M_1 , ..., M_x , ..., M_n is a conception of Michelangelo, there is an overwhelming temptation to consider some more relevant to Michelangelo than others. Had Michelangelo not carved the *Pietà*, *David*, or *Moses*, or had he not frescoed the *Sistine Chapel*, nobody would consider the individual now universally known as *Michelangelo* as Michelangelo. What I mean is that the name *Michelangelo* is used not only for an individual (who could have been a different person and is therefore independent of his accomplishments), but also for a conceptual content individuated by reference to the individual: the content is known by accessing relevant descriptions of the man. Which of these descriptions is required to use the name *Michelangelo* in a proper sense, that is, to refer to the greatest Florentine Neoplatonist? Since at least some of M_0 , M_1 , ..., M_x ,

⁸ Varzi, "Change, Temporal Parts, and the Argument from Vagueness," 497-98.

..., M_n do not take part in the conceptual content individuated by the name *Michelangelo*, they are evidently not essential to being Michelangelo.

For example, Michelangelo completed the *Pietà*, his first universally known work, in 1499, when he was 24 years old. Does any M_x where x is less than 24 really belong to the conceptual content individuated by reference to Michelangelo? Suppose that the *Pietà* is not essential to Michelangelo's artistic production and that the *Sistine Chapel* fresco, created between 1508 and 1512, is the only necessary work. This being the case, belonging or not to the conceptual content individuated by referring to Michelangelo is marked by a different M_x .

My conclusion, in line with the interpretive thinking of art historians and critics, is that $M_0, M_1, \dots, M_x, \dots, M_n$ form a set of vague descriptions. All such scholars debate about the same individual, evaluating his life and work and providing interpretations for his development and artistry. Each description, however, gives a very different account of the same object, and the differences are made possible because *Michelangelo* is a rigid designator. This does not remove the differences, however; the conceptual content individuated by referring to Michelangelo requires a precisification. Each book, essay, and discussion about him satisfies that requirement exactly.

Notoriously, van Inwagen argues for the claim that attributing a proper name by baptism dispenses from providing a description of the named thing; and that such a fact gives a reason in support of ontic vagueness in face of the semantic one.⁹ However, baptism is a performative act which requires understanding a wide extent of descriptive conditions (for example, anything which is necessary for individuating the baptised thing). Consequently, the possibility of a baptism without description is deceptive.

The moral of the story is that vague identity claims reveal a conflict between semantic intuitions concerning designation. Once a term is rigidly introduced by baptism for referring to a thing, it is associated with a series of descriptions of that thing. Although the baptism confers rigidity, the descriptions are counterfactually variable. This variability leaves room for different choices as to which of these descriptions is the cutting line between belonging or not to the conceptual content rigidly designated by the relevant term. If this variability admits precision in giving strict definitions, the work of theoreticians pushes vagueness away, and settles the dispute. On the contrary, whenever different intuitions about the precisification of meaning conflict and compete with each other in a manner which cannot find conclusive reasons in support of any of them, the plurality of slightly different

⁹ Peter van Inwagen, "How to Reason About Vague Objects," *Philosophical Topics* 16, 1 (1988): 255-284.

descriptions for the same conceptual content to which a certain name refers generates a vague approach to the relevant thing. This being the case, the logic of vagueness and its linguistic expression are not able to individuate whether vagueness is a semantic or ontic fact. A supplement of ontological reasoning should be necessary here.

HARDER, BETTER, FASTER, STRONGER: EPISTEMIC STANDARDS AND MORAL BELIEFS

Nicole DULAR

ABSTRACT: Much work in moral epistemology is devoted to explaining apparent asymmetries between moral and non-moral epistemology. These asymmetries include testimony, expertise, and disagreement. Surprisingly, these asymmetries have been addressed in isolation from each other, and the explanations offered have been piecemeal, rather than holistic. In this paper, I provide the only unified account on offer of these asymmetries. According to this unified account, moral beliefs typically have a higher epistemic standard than non-moral beliefs. This means, roughly, that it is typically more difficult for agents to receive the relevant positive epistemic credit (e.g. knowledge) for moral beliefs than for non-moral beliefs. After presenting this account, I consider two alternative unified accounts. According to the first alternative, moral matters are more cognitively demanding; according to the second, moral beliefs have more defeaters. I argue that neither of these alternative accounts succeed, and that my higher standards account is the best unified explanation.

KEYWORDS: epistemic standards, moral testimony, moral expertise,
moral disagreement

Introduction

A quick survey of recent literature in moral epistemology will tell you that many think that moral beliefs are epistemically special. More particularly, one will find many papers dedicated to discussing noteworthy asymmetries between certain areas in our moral and non-moral epistemology, like testimony, expertise, and disagreement. These differences are often viewed as obstacles or hurdles moral beliefs face on their way to moral knowledge that non-moral beliefs don't face. For example, while non-moral knowledge is thought to be easily achieved via testimony, moral testimony is thought to be epistemically problematic, morally problematic, or both.¹ In the same vein, while non-moral expertise is obvious,

¹ Roger Crisp, "Moral Testimony Pessimism: A Defense," *Proceedings of the Aristotelian Society Supplementary Volume* 88, 1 (2014): 129-143; Nicole Dular, "Moral Testimony under Oppression," *Journal of Social Philosophy* 48, 2 (2017): 212-236; Allison Hills, "Moral Testimony and Moral Epistemology," *Ethics* 120, 1 (2009): 94-127; Robert Hopkins, "What is Wrong with

moral expertise is highly controversial and doubtful at best, and moral disagreement threatens skepticism in a way that non-moral disagreement fails to.

What explains these puzzles? Perhaps different things explain each: the proper explanation of the puzzle concerning testimony will in turn differ from the proper explanation for the puzzle concerning expertise which will differ from the proper explanation of the puzzle concerning disagreement. In fact, those who have sought to explain these puzzles in moral epistemology have done just that, seeking to explain them individually rather than collectively.² I am not interested here in these piecemeal accounts. Rather, I am interested in the possibility of giving a *unified* explanation of all of these puzzles. As I'll argue, we can give such a unified explanation. The unified explanation I articulate here is an elegant, simple explanation that utilizes a familiar epistemic mechanism. Given that, all things considered, a unified account ought to be preferred, and provided that this account can adequately explain the puzzles and explain them better than alternative unified accounts, we have reason to prefer the account I give.

Moral Testimony?," *Philosophy and Phenomenological Research* 74, 3 (2007): 611-634; Robert J. Howell, "Google Morals, Virtue, and the Asymmetry of Deference," *Nous* 48, 3 (2014): 389-415; Sarah McGrath, "The Puzzle of Pure Moral Deference," *Philosophical Perspectives* 23, 1 (2009): 321-344; Andreas L. Mogensen, "Moral Testimony Pessimism and the Uncertain Value of Authenticity," *Philosophy and Phenomenological Research* 92, 1 (2015): 1-24; Philip Nickel, "Moral Testimony and its Authority," *Ethical Theory and Moral Practice* 4, 3 (2001): 253-266.

² For accounts which deal only in moral testimony, see Crisp, "Moral Testimony Pessimism," Hills, "Moral Testimony," Hopkins, "What is Wrong," Howell, "Google Morals," McGrath, "Pure Moral Deference," Mogensen, "Moral Testimony Pessimism," and Nickel, "Moral Testimony"; for accounts which deal only in moral expertise, see Sarah McGrath, "Skepticism about Moral Expertise as a Puzzle for Moral Realism," *Journal of Philosophy* 108, 3 (2011): 111-137 and Gilbert Ryle, "On Forgetting the Difference between Right and Wrong," in *Essays in Moral Philosophy*, ed. A. Melden (Seattle: University of Washington Press, 1958), 147-159; for accounts which deal only in moral disagreement, see William Tolhurst, "The Argument from Moral Disagreement," *Ethics* 97, 3 (1987): 610-621. Although no accounts exist which seek to explain all three puzzles together, some accounts consider two of the puzzles in tandem, looking to the bearing one puzzle may have on explaining the other (but not giving an account of what explains them both): for example, Ben Cross, "Moral Philosophy, Moral Expertise, and the Argument from Disagreement," *Bioethics* 30, 3 (2016): 188-194 argues that the puzzle of moral disagreement undermines the possibility of moral expertise; Julia Driver, "Autonomy and the Asymmetry Problem for Moral Expertise," *Philosophical Studies* 128, 3 (2006): 619-644 considers the puzzle of our resistance to accepting the testimony of supposed moral experts, and Sarah McGrath, "Moral Disagreement and Moral Expertise," in *Oxford Studies in Metaethics Vol. 4*, ed. Russ Shafer-Landau (Oxford: Oxford University Press, 2008), 87-108 looks to moral disagreement within the context of there being no moral experts.

This paper will proceed as follows. First, I will look more closely at these longstanding puzzles of testimony, expertise, and disagreement, and the existing piecemeal explanations on offer. Then, I will provide my unified explanation, the Higher Standards account, which holds that moral beliefs typically have a higher epistemic standard than non-moral beliefs. After providing my unified account and showing how it explains the puzzles, I consider two competing unified accounts and argue that both are unacceptable. Finally, I consider and respond to two objections to my own account.

1. The Oddity of Moral Epistemology

Here, I'll explain briefly why each of the three puzzles noted above has been thought to be especially *puzzling*. In the next section, I'll explain how to deal with these puzzles in a unified way.

One area of moral epistemology that has recently received a great deal of attention is moral testimony, and for good reason: our judgments regarding moral and non-moral testimony exhibit a striking asymmetry. While we think it's perfectly acceptable to form non-moral beliefs solely on the basis of others' reports, we balk at instances of forming moral beliefs solely on another person's say-so. Consider:

Eleanor has always enjoyed eating meat but has recently realized that it raises some moral issues. Rather than thinking further about these, however, she talks to a friend, who tells her that eating meat is wrong. Eleanor knows that her friend is normally trustworthy and reliable, so she believes her and accepts that eating meat is wrong.³

Danielle hears about an upcoming demonstration protesting Israel's war in Gaza. Although she knows the causes of the war and knows that civilians are dying from IDF bombing, Danielle is unsure whether the war is just. She doesn't try to think through the matter for herself. Instead, she asks a reliable and trustworthy friend, who says the war is immoral. Danielle accepts her friend's claim and joins the protest. Asked by a journalist why she is demonstrating, Danielle says she knows the war is wrong because her friend told her so.⁴

Here, many object to Eleanor's and Danielle's reliance on their friends in forming their moral beliefs: there is something *prima facie* wrong about Eleanor and Danielle forming their moral beliefs solely on the basis of their friends' say-so. Importantly, these judgments don't seem to be confined to the specific moral subject matter (e.g. eating meat) or sporadic; as Sarah McGrath notes, "the attitude

³ Hills, "Moral Testimony," 91.

⁴ Mogensen, "Moral Testimony Pessimism," 1.

that pure moral deference is more problematic than non-moral deference is widespread, even if not universal, in our culture.”⁵

Moral testimony isn’t the only area in moral epistemology that presents unique epistemic challenges; consider *expertise*. While it’s obviously true that there are experts on all kinds of non-moral subjects, moral experts are thought to be at best few and far between, and at worst entirely non-existent.⁶ Moreover, while it’s usually clear what’s required for non-moral expertise, there’s confusion and disagreement over what is even required for moral expertise. To put it most pessimistically: if, contrary to appearances, there even are any moral experts, we will be seriously hard pressed to find them.⁷

And, if moral testimony and expertise weren’t enough, moral disagreement poses its own unique challenges. Unlike disagreement in non-moral domains, moral disagreement is thought to be especially intractable, as it persists even when both parties appear to share the same (non-moral) evidence. Because of its intractability and persistence, the mere fact of moral disagreement appears to lead directly to moral skepticism. For example, Tolhurst argues that it makes our moral beliefs never justified,⁸ while McGrath and Vavova both argue that disagreement leads to skepticism about a certain subset of our moral beliefs.⁹ Note that no such route to non-moral skepticism (about the existence of global warming, say) is generally thought to be available. Worse, moral disagreement seems to be more widespread than non-moral disagreement.

This way in which moral disagreement appears to lead to moral skepticism will be my focus here regarding the epistemic asymmetry of moral and non-moral disagreement. Even so, there two closely related questions regarding moral disagreement that I’m not interested in pursuing here. I’ll mention them only to set them aside for the remainder of the paper. First, the question of (a) why moral disagreement is so widespread and intractable, and, second, the question of (b) whether we should be “steadfast” and retain our moral beliefs when faced with such disagreement. I set these related issues aside and focus on the question of how moral disagreement can lead to moral skepticism for present purposes because unlike the issue of skepticism, (a) and (b) do not directly concern notable *epistemic*

⁵ McGrath, “Skepticism about Moral Expertise,” 323.

⁶ McGrath, “Skepticism about Moral Expertise,” 323; McGrath, “Moral Disagreement;” Ryle, “On Forgetting.”

⁷ Michael Cholbi, “Moral Expertise and the Credentials Problem,” *Ethical Theory and Moral Practice* 10, 4 (2007): 323-334.

⁸ Tolhurst, “Moral Disagreement.”

⁹ McGrath, “Moral Disagreement;” Katia Vavova, “Moral Disagreement and Moral Skepticism,” *Philosophical Perspectives* 28, 1 (2014): 302-333.

asymmetries in *moral* epistemology. I take (a) to be a metaphysical metaethical question, as the widespread and persistent nature of moral disagreement typically requires metaphysical explanations, such as that either moral relativism or expressivism is true.¹⁰ Although (b) is an epistemic question, I take it to be a question about the correct response to peer disagreement in general, not a question about moral epistemology in particular. In other words, it's unlikely that the correct response to peer disagreement about morality differs the correct response to peer disagreement about non-moral matters. In any case, I won't pursue either of these questions here.

Now, while moral epistemologists have offered explanations of these three asymmetries between moral and non-moral epistemology, what is striking is that all extant approaches have been piecemeal in nature: such accounts aim to explain only why moral testimony is especially problematic, or why moral expertise is especially difficult, or why moral disagreement is especially bad news for moral knowledge. For example, proposals to explain moral testimony appeal to problems it creates for moral agency,¹¹ or moral understanding (the true "aim" of moral beliefs),¹² or that we can't identify reliable testifiers.¹³ Likewise, explanations of the puzzle of moral expertise have pointed to difficulties in identifying experts¹⁴ or to the widespread presence of disagreement as undermining the possibility of moral experts.¹⁵ Lastly, accounts of moral disagreement have claimed that the explanation of why moral disagreement leads to skepticism is that we should all be conciliationists about disagreement in general.¹⁶ But when each of the issues of moral testimony, moral expertise, and moral disagreement are taken together as a whole, the phenomenon to be explained changes its shape and becomes quite striking: it seems that there's not one special problem with moral testimony, one special problem with moral expertise, and one special problem with moral

¹⁰ For examples of these types of arguments, see: Gilbert Harman, "Moral Relativism," in *Moral Relativism and Moral Objectivity*, ed. Gilbert Harman and Judith Jarvis Thomson (Cambridge, MA: Blackwell Publishers, 1996), 1-64; Jesse Prinz, *The Emotional Construction of Morals* (New York: Oxford University Press, 2007); David Wong, *Natural Moralities: A Defense of Pluralistic Relativism* (New York: Oxford University Press, 2006).

¹¹ Crisp, "Moral Testimony Pessimism;" Hills, "Moral Testimony;" Hopkins, "What is Wrong;" Howell, "Google Morals;" Mogensen, "Moral Testimony Pessimism;" Nickel, "Moral Testimony."

¹² Hills, "Moral Testimony."

¹³ McGrath, "Pure Moral Deference."

¹⁴ Cholbi, "Moral Expertise;" Driver, "Moral Expertise."

¹⁵ Cross, "Moral Philosophy."

¹⁶ Vavova, "Moral Disagreement."

disagreement. Instead, it seems there's some special problem with moral epistemology *as a whole*.

Of course, some think that our judgments concerning the asymmetry of moral testimony, expertise, and disagreement with their non-moral counterparts are illusory, preferring instead to offer debunking explanations of these judgements.¹⁷ My purpose in this paper is not to take issue with the asymmetry judgments themselves. Rather, I'll simply assume things are as they appear to be. Supposing that there are these puzzling differences, we are faced with two options: either go piecemeal, and explain each puzzle independently, or go wholesale, and offer a unified account that explains them all together. Again, what's notable is that all approaches to these puzzling asymmetries between moral and non-moral beliefs (including the debunking ones) have taken the first option, offering *disunified, piecemeal* explanations.¹⁸ What hasn't been attempted, though, is taking the second option and going wholesale in our explanation. My aim in this paper is to do just that, taking the second, unexplored option, and providing a *unified* account.

In the next section, I will lay out my unified account. Importantly, my account has advantages over the piecemeal accounts currently on offer. Beyond the fact that, all things considered, unified explanations ought to be preferred to disunified ones, my account avoids positing any exceptional features of moral beliefs that some other piecemeal accounts have relied on, like the idea that moral beliefs have a distinct "aim" that non-moral beliefs don't. Rather, my account relies on a familiar epistemic mechanism that is commonplace and widely discussed: epistemic standards and how they shift. According to my account, moral beliefs typically have a higher epistemic standard than non-moral beliefs. This means,

¹⁷ Driver, "Moral Expertise;" Jason Decker and Daniel Groll, "Moral Testimony: One of These Things is Just Like the Other," *Analytic Philosophy* 54, 4 (2014): 54-74; Jason Decker and Daniel Groll, "The (In)significance of Moral Disagreement for Moral Knowledge," in *Oxford Studies in Metaethics, Volume 8*, ed. Russ Shafer-Landau (New York: Oxford University Press, 2013), 140-167; Karen Jones "Second-hand Moral Knowledge," *Journal of Philosophy* 96, 2 (1999): 55-78; Karen Jones and Francois Schroeter, "Moral Expertise," *Analyse und Kritik* 34, 2 (2012): 217-230; Andrew Reisner and Joseph Van Weelden, "Moral Reasons for Moral Beliefs: A Puzzle for Moral Testimony Pessimism," *Logos and Episteme* 4 (2015): 429-448; Peter Singer, "Moral Experts," *Analysis* 32, 4 (1972): 115-117; Paulina Sliwa, "In Defense of Moral Testimony," *Philosophical Studies* 158, 2 (2012): 175-195.

¹⁸ To be clear: while some have considered two of these puzzles together (e.g., Kieran Setiya, *Knowing Right from Wrong* (Oxford: Oxford University Press, 2012), addresses both disagreement and testimony, and Cross, "Moral Philosophy" considers how the presence of disagreement bears on expertise), there exists no account that explains all three in a wholly unified manner.

roughly, that the standard agents must meet in order to receive the relevant positive epistemic credit (e.g., knowledge or justification) is typically more stringent for moral beliefs than the corresponding standard is for non-moral beliefs. To be clear, I won't be arguing for a universal claim: that every single moral belief will have a higher epistemic standard compared to any other non-moral belief. Such a universal claim is too strong to be plausible. Rather, my claim will be that this is *typically* the case, and as such it is a characteristic and noteworthy feature of moral epistemology as such. Importantly, one need not endorse such a universal claim to adequately explain the asymmetries between particular areas of moral and non-moral epistemology, since, as we've seen, these concern *general* issues with particular aspects of moral epistemology. For example, the testimony-involving asymmetry is not that for every single possible instance of non-moral testimony, any possible instance of moral testimony will be more problematic than any possible instance of non-moral testimony. That would be quite implausible; rather, it is that moral testimony *in general* is (more) problematic.¹⁹ In order to assess this account, we should first turn to the concept of an epistemic standard.

2. The Higher Standards Account

2.1. Epistemic Standards

In very basic terms, we can think of an epistemic standard as marking how good of an epistemic position an agent needs to be in to count as *knowing* or as *having a justified belief*. The idea of an epistemic standard captures the intuitive thought that in order to determine whether an agent's belief is justified or counts as knowledge, we need to know not just how much evidence they *have*, but how much they *need*.

This concept of an epistemic standard allows us to capture the thought that in some areas of inquiry, or in some contexts, what's required for knowledge or justification can *change*: it's not that knowledge of every kind of fact requires the

¹⁹ The same can be said for the other aspects of moral epistemology that have received widespread attention, namely expertise, and the effect disagreement has in undermining knowledge or leading to skepticism. Expertise by definition concerns a general ability, or knowledge of a range of facts about a particular topic, not perfect ability or knowledge of every single fact about a particular topic. Likewise, the phenomenon regarding moral disagreement concerns how it in general leads to skepticism, not how every single instance of moral disagreement undermines the status of knowledge for every single moral belief every single person has. I further explain how my account of there typically being a higher epistemic standard for moral beliefs explains puzzling asymmetries in moral epistemology in section 2.2.

same strength of evidence. This is just to say that sometimes at least, the epistemic standards shift.²⁰

This shiftiness of epistemic standards has been utilized by contextualists in epistemology to explain otherwise surprising patterns in our knowledge attributions. For it seems that, while we may want to deny large-scale skepticism wherein agents always know little to nothing at all, we may also want to allow for small-scale skepticism, wherein agents fail to know particular propositions in particularly demanding circumstances. For example, while it seems perfectly innocuous to say that I know that I have hands when I am walking to class, once I find myself embedded in a classroom discussion about skepticism it seems correct to deny that I know I have hands. Contextualists explain these shift judgments by appealing to epistemic standards: from the walk to the classroom to the discussion of skepticism within the classroom the epistemic standard has shifted (more specifically it has gotten more strict).²¹ In this case, while my perception of having hands was good enough to make my belief that I have hands knowledge outside of the classroom, this evidence is no longer sufficient to make my belief knowledge once inside the classroom's skeptical walls with its stricter epistemic standard.

That is the intuitive idea. But we can get a bit more specific. We can say that an epistemic standard specifies a range of possibilities that an agent may ignore or fail to rule out while still counting as knowing or having a justified belief.²² These possibilities specify ways the world could be in which not-*p* is true (when one's belief is *p*). Importantly, this means that for any given belief, there is more than one epistemic possibility: we don't divide up the epistemic possible worlds simply into two worlds, *p* and not-*p*, where one of these is the actual world. Rather, epistemic possibilities are individuated by *ways* in which your belief could be false.

²⁰ I use an evidentialist model of standards here for the sake of simplicity. Nothing in my argument hangs on this assumption.

²¹ Stuart Cohen, "Knowledge and Context," *Journal of Philosophy* 83, 10 (1986): 574-583.

²² Strictly speaking, this is actually where contextualists and fallibilists—who also appeal to epistemic standards—part ways in their understanding of what a standard specifies. Fallibilists will say that an agent does not need to rule out every possibility, while contextualists will say that they do; the difference is how each is quantifying over 'every.' For the fallibilist, 'every' really does pick out every single possibility, while for the contextualist 'every' picks out a certain subset of every single possibility, for example every *salient* possibility. This is perhaps why some contextualists hold that contextualism is an infallibilist position (see David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74, 4 (1996): 549-567.), while others hold it to be fallibilist in nature (see Mark Heller, "The Proper Role for Contextualism in Anti-Luck Epistemology," *Nous* 33, 13 (1999): 115-129.). In the end, though, each camp seems to agree on this general statement: out of all the total possibilities, in order to know an agent must be able to rule out only all of those possibilities in a subset of these total possibilities.

For example, there are many possible worlds in which your belief that you have hands is false: you could be hallucinating, you could be dreaming, etc. But only some of the ways the world could be—only some of these possible worlds—are relevant to the epistemic status of your beliefs in the actual world. This is because of some relation they bear to you, and that you bear to them: they are salient, or relevant, etc. Provided you are able to rule out that set of worlds where your belief would be false, your beliefs enjoys the relevant positive epistemic status (e.g. knowledge, justification). Overall, the rigor of an epistemic standard can be specified in one of two ways: sometimes, a more rigorous standard specifies *more* possibilities that one must be able to rule out, while other times it specifies possibilities that are simply *harder* to rule out. My account allows for both of these interpretations of rigor.

Like rigor, the notion of “ruling out” possibilities can be understood in a number of ways. On a probabilistic model, this could mean either that some possibilities are made *more improbable*, or that *more possibilities* are made improbable. My claim is just that for moral beliefs, the epistemic standard shifts, becoming more rigorous and thus requiring more in at least one of these two ways. Importantly, this view of standards is also compatible with both internalist and externalist theories of justification and knowledge. For example, if one were a reliabilist, the upwards shift in the rigor of the standard would require one to have more safety or sensitivity. If one were an evidentialist, one would be required to possess stronger evidence that rules out more possibilities. What’s important for my claim is that what it takes to have an epistemic state (justification, knowledge) depends on the rigor of the standard, and that morality makes this rigor increase.

Additionally, my account is neutral between competing accounts of how standards are fixed.²³ For example, some hold that this range is flexible, picking out different worlds in different contexts, while others hold that the same range of worlds is picked out in all contexts.²⁴ Articulating the causes of the shiftiness of epistemic standards in general, and the shiftiness of standards for moral beliefs in particular, is a large project unto itself. Happily, it’s mostly outside the scope of the current paper. This is because there are two independent questions: *whether* moral beliefs typically have a higher standard and *what* exactly fixes standards. These questions are obviously related, since one’s answer to the latter might determine

²³ To be clear: my account of what an epistemic standard is neutral along these lines; however, invariantism regarding epistemic standards (that is, standards for any and all kinds of beliefs) is incompatible with my argument for the higher standard for moral beliefs.

²⁴ The former being contextualists and subject sensitive invariantists, and the latter being invariantists.

one's answer to the former. But answers to the two can come apart in the sense that many can agree that moral beliefs have a higher standard while completely disagreeing about what fixes the standard. For example, many can agree that skeptical scenarios have a higher epistemic standard than non-skeptical ones while disagreeing over what functions to make the standard stricter in skeptical scenarios (e.g. whether contextualism or subject sensitive invariantism is the best account). However, to preserve the credibility of my claim that moral beliefs typically have a higher standard it is important that there at least be some initially plausible models available, so I will briefly address this issue here.

One possible model of how standards are fixed is the well-known stakes-model, wherein an epistemic standard is determined in part by the practical stakes, or the costs of one's belief turning out to be false.²⁵ Such a standards-fixing model is taken up elsewhere, where it is said that there are certain practical stakes are unique to moral beliefs (for example, the costs of being the target of certain reactive attitudes) such that when we account for these stakes, such a model does a good job of tracking how most moral beliefs have a higher epistemic standard and how the ones that intuitively don't, don't.²⁶ Although articulating further details of this model would take us too far afield here, I hope this gives the intuitive, initially plausible flavor of the model. Of course, if this particular model does not sound appealing, one needn't reject my claim that moral beliefs typically have a higher standard: again, these are distinct claims, and so we can agree that moral beliefs typically have a higher standard while disagreeing over the correct account of what fixes those standards. The claim that moral beliefs have a higher epistemic standard does *not* depend on the success of my—or any—particular standards-fixing model. For example, we could instead adopt a kind of Relevant Alternatives Contextualist view, where the possibilities that one must be able to rule out are those that are presupposed or otherwise entered into the conversational score, coupled with a view that moral beliefs presuppose more or more difficult to rule out possibilities.²⁷ Again, although I lack the space here to adequately address which particular standards-fixing models are the best accounts of the typical higher standard for moral beliefs, such plausible models are available. Given the

²⁵ Jeremy Fantl and Matt McGrath, *Knowledge in an Uncertain World* (Oxford: Oxford University Press, 2009); Robin McKenna, "Interests Contextualism," *Philosophia* 39, 4 (2011): 741-750; Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005).

²⁶ Nicole Dular, "Moral Stakes, Higher Standards," (unpublished manuscript).

²⁷ Michael Blome-Tillman, "Knowledge and Presuppositions," *Mind* 118, 470 (2009): 241-294.

availability of such models and their initial plausibility, the credibility of the claim I make here that moral beliefs have such a higher standard should remain intact.

2.2. A Unifying Explanation

With this conception of epistemic standards and the idea that the epistemic standard is typically stricter for moral than for non-moral beliefs in hand, we can approach our original problem. I'll now briefly explain how my Higher Standards account resolves the three puzzling featured in moral epistemology with which we began.

First, consider moral testimony and the default judgment that it is an illegitimate way to gain moral knowledge. According to my account, in order to have moral knowledge the requirement that an agent rule out possible worlds is relatively stringent: an agent either needs to rule out a significant number of possible worlds or to rule out a set of worlds that is harder to rule out. The reason why agents are unable to gain moral knowledge from testimony is because merely forming one's belief on the basis of another's report does not provide one with the ability to rule out all of the possibilities that one would need to in order to have (moral) knowledge. Although testimony may equip one with true moral beliefs, it does not equip one with the ability to rule out the demanding set of possible worlds that one needs to in order to have moral knowledge.²⁸

²⁸ One may wonder how far my Higher Standards account goes in explaining not just asymmetries in judgments about cases of pure moral and non-moral deference (where speakers do not inform hearers of any of the reasons for the truth of their belief) but also in explaining asymmetries in judgments about cases of impure moral and non-moral deference (where hearers come to adopt not only the speaker's belief, but also their reasons in support of the truth of their belief). The worry is that since my account explains the asymmetry in terms of being in a position to rule out possibilities, in cases of impure moral deference the hearer would be able to rule out all of the same possibilities as the speaker, since they possess the same reasons for the belief; but, the asymmetry remains even in these cases, as we still judge that the hearer lacks justification or knowledge while the speaker does not. However, my Higher Standards view is amenable to preserving this asymmetry of impure testimony: it can do so by adopting a more robust interpretation of what "ruling out" requires. For example, on some contextualist views, ruling out would require more than just possessing evidence that makes certain propositions improbable to a certain degree. Rather, it requires that one is able to engage with others in a certain way, for example by appealing any objections they may have about the truth of your belief. For this more robust understanding of "ruling out", see David Annis, "A Contextualist Theory of Epistemic Justification," *American Philosophical Quarterly* 15, 3 (1978): 213-219, and Carl Wellman, *Challenge and Response: Justification in Ethics* (Southern Illinois University Press, 1971) on the Challenge-Response Model.

Next, consider the apparent lack of moral expertise. According to my account, the standard for moral expertise is stricter than the standard for expertise in other, non-moral domains. This means that the kind of epistemic credentials one would need to have in order to count as an expert are greater for moral expertise. For example, one would need to be able to rule out a comparatively large amount of possibilities for a comparatively large amount of moral beliefs to count as an expert. The reason why moral experts are either scarce or entirely non-existent is because few or perhaps none of us have the ability to do this.

Lastly, my model can explain how disagreement may, after all, lead to skepticism. One way it could do this is by functioning to make relevant new possibilities. For example, it may function to make relevant possibilities like making a mistake in reasoning, or succumbing to a bias. The more widespread a case of disagreement over some moral proposition m , the more possibilities must be ruled out in order to qualify as having knowledge that m . Provided that I cannot rule these out, I fail to secure knowledge. Since standards are understood in terms of possibilities that must be ruled out, moral disagreement leads to skepticism by making more possibilities relevant, and thus by making the epistemic standard more stringent.

Now that we're clear on how my Higher Standards account explains these problematic asymmetries, we should look to see how alternative unified accounts would explain the asymmetries. Again, since in this paper I am seeking an explanation of the apparent oddity of moral epistemology that would vindicate our commonsense judgments about moral testimony, moral expertise, and moral disagreement, I will not be considering debunking explanations of that oddity. As alternative explanations, the accounts to consider are those that posit a mechanism other than the one I appeal to, namely epistemic standards. In the next section, I will consider such rival accounts.

3. Alternative Explanations

3.1. Morality is Hard

One explanation that moral epistemology in general is more problematic than non-moral epistemology is that moral matters are just so exceedingly difficult to figure out. It's just so much more difficult, the thought goes, to determine moral matters such as whether abortion or eating meat is morally permissible than whether the bus runs on Saturdays. It's a very difficult task to do the work that is necessary to adequately settle moral questions: one must consider arguments for and against, checking for falsities, fallacies, counterexamples, and more. Both the kind of

reasoning and time required to consider such questions is large and looming. Morality is hard.

Of course, I agree that morality is hard: this is something that my Higher Standards account explains. In order for this view to be a real competitor, it can't simply amount to the view that moral matters are difficult, since the Higher Standards account may admit this, and then just explain this fact in terms of a more rigorous epistemic standard for morality. Instead, this account must explain what *makes* moral matters epistemically difficult. Moreover, it must do so by appeal to a mechanism other than the one I've identified in order to be a genuine rival.

There are two mechanisms that this rival account might point to. One way of thinking about the "morality is hard" view is that settling moral questions requires a large amount of time; alternatively, one may think that the kind of reasoning required to settle moral questions is exceedingly demanding. Using E to stand for the evidence base that's required to have a justified belief, the view might be either (a) that it is harder to obtain E, i.e. one generally needs to spend more time working in order to obtain E, or (b) that it is harder to draw the correct conclusion on the basis of E, i.e. that the kind of reasoning required to work through one's evidence in order to arrive at a justified belief is of a high level or is quite complex (e.g. it involves the use of difficult mathematical formulas), or both (a) and (b).

Let's take option (a) first. Given this mechanism, one would say that the reason why moral knowledge or justification is harder to obtain is that one needs more time working through or thinking about moral issues in order to successfully arrive at knowledge. More specifically, many agents considering moral questions just haven't obtained E yet (or, more minimally, that they've been able to obtain less of E than the amount of E they're typically able to obtain within the same time for the E that corresponds to various non-moral beliefs).²⁹ Taking option (b) instead, one would say that the reason why moral knowledge or justification is harder to obtain is that moral issues require one to engage in more demanding or complex forms of reasoning in order to successfully arrive at knowledge. More specifically, many agents considering moral questions just haven't successfully used the kind of higher level reasoning required to adequately draw conclusions on the basis of E. Lastly, if one held both (a) and (b), one would say that the reason why moral knowledge or justification is harder to obtain is that moral issues both require greater time and more complex reasoning in order to successfully arrive at a justified belief or knowledge.

²⁹ For example, one could think that one needs normative evidence to justify a normative belief, and it is generally harder to acquire normative evidence (than descriptive evidence).

In general, this unified account could explain the initial asymmetries in the following way. If moral beliefs are hard with respect to (a) and (b), and moral expertise requires one to have a high amount of evidence and evaluate it extremely well when reaching certain moral beliefs, then moral expertise would be hard to come by. Likewise, given (a) and (b) reliable testifiers would be hard to come by. And, lastly, if it is difficult to assess moral claims in the ways (a) and (b) outline, moral disagreement can lead to skepticism by causing one to lose the evidence one may have had or undermining one's ability to work through the now-competing evidence one has.

Are either of these mechanisms a good explanation of the epistemic difficulty of morality? I think that they are not. Remember here that in order for this rival explanation to explain why moral beliefs have certain epistemic puzzles that non-moral beliefs don't, the mechanisms it points to need to be distinctive of moral beliefs. This is because the explanation we are seeking is one that explains how there are certain systematic *differences* between moral and non-moral epistemology. The reason why this rival account fails is simply because the mechanisms it picks out are not distinctive. To see why, consider the following pair of moral and non-moral beliefs:

(NM2): Daria is a college freshman taking an applied ethics course and after one month in the course has just been told that many animals were killed last year for their meat, as well as the fact that many animals (e.g. mice, rabbits, and moles) are killed each year in producing and maintaining crops for food that all vegetarians depend on. Daria considers the question of whether being vegetarian kills more animals than being a meat-eater does. After consulting a few reliable yet neutral sources (e.g. peer-reviewed scientific journals, not PETA) on each side of the debate and crunching the numbers, Daria forms the belief that being vegetarian kills more animals than being a meat-eater.

(M2): Daria is a college freshman taking an applied ethics course and after one month in the course has learnt about arguments both for and against eating meat, considering only arguments for its permissibility and impermissibility (not its obligatoriness), and considering the same quantity (e.g. one each) and quality (e.g. both valid, with plausible premises) of arguments for each side, from a credible yet neutral source (e.g. the Stanford Encyclopedia of Philosophy). Daria considers the question of whether eating meat is morally permissible or morally impermissible. Without consulting anyone else, and after carefully considering the arguments, Daria forms the belief that eating meat is morally permissible.

In these cases, it's clear that the non-moral belief is difficult with respect to (a): Daria would need to spend a lot of time working collecting the relevant data about the statistics of animal deaths in crop cultivation and meat farms. It's also the case that each belief is difficult with respect to (b): Daria would need to engage in

some high-level reasoning such as higher-level math to work through all of the information on statistics he had gathered. And, as this account stipulates, the moral belief is likewise difficult with respect to (a) and (b). Yet, it seems that the moral belief still lacks the same kind of epistemic credit that the non-moral belief has (for example, it appears to be less justified).³⁰ Moreover, upon reflection is it simply not true that morality is the only domain of inquiry that requires a great amount of time or complex reasoning to arrive at knowledge or justified beliefs within that domain: various complex scientific questions also require these. So, even though this account is unified, it does not succeed in accounting for the *asymmetries* of moral and non-moral epistemology.

However, defenders of this alternative account might object. They might insist that the kind of reasoning required for moral beliefs is always going to be more demanding or complex than that required for any other domain of inquiry, as it's of its own special kind, unlike any other type of reasoning used in any other domain. For example, perhaps moral reasoning requires a special kind of sense or faculty that other domains don't, the operation of which is itself extremely complex. But it's terribly *ad hoc* to posit a special kind of reasoning just to save this account. Moreover, this seems to just put a *name* to the problem, rather than offering an *explanation* of it. We started by observing that moral knowledge is hard to come by. It won't do to end simply by observing that the kind of reasoning that leads to moral knowledge is also itself hard to come by. We would still want to know why this is.

We've just seen why this Morality is Hard explanation fails. In the next section, I'll explain why the other competing explanation won't work either.

3.2. Morality's Many Defeaters

Another unified explanation claims that the reason moral beliefs lack the kind of epistemic credit non-moral beliefs enjoy is that moral beliefs typically come with more defeaters than non-moral beliefs do. There are two ways of understanding this defeaters account. On one way of understanding it, the accounts turns out not

³⁰ At this point one may object that we would not have the judgment that the moral belief is less justified here if the non-moral belief were to be some controversial scientific claim. First, notice that the non-moral belief presented is controversial: Daria is confronting conflicting accounts of the number of animals killed. Second, in order for the cases to be analogous, if the controversial scientific claim considered is abstract and general, so must the moral claim, which would force us to consider a new moral case as well (e.g., if we are to consider a controversial scientific theory we would need to consider a controversial moral theory); here, both beliefs are comparative and concrete in nature.

to be a genuine rival to my Higher Standards account. On another understanding, although it is a genuine rival, it results in counterintuitive conclusions, and so ought to be rejected. First, let me briefly explain the relevant notion of defeaters in play.

Defeaters come in roughly two kinds: *rebutting* and *undercutting* defeaters.³¹ On an evidentialist picture, *rebutting defeaters* are those that serve as a reason to believe a proposition that's incompatible with one's conclusion from the evidence (e.g. d is a defeater that warrants not-p (on the basis of E) when one was originally warranted in concluding p on the basis of E), while *undercutting defeaters* serve as reason to believe that E does not actually itself warrant p, without providing reason to believe the negation of p. Given this characterization, one way to understand defeaters is as a kind of higher-order evidence, that is, evidence about the character of one's (first-order) evidence.³² For example, consider your belief that the apple is red that you formed on the basis of your perception of the apple appearing red to you. Your belief would be accompanied by the first type of (rebutting) defeater if you were told that you were given an inverted color spectrum drug: in this case, the fact that you were given such a drug means that you now have, on the basis of your perception, a reason to believe that the apple is green, not red. It is evidence that your original first-order evidence—your perception—actually does not warrant p (that the apple is red), but rather warrants a proposition incompatible with p (that the apple is green). In this case we can say that your total evidence consisting of E+d warrants not-p. Your belief would be accompanied by the second type of (undercutting) defeater if you were told that there's a 50/50 chance that you were given an inverted color spectrum drug: in this case, your original evidence for your belief that the apple is red (your visual perception) would be insufficient evidence for your original belief, such that you ought to abstain from believing what color the apple is. In this case we can say that your total evidence consisting of E+d fails to warrant p.

Now, for the opponent who wants to claim that the grounds of the asymmetries in moral epistemology is that moral beliefs typically have more defeaters than non-moral beliefs, they must not only point to defeaters that accompany moral beliefs, but also point to ones that are *specific* to moral beliefs such that non-moral beliefs either don't also typically have them or don't typically

³¹ John Pollock, *Contemporary Theories of Knowledge* (Savage, MD: Rowman and Littlefield, 1986).

³² David Christensen, "Higher-Order Evidence," *Philosophy and Phenomenological Research* 81, 1 (2010): 185-215; Maria Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," *Philosophy and Phenomenological Research* 88, 2 (2014): 314-345.


have them to the same degree. Otherwise such defeaters would not account for the *difference* in epistemic credit between moral and non-moral beliefs. Given this constraint, there are a few considerations one might cite. One might point to the fact that there is a lot of disagreement surrounding moral claims, much more than what typically surrounds non-moral claims. Likewise, one might argue that there are more counter-arguments to consider with respect to moral claims than non-moral claims. With each of these options, one could claim that one's (first-order) evidence E doesn't yield a justified moral belief or knowledge because any of these considerations would serve as a kind of defeater for E, either in the sense that it makes E insufficient to warrant the belief that p, or that it makes E warrant the belief that not-p: either way, one's total evidence consisting of E+d fails to make one epistemically justified in believing p or knowledge that p. For example, consider a case where I originally believe that eating meat is morally permissible, but then come across another rational person (perhaps even with all the same non-moral evidence that I have) who disagrees with me and who instead believes that eating meat is morally impermissible. One could claim that that's a reason to think that my original evidence E is not sufficient to justify me in believing that eating meat is morally permissible, such that I should abstain from believing it. In this case, the fact of this disagreement undercuts my (first-order) evidence E to believe that eating meat is morally permissible; thus, my total evidence consisting of E+d would fail to make my belief that eating meat is morally permissible epistemically justified. In this way, even if an agent had roughly the same amount of first-order evidence for both her moral and non-moral beliefs, her moral belief would be less justified because there would be more defeaters present, and so more reasons that make it the case that E is not sufficient to warrant her moral belief. The total evidence the agents typically have for moral and non-moral beliefs is not the same.

At this point we need to consider precisely how defeaters function to make one's evidence insufficient to warrant one's belief that p. On one understanding, defeaters (or, more specifically, the fact of disagreement in the moral case) function by raising a specific possibility that my belief is false. For example, maybe eating meat is morally impermissible after all, given that (so many) reasonable others think so; perhaps I made a mistake in my reasoning, or succumbed to bias. On this understanding, while defeaters undermine my (first-order) evidence E for my belief that p such that my total evidence of E+d is no longer sufficient to justify p, they do this by introducing additional ways in which my belief could be false, that is, possibilities. On this account, defeaters just introduce or make relevant certain kinds of possibilities, ones that are not ruled out by one's evidence (given that, if it could be ruled out, it wouldn't render E insufficient to justify p).

For example, consider our previous example involving the belief that the apple is red, where one's evidence consists of the perception of the apple appearing red, and the defeater that's present is the fact that there's a 50/50 chance one was given an inverted color-spectrum drug. On the proposed understanding of what defeaters are, the fact that there's a 50/50 chance that one was given an inverted color-spectrum drug introduces a new possibility that the apple is not red (more specifically, that it's green). However, since one's evidence—namely, one's perception—is not able to rule out this possibility, one's belief fails to be justified or count as knowledge.

At this point, talk of possibilities should sound familiar to the attentive reader. This is because epistemic standards were originally understood as specifying possibilities that must be ruled out in order for a subject's belief to count as justified or knowledge. Remember again that this is just to say that the more rigorous the standard, the greater the set of possibilities. So, if defeaters are just relevant possibilities—specifically, ones that one's evidence is unable to render sufficiently improbable—then one who holds that there are generally more defeaters for moral beliefs than non-moral beliefs is committed to the view that moral beliefs generally have higher epistemic standards.

To further understand how this 'More Defeaters' view is not a rival view to my favored 'Higher Standards' view, consider the following model.



On this model, let the box indicate the set of all epistemic possibilities. Let the 'P' circle indicate the possible worlds in which p is true, and the 'E' circle indicate the

worlds that are compatible with one's evidence; all of the space outside of these circles consists of not-*p* worlds. Using our case, we can understand the 't1' line as indicating the epistemic standard at the time before the defeater was introduced (before you were told that there's a 50/50 chance you were given an inverted color spectrum drug), while the 't2' line indicates the epistemic standard at the time after the defeater was introduced. The epistemic standard at t1 indicates all of the possible worlds one needs to rule out at t1 in order to count as having a justified belief that *p* (namely all of those worlds above the 'standard at t1' line), while the epistemic standard at t2 indicates all of the possible worlds one needs to rule out at t2 in order to count as having a justified belief that *p* (all of the worlds above the 'standard at t2' line). The standard at t1 is pretty low: it indicates, roughly, that one can fail to rule out all of the not-*p* worlds that fall below it while still having a justified belief that *p*. However, at t2 the standard increases, becoming more stringent, thus indicating, roughly, that one can fail to rule out only those not-*p* worlds that fall below it while still having a justified belief that *p*. Importantly, though, while at t1 (pre-defeater) there are no not-*p* worlds that are compatible with your evidence (that is, there are no worlds that are inside the *E* circle but outside the *P* circle), at t2 (post defeater) there are; this means that while your belief meets the epistemic standard at t1, it fails to meet it at t2, such that while you have a justified belief or know that *p* at t1, you have an unjustified belief or fail to know that *p* at t2.

It should be clear, then, that this particular interpretation of the More Defeaters view is not a rival account to my Higher Standards account. Rather than denying that moral beliefs enjoy higher epistemic standards than non-moral beliefs, this More Defeaters view is just articulating a specific way in which the standard is higher, or how it is that the standard is higher for moral beliefs (or, more specifically, what makes a possibility one an agent must be able to rule out). But, again, they are not disagreeing about the fact that the epistemic standard is higher for moral beliefs.

However, there remains an interpretation of the More Defeaters view that is a genuine competing alternative to my Higher Standards view. On this alternative understanding, defeaters (or, more specifically, the fact of disagreement in the moral case) function to make one's evidence insufficient to warrant one's belief that *p* by directly affecting one's evidence. It is not that the standard becomes more rigorous, but just that one falls farther from it given the reduced strength of one's evidence. On this account, the epistemic standards for moral and non-moral beliefs could be exactly the same and remain fixed, but yet moral beliefs are more

epistemically problematic because one's evidence is typically comparatively worse in the moral domain.

Importantly, for this view to capture cases of comparative lack of justification and not just knowledge for moral beliefs, it would have to be the case that the relevant defeaters are recognized or possessed by the agent. This is because although some hold that the simple existence of defeaters—in this case, the simple existence of moral disagreement—is enough to undermine knowledge, it is widely held that in order to affect justification, the agent herself must be confronted with the defeater or made aware of it.³³

The problem with this account is that while it seems correct to say that justification is undermined by defeaters only when agents are cognizant of them for non-moral cases, in the moral case lack of awareness of the defeater leads to counterintuitive results. For example, this understanding of the More Defeaters view would implausibly conclude that in cases where agents just aren't aware of such disagreement concerning a moral issue (for example, because they live in very isolated homogeneous communities, or never bothered to ask anyone else their opinion on the matter), their moral beliefs would not suffer a loss of justification. Likewise, if all that is required to be a moral expert is to have a sufficiently high volume of justified moral beliefs, then one could become a moral expert quite easily. But this is very counterintuitive. So, while this understanding of defeaters can explain some cases, it cannot explain all the puzzles that would need to be explained.

In the end, then, the More Defeaters view either is not a genuine rival to my Higher Standards view, or is rife with counterexamples, and so ought to be rejected.

4. Different but Equal?

Even if the first understanding of the More Defeaters view is not incompatible with my favored Higher Standards view, we might still wonder why one should favor my account. After all, if both accounts explain initial puzzles about moral beliefs, and do so by appealing to epistemic possibilities, then why should we say that what explains this difference is that moral beliefs have a higher epistemic standard, rather than that they are accompanied by more defeaters?

³³ Defeaters that undermine justification are commonly referred to as “mental state defeaters,” as opposed to “propositional defeaters” which are not believed by the agent and only undermine knowledge. On mental state and propositional defeaters see Michael Bergmann, *Justification without Awareness* (New York: Oxford University Press, 2006) on mental state and propositional defeaters.

For example, some may think that my Higher Standards view sacrifices important intuitions regarding the relation between evidence and defeaters by always viewing defeaters as relevant possibilities. On my view, the relationship between evidence and defeaters involves the introduction of new possibilities. This makes it seem as though while one's epistemic position worsens, one's evidence doesn't worsen at all—that is, one's epistemic position worsens despite one's evidence not worsening at all. But this seems to sacrifice a very intuitive thought that one's evidence gets worse with the presence of defeaters. Instead of raising epistemic standards, defeaters are typically conceptualized under the second interpretation of the More Defeaters view, wherein they render one's belief insufficiently justified by just simply reducing the strength of what serves as one's justification, for example one's evidence. Intuitively, we think that when one is told that there's a 50/50 chance that one was given an inverted color spectrum drug, it's not just that one's belief now fails to be justified, but that one's evidence has gotten *worse*, and fails to be justified *because* one's evidence has gotten worse. On a probabilistic model of evidence, the thought is as follows: while initially one's evidence may have made *p* probable to degree .9, when a defeater is introduced one's evidence now makes *p* probable to degree .5. However, as noted, this understanding of how one's evidence has gotten worse when a defeater is present is compatible with epistemic standards remaining at the same level. So, it might seem as though my Higher Standards account cannot account for the commonsensical thought that when defeaters get introduced one's evidence becomes worse.

While I agree that it would be problematic for my view if it was unable to account for this commonsensical thought, I don't believe that it faces this problem. To see this, we should return to our model. On a standard probability model, a defeater just functions to make *E* smaller (in other words, by making the not-*p* space bigger), where a certain probability is specified for an epistemic standard, and the probability that *p* is determined as follows (assuming for simplicity only finitely many possible worlds):

$$\text{Pr}(p) = \text{number of } p\text{-worlds in } E / \text{total number of worlds in } E$$

There is, however, an alternative way to think of how defeaters affect probability. On my model, it's true that when a defeater is introduced, the degree to which one's evidence makes *p* probable decreases. Rather than utilizing the above standard model of probability, though, my fallibilist view amends it as follows:

$$\text{Pr}(p) = \text{number of } p\text{-worlds in } E \text{ above } t_n / \text{total number of worlds in } E \text{ above } t_n$$

While on this model of probability it's true that one's evidence is *worse* in the sense of yielding a lower probability of p at t_2 (post-defeater) than at t_1 (pre-defeater), it has gotten worse precisely because the standard has gone up. So, this alternative model can show how the probability of p given one's evidence has gotten worse when a defeater is present in a way that doesn't make the raising of epistemic standards irrelevant. Since my proposed way of understanding defeaters in terms of possibilities can accommodate the sense in which one's evidence has gotten *worse* when a defeater is introduced, it ought not be abandoned

Another reason to favor my Higher Standards account is if it explains some cases that this interpretation of the More Defeaters account doesn't. Some of this may turn on the precise theoretical explanation for the higher epistemic standard; for example, if we endorse a kind of impurist view wherein the practical stakes of holding a belief affects the degree of justification the belief has, then the More Defeaters view would be an insufficient explanation of the degree of justification. To see why this would be the case, take the classic bank cases as an example.³⁴ Here, the proposition that the bank could've changed its hours isn't properly characterized as a defeater, since it's not properly characterized as higher-order evidence (that is, it's not evidence that your first order evidence (that you were at the bank last Saturday) does not warrant your belief (that the bank is open on Saturdays)). Rather, something like the proposition that you were only dreaming that you were at the bank last Saturday would be higher-order evidence. If we should conceive of the way justification is determined for moral beliefs as analogous to the bank cases (namely where the possibilities an agent must be able to rule out in order to have a justified moral belief is partly determined by what's practically at stake in holding the belief), then this More Defeaters view will be ruled out as the best explanation.

Moreover, it can also be said that in so far as defeaters introduce just *one* type of possibilities, or hold that possibilities can be introduced in just *one* way, my Higher Standards view will be able to explain more cases, and more diverse cases, as possibilities are introduced in multiple ways (the presence of disagreement isn't the only way to introduce a possibility). These are all reasons to favor my Higher Standards account over the first interpretation of the More Defeaters account, even if the More Defeaters view is not a genuine rival to my favored Higher Standards view.

³⁴ Keith DeRose, "Solving the Skeptical Problem," *The Philosophical Review* 104, 1 (1992): 1–52.

6. Conclusion

Moral testimony, moral expertise, and moral disagreement have all been thought to be distinctively problematic—that is, problematic in ways non-moral testimony, non-moral expertise, and non-moral disagreement are not. Previous explanations of their problematic nature have been piecemeal in nature, seeking to explain why each issue is problematic in isolation. In this paper, I've offered a unified explanation of the problematic nature of these issues, the Higher Standards account, thus departing from previous explanatory accounts of these phenomena. According to this unified account, the relative epistemically problematic nature of moral testimony, moral expertise, and moral disagreement is explained by the fact that moral beliefs typically enjoy a higher epistemic standard than non-moral beliefs. After first explaining my Higher Standards account, I considered two rival unified accounts that would explain the problematic nature of moral testimony, moral expertise, and moral disagreement, namely the Morality is Hard view and the More Defeaters view. I argued that these accounts were either rife with counterexamples, were ad hoc, or reduced to a variant of my view, concluding that my Higher Standards account is the best unified explanation on offer.³⁵

³⁵ Acknowledgements: I am especially grateful to Hille Paakkunainen and Nathaniel Sharadin for their many written comments on multiple drafts of this paper. I also thank Teresa Bruno-Nino, Janice Dowell, Matthias Jenny, David Sobel, Preston Werner, and the Women's Group of the philosophy department at Syracuse University for helpful comments and conversations, as well as audiences at the 2015 Rocky Mountain Ethics Congress and the 2016 Pacific Division Meeting of the American Philosophical Association.

STAKES-SHIFTING CASES RECONSIDERED—WHAT SHIFTS? EPISTEMIC STANDARDS OR POSITION?

Kok Yong LEE

ABSTRACT: It is widely accepted that our initial intuitions regarding knowledge attributions in stakes-shifting cases (e.g., Cohen's Airport) are best explained by standards variantism, the view that the standards for knowledge may vary with contexts in an epistemically interesting way. Against standards variantism, I argue that no prominent account of the standards for knowledge can explain our intuitions regarding stakes-shifting cases. I argue that the only way to preserve our initial intuitions regarding such cases is to endorse position variantism, the view that one's epistemic position may vary with contexts in an epistemically interesting way. Some had argued that epistemic position is incompatible with intellectualism. In reply, I point out that position variantism and intellectualism are compatible, if one's truth-relevant factors with respect to *p* can vary with contexts in an epistemically interesting way.

KEYWORDS: contextualism, stakes-shifting cases, relevant alternative, epistemic standard, intellectualism, sensitivity

1. Introduction

Let us say that *S* knows that *p* only if *S*'s *epistemic position* with respect to *p* satisfies the *standards for knowledge* in play. I take *S*'s epistemic position with respect to *p* to be a *placeholder* indicating the properties the having enough of which will render *S*'s true belief that *p* knowledge.¹ It is customary to talk about the *strength* of *S*'s epistemic position with respect to *p* or to compare the strength of *S*₁'s epistemic position with respect to *p* to the strength of *S*₂'s epistemic position with respect to *q*. The standards for knowledge specify how strong *S*'s epistemic position with respect to *p* has to be in order for *S* to know that *p*.

We may call the view that the standards for knowledge may vary with contexts in an epistemically interesting (non-trivial) way 'standards variantism.' *Epistemic contextualism* is a kind of standards variantism. On contextualism, the

¹ See Keith DeRose, *The Case for Contextualism* (Oxford: Clarendon Press, 2009), 7.

standards for knowledge vary with the *attributor* context.² What John MacFarlane calls *relativism* is also a kind of standards variantism.³ On relativism, the standards for knowledge vary with the *assessor* context.

Many have argued that standards variantism (contextualism in particular) is directly supported by ordinary cases of knowledge attributions of the following sort:

Airport. Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, "Yes, I know—it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to check with the airline agent.⁴

Suppose that Smith believes truly that the flight stops in Chicago. Intuitively, both Smith's knowledge attribution "I know the flight stops in Chicago" and Mary and John's attribution "Smith does not know that the flight stops in Chicago" seem true. This intuition is puzzling since Smith's belief that the flight stops in Chicago is held fixed, and the only difference between Smith's situation and Mary and John's is that the stakes in whether the flight stops in Chicago are high for Mary and John but low for Smith. Stakes are pragmatic, non-*truth-relevant*, factors that seem to have no direct bearing on whether one knows or not (at least initially).

Many have claimed that standards variantism offers the best explanation of our intuitions regarding knowledge attributions in *stakes-shifting cases*⁵ such as *Airport*; stakes-shifting cases are supposed to provide *prima facie* support for standards variantism. Call this *the standards-variantist assumption*.

² See David Kaplan, "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals," in *Themes from Kaplan*, eds. Joseph Almog, John Perry, and Howard Wettstein (Oxford: Oxford University Press, 1989), 481–566.

³ See John MacFarlane, "The Assessment Sensitivity of Knowledge Attributions," in *Oxford Studies in Epistemology*, eds. Tamar Szabó Gendler and John Hawthorne, vol. 1 (Oxford: Clarendon Press, 2005), 197–233.

⁴ Stewart Cohen, "Contextualism, Skepticism, and the Structure of Reason," *Philosophical Perspectives* 13 (1999): 58.

⁵ This term is from Jonathan Schaffer, "The Irrelevance of the Subject: Against Subject-Sensitive Invariantism," *Philosophical Studies* 127, 3(2006): 87–107.

In this paper, I argue against the standards-variantist assumption. My first thesis is that preserving our initial intuitions regarding knowledge attributions in stakes-shifting cases requires us to endorse *position variantism*, the view that the subject's epistemic position with respect to *p* varies with stakes-shifting cases. I will first reconstruct the main argument for the standards-variantist assumption (Section 2). I will then argue that no prominent account of the standards for knowledge can account for our stakes-shifting cases (Sections 3-6).

One might argue that position variantism is incompatible with *intellectualism*,⁶ a view that the factors that turn one's true belief into knowledge are exclusively truth-relevant. I reject this argument. More precisely, my second thesis is that position variantism is compatible with intellectualism if one's truth-relevant factors with respect to *p* vary with stakes-shifting cases. I will point out that the last view is far from being implausible. I will first argue that the position-variantist explanation can account for our intuitions regarding stakes-shifting cases (Section 7). I will also consider an important objection to position variantism (Section 8).

2. The Argument for the Standards-Variantist Assumption

Contextualists argue that stakes-shifting cases provide *prima facie* support for contextualism. For instance, Keith DeRose has claimed that stakes-shifting cases are "the best ground"⁷ for the contextualist theory:

[Stakes-shifting cases] provide us with the best possible type of evidence you could ask for that 'know(s)' is context-sensitive in at least roughly the way contextualists claim it is.⁸

Similarly, Stewart Cohen also claims that:

[*Airport*], and others like it, strongly suggests that ascription of knowledge are context-sensitive. The standards that determine how good one's reasons have to be in order to know are determined by the context of ascription.⁹

⁶ The term is from Jason Stanley, *Knowledge and Practical Interests* (Oxford: Clarendon Press, 2005). Stanley's formulation is slightly different from the one offered here. Jeremy Fantl and Matthew McGrath have discussion a very similar view, which they call 'purism about knowledge.' See Jeremy Fantl and Matthew McGrath, *Knowledge in an Uncertain World* (Oxford: Oxford University Press, 2009). Also see Footnote 35 below.

⁷ DeRose, *The Case for Contextualism*, 47.

⁸ DeRose, 67.

⁹ Cohen, "Contextualism, Skepticism, and the Structure of Reason," 59.

DeRose's and Cohen's remarks are misleading at best. Even if stakes-shifting cases did indicate the variability of the standards for knowledge, they would remain neutral to different forms of standards variantism. Specifically, they do not support DeRose and Cohen's favorite type of standards variantism (i.e., contextualism) over other types of standards variantism (e.g., nonindexical contextualism, relativism¹⁰).

Other standards variantists regard stakes-shifting cases as directly supporting standards variantism. Here are MacFarlane's remarks:

If I was speaking literally both times [in stakes-shifting cases] and didn't make a mistake, then *presumably the standards I must meet in order to count as "knowing" must have changed*. I met the laxer standards that were in play [in the first case], but not the stricter ones that come into play [in the second case].¹¹

Standards variantists often take the standards-variantist assumption for granted, so much so that few have bothered to justify the assumption. Most effort, rather, has been dedicated to showing that stakes-shifting cases, when elaborated, support one type of standards variantism over another.

Many non-standards variantists also are sympathetic to the standards-variantist assumption. Richard Feldman, for instance, has hypothesized the variability of the standards for knowledge:

It may be that knowledge attributions are context dependent. Perhaps the ordinary standards for knowledge are somehow flexible. Perhaps, setting aside the typical skeptical problems for a moment, it is sometimes true to say that a person knows a proposition and sometimes true to deny that the person knows that same proposition. Thus, for example, maybe the standards for knowledge shift in such a way that in casual conversation just prior to an election for which there are reliable polls indicating a clear winner, it is correct to say that we know what the outcome will be. Maybe in other contexts stricter standards apply and it is not correct to say that. That makes contextualism correct.¹²

Why has the standards-variantist assumption been so widely accepted? The remainder of this section elaborates what I take to be the main argument for the standards-variantist assumption.

¹⁰ For nonindexical contextualism, see Berit Brogaard, "In Defence of a Perspectival Semantics for 'Know,'" *Australasian Journal of Philosophy* 86, 3 (2008): 439–59 and John MacFarlane, "Nonindexical Contextualism," *Synthese* 166 (2009): 231–50. For relativism, see MacFarlane, "The Assessment Sensitivity of Knowledge Attributions" and Mark Richard, "Contextualism and Relativism," *Philosophical Studies* 119 (2004): 215–42. Both nonindexical contextualism and relativism are able to explain (most) stakes-shifting cases in much the same way as contextualism does.

¹¹ MacFarlane, "The Assessment Sensitivity of Knowledge Attributions," 201; my italics.

¹² Richard Feldman, "Contextualism and Skepticism," *Philosophical Perspectives* 45 (1999): 111.

Let us begin with analyzing the stakes-shifting cases, which are abundant in recent literature.¹³ Such cases are designed to share two features: (a) the same knowledge-attributing sentence, say, ‘S knows that p,’ seems true when uttered in a context C₁ but false when uttered in another context C₂, and (b) C₁ and C₂ are basically identical except that someone’s stakes in whether p are high in C₂ but low in C₁. Let us call C₁ and C₂ ‘LOW’ and ‘HIGH’ respectively (indicating low-stakes and high-stakes contexts respectively).

Let us focus on *Airport*. We intuitively think that:

- (1) Smith’s utterance “I (Smith) know that the flight stops in Chicago” is true in LOW, and Mary and John’s utterance “Smith does not know that the flight stops in Chicago” is true in HIGH.¹⁴

On the standards-variantist assumption, (1) directly supports standards variantism; standards variantism provides the best explanation of (1).

The truth value of “S knows that p” is determined by whether S knows that p or not. Traditionally, whether S knows that p or not is regarded as depending on (a) whether S believes that p or not, (b) whether p is true or not, and (c) whether S’s epistemic position with respect to p satisfies the standards for knowledge in play or not.

(A) and (b) are not the issues here, since they remain constant across stakes-shifting cases. For instance, Smith believes truly that the flight stops in Chicago in both LOW and HIGH. Our target is (c). That is, the variability of the truth value of “S knows that p” in stakes-shifting cases is generated by S’s epistemic position with respect to p satisfying the standards for knowledge in LOW but not in HIGH.

The following, hence, is plausible:

- (2) If (1), then Smith’s epistemic position with respect to *The flight stops in*

¹³For instance, the *Bank Case* from Keith DeRose, “Contextualism and Knowledge Attributions,” *Philosophy and Phenomenological Research* 52, 4 (1992): 913–29 and the *Train Case* from Jeremy Fantl and Matthew McGrath, “Evidence, Pragmatics, and Justification,” *The Philosophical Review* 111, 1 (2002): 67–94. Also see Stanley, *Knowledge and Practical Interests*, 3–5.

¹⁴ Joshua May et al., “Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study,” *Review of Philosophy and Psychology* 1, 2 (2010): 265–73; Wesley Buckwalter, “Knowledge Isn’t Closed on Saturday: A Study in Ordinary Language,” *Review of Philosophy and Psychology* 1, 3 (2010): 395–406. For reply, see Keith DeRose, “Contextualism, Contrastivism, and X-Phi Surveys,” *Philosophical Studies* 156 (2011): 81–110, and Jonathan Schaffer and Joshua Knobe, “Contrastive Knowledge Surveyed,” *Noûs* 46, 4 (2012): 675–708. It is not my intention to settle the issue here. Those who do not think that (1) is our intuitive judgment regarding *Airport* are invited to consider the present thesis as a conditional claim: if (1) is indeed our intuitive judgments, then it is best explained by position variantism.

Chicago satisfies the standards for knowledge in play in LOW but fails to do so in HIGH.

The consequent of (2) can be satisfied in two prominent ways:

- (2a) Smith's epistemic position with respect to *The flight stops in Chicago* varies with LOW and HIGH.
- (2b) The standards for knowledge in play vary with LOW and HIGH.

Hence, I assume that:

- (3) If Smith's epistemic position with respect to *The flight stops in Chicago* satisfies the standards for knowledge in play in LOW but fails to do so in HIGH, then either (2a) Smith's epistemic position with respect to *The flight stops in Chicago* vary with LOW and HIGH, or (2b) the standards for knowledge in play vary with LOW and HIGH.

It is tempting to deny the first disjunct of the consequent of (3):

- (4) Smith's epistemic position with respect to *The flight stops in Chicago* does not vary with LOW and HIGH.

(4) seems plausible. After all, Smith's relation to *The flight stops in Chicago* is held fixed across LOW and HIGH—Smith's evidence for *The flight stops in Chicago*, for instance, seems to remain the same in LOW and HIGH (however, I will argue that we should deny (4). See Section 7).

(1)-(4) entail:

- (5) The standards for knowledge in play vary with LOW and HIGH.

(1)-(5) constitute a very strong argument for the standards-variantist assumption. The argument is valid, and all of its premises seem plausible. No wonder the standards-variantist assumption is widely accepted.

The argument is not sound, however. The following four sections examine all prominent accounts of the standards for knowledge. I argue that all of them fail to support (5). Admittedly, such an argument strategy does not offer any conclusive objection against standards variantism. But if what is said below is correct, it will at least show that standards variantism is ill-motivated.

3. The Nature of the Standards for Knowledge

On the standards-variantist assumption, the difference in linguistic dispositions between Smith and Mary/John is best explained by the difference in the standards for knowledge between LOW and HIGH. This provides us with a basis to test the standards-variantist assumption. That is, if our intuitions concerning Smith's and Mary and John's knowledge-attribution dispositions manifest the systematic

differences predicted by the variation of the standards for knowledge, standards variantism offers (or at least is in a position to offer) a proper explanation for stakes-shifting cases; otherwise, not.

A question immediately arises: what is the nature of the standards for knowledge? Without knowing what the standards for knowledge are, the standards-variantist assumption is simply non-evaluable. Here, my strategy is to examine all (prominent) accounts of the standards for knowledge, evaluate them individually, and see whether or not they can account for our intuitions concerning knowledge attributions in stakes-shifting cases.

Jonathan Schaffer has offered a useful framework for my project.¹⁵ As Schaffer points out, the term ‘the shift of the standards for knowledge,’ when used by contextualists, may mean three different things: in his terms, the shift of ‘threshold,’ ‘standard,’ or ‘alternative.’ Schaffer focuses mainly on contextualism, but his framework can be extended to cover standards variantism in general without losing its plausibility. At any rate, this is how I will proceed. A caveat: the following will focus on how *contextualists* characterize epistemic standard rather than how other *non-contextualist* epistemologists characterize epistemic standard. While it is worth extending the following argument to cover what non-contextualists epistemologists have to say on this topic, this goes beyond the scope of the present study. As will become clear, we will have enough on our plate.

Following Schaffer’s framework, I will distinguish two main accounts of the standards for knowledge, which I call the *general* and *particular accounts*. According to the general account, the shift of the standards for knowledge will affect any proposition with a certain property. The effect of the standards for knowledge over the logical space is, in Schaffer’s term, “globally encompassing.” The particular account contends that the shift of the standards for knowledge affects only a specific set of propositions that does not form a globally encompassing logical space. The effect of the standards on the logical space is, in Schaffer’s term, “pointlike.”

As a start, it is useful to describe the structure of my argument: according to our intuitions, a certain knowledge attribution, say, “S knows that p,” is true in LOW but false in HIGH. In principle, we can find a proposition q ($q \neq p$) such that q is an *epistemic counterpart of p for S* in LOW, while q is an epistemic counterpart of p for S in C if and only if the strength of S’s epistemic position with respect to q is on epistemic par with the strength of S’s epistemic position with respect to p in C. Suppose that S believes truly that q. Intuitively, we take “S knows that q” to be true in LOW as well. I will construct a certain epistemic counterpart q

¹⁵ See Schaffer, “What Shifts? Thresholds, Standards, or Alternatives?”

of p for S in LOW such that we intuitively think that “ S knows that p ” is false in HIGH, but that “ S knows that q ” is true in HIGH. I then argue that neither the general nor particular account is able to account for our intuitions regarding such knowledge attributions. More precisely, the general account has no resources for predicting that “ S knows that q ” is true in HIGH, while the particular account has no non-arbitrary way of making the prediction.

4. The Linear Account

There are two types of the general account. Call them the *linear* and *spherical accounts*. The linear account takes the standards for knowledge as specifying a *threshold* for S 's epistemic position with respect to p such that S can be truthfully described as “knows that p ” only if S 's epistemic position with respect to p meets the threshold.

Two prominent linear accounts suggest themselves. According to the *evidentialist account*, the standards for knowledge set the threshold for the strength of evidence such that the strength of S 's evidence for p must reach a certain threshold in order for S to be counted as “knows that p .” On the *reliabilist account*, by contrast, the standards for knowledge set the threshold for the degree of reliability of belief-forming processes such that the degree of reliability of the process that forms S 's belief that p must meet the threshold in order for S to be counted as “knows that p .” I will focus on the reliabilist account, but the same point also applies to the evidentialist account.

Suppose that, in LOW, the degree of reliability of the belief-forming process of Smith's belief that the flight stops in Chicago is .80 (‘the degree of reliability of Smith's belief is .80’ in short). Moreover, suppose that, according to the standards for knowledge in play, the threshold of the degree of reliability in HIGH is .95, while the threshold in LOW is .75. The reliabilist account typically explains *Airport* as follows:

Smith's knowledge claim “I know that the flight stops in Chicago” is true in LOW since the degree of reliability of Smith's belief that the flight stops in Chicago meets the (low) threshold for the degree of reliability. By contrast, Mary and John's denial of knowledge “Smith does not know that the flight stops in Chicago” is also true in HIGH since the degree of reliability of Smith's belief that the flight stops in Chicago is held fixed, and the latter does not meet the (high) standards in play in HIGH.

To test the reliabilist account's explanation, let us consider the epistemic counterparts of *The flight stops in Chicago* for Smith in LOW, i.e., Smith's beliefs whose degree of reliability is also .80 in LOW.

Stakes-Shifting Cases Reconsidered—What Shifts? Epistemic Standards or Position?

Consider:

Airport'. Everything is like *Airport* except that Mary and John and Smith are friends, and that Mary and John know that Smith is a Lakers fan. Suppose that Mary knows that the Lakers won yesterday, and she also knows that Smith believes that the Lakers won yesterday since she saw him reading today *The New York Times*. Suppose that, in LOW, Smith's epistemic position with respect to *The flight stops in Chicago* is epistemically equivalent to his epistemic position with respect to *The Lakers won yesterday*. Suppose that, on their way to find the airline agent to check on the flight schedule, Mary and John are chatting. John asks Mary whether Smith knows that the Lakers won yesterday.

Consider two possible situations:

Airport(a). Mary says to John, "Yes, Smith knows the Lakers won yesterday."

Airport(b). Mary says to John, "No, he doesn't. Smith believes truly that the Lakers won yesterday, but he does not know."

Intuitively, *Airport(a)*, rather than *Airport(b)*, is the natural reply. I contend that a correct account of the standards for knowledge should be able to handle *Airport'*.

Unfortunately, however, the way the reliabilist account handles *Airport* cannot be employed to account for *Airport'*. Let us elaborate.

Airport' is so stipulated such that:

- (6) In LOW, Smith's epistemic position with respect to *The flight stops in Chicago* is epistemically equivalent to his epistemic position with respect to *The Lakers won yesterday*.

On the reliabilist account, the strength of one's epistemic position with respect to p should be characterized by the degree of reliability of one's belief that p. Hence, from (6), the reliabilist account implies:

- (7) In LOW, the degree of reliability of Smith's belief that the flight stops in Chicago is the same as the degree of reliability of Smith's belief that the Lakers won yesterday.

It is worth noting that standards variantism is compatible with position variantism. Proponents of standards variantism, however, usually assume that one's epistemic position does not vary with stakes-shifting cases—most (if not all) standards variantists are position invariantists. This assumption seems natural given that the standards-variantist assumption implies that one's epistemic position is *not* responsible for the shift in our intuitions regarding knowledge attributions in stakes-shifting cases. At any rate, proponents of standards variantism have endorsed:

- (8) Smith's epistemic position with respect to p does not vary with LOW and

Kok Yong Lee

HIGH.

(7) and (8) entail:

- (9) In HIGH, the degree of reliability of Smith's belief that the flight stops in Chicago is the same as the degree of reliability of Smith's belief that the Lakers won yesterday.

Arguably, (9) leads to:

- (10) In HIGH, Smith knows that the flight stops in Chicago if and only if Smith knows that the Lakers won yesterday.

But, according to our initial intuitions regarding *Airport'*, it seems that:

- (11) In HIGH, Smith does not know that the flight stops in Chicago.

From (10) and (11), we can conclude:

- (12) In HIGH, Smith does not know that the Lakers won yesterday.

This argument is valid. Proponents of the reliabilist account, as I see it, must accept all the premises. However, (12) indicates that *Airport(b)* rather than *Airport(a)* is the natural follow-up of *Airport'*. Counterintuitive.

Since infinitely many counterexamples can be constructed along this line and against another linear account such as the evidentialist account, the prospect of the linear account is dim. Without further ado, let us consider the other type of the general account.

5. The Spherical Account

On the spherical account, possible worlds can be ordered to form a certain "sphere" with the actual world serving at its center. The distance between possible worlds is often regarded as a function of the (overall) similarity between them.¹⁶ The strength of S's epistemic position with respect to p is determined by the number of possible worlds in (or the area of) a possible-world sphere in which S can maintain a certain (epistemic) relation R to p. The stronger S's epistemic position with respect to p, the larger the number of possible worlds in (or the area of) the possible-world sphere in which S can maintain a relation R to p in the possible-world sphere.

The standards for knowledge specify an *area* of a sphere, which can be used to evaluate different spheres. A sphere may have an area identical to, smaller than, or larger than the one specified by the standards for knowledge. S's epistemic

¹⁶ See Duncan Pritchard, *Epistemic Luck* (New York: Clarendon Press, 2005); Wolfgang Freitag, "Safety, Sensitivity and 'Distant' Epistemic Luck," *Theoria* 80 (2014): 44–61.

position with respect to p satisfies the standards for knowledge if and only if S can maintain a relation R to p in the sphere whose area is no smaller than the area specified by the standards. Using the notion of an area of a sphere, the same standards for knowledge can be applied to different subjects and/or different propositions. For instance, to say that the strength of S_1 's epistemic position with respect to p is stronger than the strength of S_2 's epistemic position with respect to q (where $S_1 \neq S_2$ and $p \neq q$) is tantamount to saying that the area of the sphere with respect to which S_1 can maintain a relation R to p is larger than the area of the sphere with respect to which S_2 can maintain a relation R to q .

Different spherical accounts will formulate the (epistemic) relation R differently. In this section, I will only examine the *sensitivity account*, which is the most prominent spherical account. On this account, S 's epistemic position is determined by the *sensitivity* of S 's belief that p with respect to a sphere s , where S 's belief that p is sensitive with respect to s if and only if S believes truly that p in the center world and, for all possible worlds w in s , (i) p is false in w , and (ii) S does not believe that p in w .¹⁷ That the strength of S_1 's epistemic position with respect to p is stronger than S_2 's epistemic position with respect to q implies that the area of the sphere with respect to which S_1 's belief that p is sensitive is larger than the area of the sphere with respect to which S_2 's belief that q is sensitive. The standards for knowledge specify an area r such that the area of the sphere with respect to which S 's belief that p is sensitive must not be smaller than r in order for S to be counted as "knows that p ." In other words, S 's epistemic position with respect to p satisfies the standards for knowledge if and only if the area of the sphere with respect to which S 's belief that p is sensitive is no smaller than the area specified by the standards. Hence, to say that the standards for knowledge are context-sensitive amounts to saying that the area of the sphere with respect to which one's belief that p must remain sensitive in order to be counted as "knows that p " may vary with contexts.¹⁸

The sensitivity account typically explains *Airport* as follows:

The flight stops in Chicago in the actual world. The possible worlds in which the flight does *not* stop in Chicago form a possible-world sphere with the actual world serving as its center. Let the possible-world sphere s' be a subset of s such that, for all w' in s' , Smith does not believe that the flight stops in Chicago in w' , and the possible-world sphere s'' be a subset of s and of which s' is a proper subset such that, for some w'' in s'' , Smith believes that the flight stops in Chicago in

¹⁷ See Robert Nozick, *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press, 1981).

¹⁸ See Keith DeRose, "Solving the Skeptical Problem," *Philosophical Review* 104, 1 (1995): 1–52.

w". It follows that Smith's belief that the flight stops in Chicago is sensitive with respect to s' but insensitive with respect to s'' .

In LOW, the standards for knowledge are relatively low such that they specify a (relatively small) area r , which is smaller than the area of s' . Since Smith's belief that the flight stops in Chicago is sensitive with respect to s' , his belief satisfies the standards for knowledge in LOW. Therefore, "Smith knows that the flight stops in Chicago" is true in LOW.

However, in HIGH, the standards for knowledge are stringent such that they specify a (relatively large) area r' , which is larger than the area of s'' . Since Smith's belief that the flight stops in Chicago is insensitive with respect to s'' , his belief does not satisfy the standards for knowledge in HIGH. Therefore, "Smith does not know that the flight stops in Chicago" is also true in HIGH.

The way the sensitivity account handles *Airport*, however, cannot be employed to account for *Airport'*. Let us elaborate.

Airport' is so stipulated such that:

- (6) In LOW, Smith's epistemic position with respect to *The flight stops in Chicago* is epistemically equivalent to his epistemic position with respect to *The Lakers won yesterday*.

According to the spherical account, to say that S_1 's epistemic position with respect to p is epistemically equivalent to S_2 's epistemic position with respect to q amounts to saying that the area of the sphere with respect to which S_1 's belief that p remains sensitive is identical to the area of the sphere with respect to which S_2 's belief that q remains sensitive. That is, (6) is tantamount to:

- (13) The area of the sphere s' with respect to which Smith's belief that the flight stops in Chicago remains sensitive in LOW is identical to the area of the sphere v with respect to which Smith's belief that the Lakers won yesterday remains sensitive in LOW.

As noted, standards variantists assume that:

- (8) Smith's epistemic position with respect to p does not vary with LOW and HIGH.

(8) and (13) entail:

- (14) The area of the sphere s' with respect to which Smith's belief that the flight stops in Chicago remains sensitive in HIGH is identical to the area of the sphere v with respect to which Smith's belief that The Lakers won yesterday remains sensitive in HIGH.

On the sensitivity account's explanation of *Airport*, the following holds:

- (15) In HIGH, Smith's belief that the flight stops in Chicago is insensitive.

(15) indicates that the area of the sphere s' with respect to which Smith's belief that the flight stops in Chicago remains sensitive is smaller than the area specified by the standards for knowledge in HIGH. However, given (14), it follows that the area of the sphere v with respect to which Smith's belief that the Lakers won yesterday remains sensitive is also smaller than the area specified by the standards for knowledge in HIGH. That is, from (14) and (15), we get:

(16) In HIGH, Smith's belief that the Lakers won yesterday is insensitive.

Since S knows that p only if S 's belief that p is sensitive, (16) thus implies:

(12) In HIGH, Smith does not know that the Lakers won yesterday.

The argument above is valid. Proponents of the sensitivity account, as I see it, must accept all the premises. (12), however, indicates that *Airport(b)* rather than *Airport(a)* is the natural follow-up of *Airport'*. Counterintuitive.

Since infinitely many counterexamples can be constructed along the similar line, the argument above indicates that the sensitivity account fails to explain our intuitions regarding cases like *Airport* (and *Airport'*). Hence, the spherical account is not promising in accounting for stakes-shifting cases, either.

Let us take stock. We have seen that the general account fails to predict the variations of the attributor's linguistic disposition in stakes-shifting cases. Both the linear and spherical accounts have suffered from a similar problem, namely, they fail to account for our intuitions regarding some epistemic counterparts q of p for S in stakes-shifting cases. In the present case, they mistakenly take *Airport(b)*, as opposed to *Airport(a)*, as the natural follow-up of *Airport'*. The moral, I take it, is this: in stakes-shifting cases, we intuitively think that the attributors' linguistic dispositions do not undergo a global variation. When the attributor counts the subject as "know that p " in LOW but "does not know that p " in HIGH, the attributor does not, at least not always, also take the subject as "does not know that q " in HIGH for all epistemic counterparts q of p for S in LOW. The fundamental flaw of the general account is its implication that the attributors exhibit a certain global variation in their linguistic dispositions in stakes-shifting cases.

6. The Particular View

Some standards variantists opt for the particular account according to which the variations of the standards for knowledge do not have a global impact on the attributor's linguistic dispositions. The most prominent particular account is the *relevant alternative account (of the standards for knowledge)* (hereafter 'the RA account'). The RA account characterizes the variations of the standards for knowledge in terms of the variations of the set of *relevant alternatives*. The idea is

that the subject S counts as “knowing that p” only if S’s epistemic position with regard to p enables S to rule out all relevant alternatives to p, specified by the standards for knowledge in play.

The RA account typically explains *Airport* as follows:

In *Airport*, Smith’s evidence is unable to rule out the proposition that *The itinerary contains a misprint*. In LOW, however, the standards for knowledge do not specify the proposition to be a relevant alternative. As a result, Smith counts as “knows that the flight stops in Chicago” even though he is unable to rule out *The itinerary contains a misprint*. By contrast, the standards for knowledge in HIGH do specify *The itinerary contains a misprint* to be a relevant alternative. Accordingly, Mary and John’s utterance “Smith does not know that the flight stops in Chicago” is true in HIGH.

Some might offer a similar RA account’s explanation for *Airport’*:

In *Airport’*, Smith’s evidence is unable to rule out the proposition, say, that *The New York Times contains a misprint*. In HIGH, however, the standards for knowledge do not specify the proposition to be a relevant alternative. As a result, Smith counts as “knows that the Lakers won yesterday” in HIGH.

The core of the RA account’s explanation of *Airport* and *Airport’* consists of:

RA1: The alternative *The flight itinerary contains a misprint* is relevant in HIGH.

RA2: The alternative *The New York Times contains a misprint* is irrelevant in HIGH.

Obviously, proponents of the RA account need to explain *why* RA1 and RA2 hold. Otherwise, the RA account’s explanation is just an *ad hoc* story tailor-made to account for whatever our intuitions are with respect to stakes-shifting cases. What we need are principles that determine whether an alternative is relevant or not. Call them ‘principles of relevance.’ The importance of specifying the principles of relevance cannot be overemphasized. “The success of the RA approach,” as Vogel puts it, “depends upon there being a principled distinction between relevant and irrelevant alternatives.”¹⁹

David Lewis was the first to articulate a complex system of principles of relevance (hereafter ‘the Lewisian system’).²⁰ For the present purposes, it is appropriate to focus on the Lewisian system, since it remains one of the most comprehensive accounts on the market.

¹⁹ Jonathan Vogel, “The New Relevant Alternatives Theory,” *Philosophical Perspectives* 13 (1999): 168.

²⁰ David Lewis, “Elusive Knowledge,” *Australasian Journal of Philosophy* 74, 4 (1996): 549–67.

A key feature of the Lewisian system is that its principles are motivated solely by our pre-theoretically intuitive judgments concerning knowledge attributions, an approach much like what Roderick Chisholm calls “particularism.”²¹ Such a particularist approach is especially dubious if our intuitions lead to theoretically incoherent principles, or if they generate conflicting judgments. Indeed, such problems crop up in the Lewisian system. As Lewis notices, some of his principles will lead to skepticism if not restrained by certain *ad hoc* conditions²²—Lewis bravely and blatantly admits that he has no solution to this problem. Lewis also considers the possibility that people may have intuitions incompatible with the Lewisian system’s verdicts, and his response, surprisingly again, is simply to give in and conclude that “we have reached a standoff.”²³

While I believe that such general problems are very serious—and they may very well be the root of all the problems the Lewisian system eventually faces—I will not pursue the issue further. Instead, I will continue the argumentative strategy of this paper and focus on the question whether or not the Lewisian system can account for stakes-shifting cases involving epistemic counterparts such as *Airport*. My answer is that it cannot.

I offer two reasons. First, the most promising Lewisian principle to handle RA2 (i.e., the *reliability principle*; see below) has failed to do its job. Second, and perhaps worse, even granted that the reliability principle can handle RA2, the principle still has an undesirable result of falsifying RA1, and there seems to be no way, not according to the Lewisian system anyway, to get rid of this undesirable result. I will elaborate them respectively.

Recall that if RA2 is to hold at all, it has to be sanctioned by the principles of relevance specifying the sufficient condition for an alternative being irrelevant. Lewis calls them “permissive principles.” He listed four such principles, but only one of them is particularly relevant to RA2. Here is how Lewis characterizes the principle: “Within limits, we are entitled to take [perception, memory, and testimony] for granted.”²⁴ We may formulate this principle as follows:

The reliability principle. If an alternative *p* is incompatible with the assumption that perception, memory, and testimony do not fail, then, defeasibly, *p* is not relevant.

²¹ Roderick Chisholm, *The Foundations of Knowing* (Minneapolis: University of Minnesota Press, 1982), 66.

²² Lewis, “Elusive Knowledge,” 556.

²³ Lewis, 561.

²⁴ Lewis, 558.

The reliability principle appears to handle RA2 quite well. The alternative *The New York Times contains a misprint* is incompatible with the assumption that Smith's testimony (i.e., *The New York Times*) does not fail. Hence, *The New York Times contains a misprint* is irrelevant in HIGH.

Appearances are deceptive, however. A closer scrutiny reveals that there is in fact no room for the Lewisian system to account for RA2. To elaborate, notice that the reliability principle is defeasible—some other principles may overthrow its verdict. Moreover, it does seem that the reliability principle's verdict on RA2 is overthrown by two other Lewisian principles.

"The possibility that actually obtains," according to Lewis, "is never properly ignored."²⁵ Hence, the following holds:

The actuality principle. If *p* is true in the actual world, then, defeasibly, *p* is always relevant.²⁶

Moreover, Lewis also thinks that if an alternative resembles another, and if "one of them may not be properly ignored, neither may the other."²⁷ This gives us:

The resemblance principle. If *p* saliently resembles *q*, and *p* is relevant (in virtue of principles other than this one), then, defeasibly, *q* is relevant.

The actuality and resemblance principles together imply that RA2 does not hold. The actual world in which *The New York Times* does not contain a misprint saliently resembles the counterfactual world in which *The New York Times* contains a misprint—after all, the only difference between these two worlds is that *The New York Times* contains a misprint in one but not in the other. By the actuality principle, *The New York Times does not contain a misprint* is relevant. Moreover, since *The New York Times does not contain a misprint* is relevant, the alternative *The New York Times contains a misprint* is also relevant, by the resemblance principle. That is, the actuality and resemblance principles overthrow the reliability principle's verdict on *The New York Times contains a misprint*.

Lewis himself has also noted that the application of the actuality principle and the resemblance principle has to be restricted; otherwise they will lead to skepticism.²⁸ Unfortunately, Lewis admits that he does not know how to give a non-ad-hoc restriction. While this is surely a vice for the Lewisian system in general, it might be interpreted as a virtue for the particular account, as proponents of the particular account can then claim that the aforementioned denial of RA2 is

²⁵ Lewis, 554.

²⁶ By actuality, Lewis means the subject's actuality, see Lewis, 554.

²⁷ Lewis, 556.

²⁸ Lewis, 556–57.

implausible as it is based on an unrestricted usage of the actuality principle and the resemblance principle. Whether this is a promising reply depends on whether a non-ad-hoc restriction on the usage of these two principles can be offered. However, this just brings us back to the problem Lewis is facing. And it seems that the prospect of solving it is dim.

Worse, even if we granted that the reliability principle is able to account for RA2, the Lewisian system still fails to account for *Airport'*. In fact, the reliability principle is too strong: not only does it rule that the alternative *The New York Times contains a misprint* is irrelevant in HIGH (i.e., RA2), it also rules that the alternative *The flight itinerary contains a misprint* is irrelevant in HIGH (i.e., the denial of RA1). For not only is the alternative *The New York Times contains a misprint* incompatible with the assumption that (Smith's) testimony (i.e., *The New Yorks time*) does not fail, the alternative *The flight itinerary contains a misprint* is also incompatible with the assumption that (Smith's) testimony (i.e., the flight itinerary) does not fail. In other words, the reliability principle is able to obtain RA2 only at the expense of RA1.

As far as I can tell, there is only one way to get around this problem, that is, to argue that the reliability principle's verdict on *The flight itinerary contains a misprint* is overturned by other principles of relevance. In what follows, I will examine two most promising candidates that can do the job. I find both of them wanting. More precisely, each of the proposed principles is either implausible on its own or too strong such that not only does it overthrow the reliability principle's verdict on *The flight itinerary contains a misprint*, it also overthrows the principle's verdict on *The New York Times contains a misprint*. In other words, either the proposed principles are implausible, or they are able to handle RA2 only at the expense of RA1.

First, on Lewis's account, "[an alternative] not ignored at all is *ipso facto* not properly ignored."²⁹ That is:

The attention principle. If an alternative p is entertained (or is not ignored), then, defeasibly, p is relevant.

The attention principle implies both RA1 and RA2. On the one hand, since Mary and John are considering whether the flight itinerary contains a misprint, *The flight itinerary contains a misprint* is relevant in HIGH (thereby overthrowing the reliability principle's verdict, as desired). On the other hand, since *The New York Times contains a misprint* has not been entertained, its status of relevance is not affected.

²⁹ Lewis, 559.

This reply is problematic since whether *The New York Times contains a misprint* is relevant or not does not seem to depend on whether the possibility of *The New York Times* containing a misprint is entertained in HIGH or not. For instance, suppose that we modify *Airport(a)* as follows:

*Airport(a)**. Mary says to John, “Yes, Smith knows that the Lakers won yesterday.
The New York Times rarely misprints the result of basketball games.”

Intuitively, *Airport(a)** is still more natural than *Airport(b)*. Yet, in *Airport(a)**, *The New York Times contains a misprint* is not ignored and should be counted as relevant in HIGH by the attention principle, contradicting RA2.

More importantly, the attention principle is implausible. It is widely agreed that merely entertaining an alternative does not automatically render it relevant.³⁰ John Hawthorne mentions a possible refinement:

The attention principle'. If an alternative *p* is *seriously* entertained, then, defeasibly, *p* is relevant.³¹

Arguably, Mary and John have entertained seriously *The flight itinerary contains a misprint* but not *The New York Times contains a misprint*. Hence, the attention principle' predicts that the former is a relevant alternative (thereby overthrowing the reliability principle's verdict, as desired), while leaving the latter's status of relevance intact. Unfortunately, the attention principle' is still implausible as serious attention is not always a sufficient condition for the relevance of alternatives. For instance, a paranoid agent could entertain remote alternatives seriously, but that does not automatically render them relevant.³²

Let us look at another proposal. Lewis also thinks that “when error would be especially disastrous, few possibilities are properly ignored.”³³ Hence:

The high-stakes principle. If the stakes of *p* are high, then, defeasibly, few alternative *q* to *p* are irrelevant.

Mary and John's stakes in whether or not the flight itinerary contains a misprint are high, but their stakes in whether or not *The New York Times* contains a misprint are low. At first sight, the high-stakes principle appears to correctly predict that *The flight itinerary contains a misprint* is a relevant alternative in

³⁰ John Hawthorne, *Knowledge and Lotteries* (Oxford: Clarendon Press, 2004), 63–64; Michael Blome-Tillmann, “Knowledge and Presuppositions,” *Mind* 118, 470 (2009): 246–47.

³¹ Hawthorne, *Knowledge and Lotteries*, 64.

³² See Vogel, “The New Relevant Alternatives Theory,” 164ff for more objections to the attention principle.

³³ Lewis, “Elusive Knowledge,” 566, Footnote 12.

HIGH (thereby overthrowing the reliability principle's verdict, as desired) while leaving the status of relevance of *The New York Times contains a misprint* intact.

The high-stakes principle's prediction of RA1, however, is problematic. Specifically, the way the principle renders *The flight itinerary contains a misprint* relevant can be exploited to show that it would be too difficult for Mary and John to obtain the knowledge of the flight's layover. Suppose that Mary and John go to check the flight schedule with the airline agent and learn that the flight stops in Chicago. Intuitively, they thus know that the flight stops in Chicago—but wait, since the stakes are very high, the alternative that the airline agent has misread the flight information on the screen should also be relevant, and Mary and John are not in a position to rule this alternative out. So, they should not know that the flight stops in Chicago after all. Suppose that Mary and John ask the airline agent to double check and are told the same answer. Do they know now? Not if they are able to rule out the alternative that the airline agent has been impatient with them and does not double-check properly—since the stakes are high, this alternative should be relevant as well. But they are in no position to rule out that alternative. So, they still do not know whether the flight stops in Chicago... This shows that the high-stakes principle is not a very plausible explanation of RA1 in the first place.

Let us take stock. In this section, we consider whether or not the RA account (or, more precisely, the Lewisian system) is able to handle stakes-shifting cases involving epistemic counterparts. Specifically, we ask the question whether or not the Lewisian system can give rise to RA1 and RA2. Our examination tells us that it cannot. On the one hand, the Lewisian system does not seem to have the theoretical resources to handle RA2 in the first place, as the reliability principle's verdict is overthrown the actuality and resemblance principles. On the other hand, granted that the reliability principle can handle RA2, the Lewisian system now has difficulties handling RA1. Either the principles that get us RA1 are implausible on their own (i.e., the attention principle, the attention principle', the high-stakes principle) or they are able to handle RA1 only at the expense of RA2 (i.e., the reliability principle, the attention principle).

I conclude that the Lewisian system (the RA account) is unable to give a non-ad-hoc explanation of why RA1 and RA2 hold. The Lewisian system (the RA account), hence, fails to account for stakes-shifting cases like *Airport'*.

7. Position Variantism and Stakes-Shifting Cases

Where are we now? We have seen that the standards-variantist assumption is supported by the argument (1)–(5). If the previous three sections are correct, the conclusion (i.e., (5)) is threatened.

Since (1)–(5) is valid, (at least) one of (1)–(4) has to be rejected. Which one? Arguably, (2) and (3) are safe. (2) relies on two intuitive ideas. Firstly, the widely accepted account of knowledge that S knows that p if and only if S believes truly that p, and S's epistemic position with respect to p satisfies the standards for knowledge in play. Secondly, S believes truly that p in both LOW and HIGH. (3) is also unproblematic, since it seems plausible that Smith's epistemic position with respect to *The flight stops in Chicago* and the standards for knowledge in play do *not* vary with both LOW and HIGH. At any rate, I will not challenge (3) here.

This leaves us (1) and (4); either the truth value of "Smith knows that the flight stops in Chicago" does not vary with LOW and HIGH (i.e., not-(1)), or Smith's epistemic position with respect to p does vary with LOW and HIGH (i.e., not-(4)). Put more generally, either we reject our initial intuitions regarding knowledge attributions in stakes-shifting cases—i.e., the truth value of "S knows that p" varies with LOW and HIGH—or we endorse what we may call *position variantism*, the view that the subject's epistemic position with respect to p may vary with contexts in an epistemically interesting way (in particular, S's epistemic position with respect to p may vary with LOW and HIGH).

I suggest that we should opt for position variantism, if only to respect our initial intuitions. I will not develop a full-fledged account of position variantism here. What I will do, rather, is to point out that position variantism is very promising in explaining stakes-shifting cases. This gives us the incentive to be serious about the view.³⁴ I will then discuss and reject one possible objection to position variantism (Section 8).

On the position-variantist explanation, the variability of the truth value of "S knows that p" in stakes-shifting cases is derived from the variability of S's epistemic position with respect to p in such cases. More precisely, S's epistemic position with respect to p is weaker in HIGH than it is in LOW such that, while the standards for knowledge in play remain constant across LOW and HIGH, S's epistemic position with respect to p satisfies the standards for knowledge in LOW

³⁴ An additional motivation of position variantism comes from the fact that this view (or something along similar lines) can offer the most plausible "contextualist" solution to skepticism. See Kok Yong Lee, "On the Standards-Variantist Solution to Skepticism," *International Journal for the Study of Skepticism* 7, 3 (2017): 173–98.

but fails to do so in HIGH. Let us, again, focus on *Airport*. A typical position-variantist explanation is as follows:

Smith's utterance "I know that the flight stops in Chicago" is true in LOW since Smith's epistemic position with respect to *The flight stops in Chicago* is strong enough to satisfy the standards for knowledge in play. By contrast, Mary and John's utterance "Smith does not know that the flight stops in Chicago" is also true in HIGH since Smith's epistemic position with respect to *The flight stops in Chicago* is not strong enough to satisfy the (same) standards for knowledge in play.

To test whether this is a plausible account, we may consider *Airport'*. As noted, we intuitively think that *Airport(a)*, rather than *Airport(b)*, is the natural follow-up of *Airport'*.

To its credit, position variantism can deliver this verdict. According to position variantism, Smith's epistemic position with respect to *The flight stops in Chicago* varies with LOW and HIGH, but this, by itself, does not imply that Smith's epistemic position with respect to *The Lakers won yesterday* (or any other epistemic counterparts of *The flight stops in Chicago* for Smith) vary with LOW and HIGH.

8. Position Variantism vs. Intellectualism

Position variantism has its criticisms. One main worry is that position variantism is incompatible with a very plausible epistemic principle:

Intellectualism. The factors that turn one's true belief into knowledge are exclusively truth-relevant.³⁵

In the present terminology, intellectualism is tantamount to the view that S's epistemic position with respect to p depends exclusively on S's truth-relevant factors with respect to p.

Notice that many are willing to assume that:

- (17) S's truth-relevant factors with respect to p do not vary with LOW and HIGH.

³⁵ See DeRose, *The Case for Contextualism*, 24. For recent criticisms to intellectualism, see Hawthorne, *Knowledge and Lotteries*; Stanley, *Knowledge and Practical Interests*; Fantl and McGrath, *Knowledge in an Uncertain World*; Matthew McGrath, "Contextualism and Intellectualism," *Philosophical Perspectives* 24 (2010): 383–405. For a recent defense of intellectualism, see Kok Yong Lee, "On Two Recent Arguments against Intellectualism," *NCCU Philosophical Journal*, forthcoming.

(17), together with intellectualism, imply that S's epistemic position with respect to p does not vary with LOW and HIGH (i.e., the denial of position variantism). In other words, intellectualism and position variantism are incompatible.

I want to draw a different conclusion, however. I think intellectualism and position variantism are compatible; we should give up (17). Arguing against (17), however, goes beyond the scope of this paper. Here, I will settle for a weaker conclusion. I will argue, instead, that the denial of (17) is not totally implausible, and that philosophers are in no position to simply take (17) for granted.

To make my point, it suffices to show that there are theories on the market which (a) are position-variantist in character, and (b) can be employed to falsify (17). That is, theories that imply:

- (18) S's truth-relevant factors with respect to p may vary with LOW and HIGH in an epistemically interesting way.

So long as such theories are plausible enough to deserve serious attention, philosophers cannot simply assume (17).

For simplicity's sake, I will focus on one such theory. Ram Neta has developed a theory of evidence whose core idea is that what counts as one's evidence for p may vary with contexts in an epistemically interesting way.³⁶ It should be noted that although Neta calls himself a "contextualist," his account does not belong to standard-variantism as defined here.³⁷

Suppose that the set of propositions $\{p_1, \dots, p_n\}$ is S's evidence for p_{n+1} .³⁸ Details aside, Neta argues that S's evidence for p_{n+1} is affected by the following rule:

- (R) If at time t , S_1 raises a hypothesis H that is an uneliminated counterpossibility with respect to S_2 's knowing that p , then S_2 's body of evidence at t is to just those p_i of $\{p_1, \dots, p_n\}$ that is introspectively available to S , at t whether or not H is true,

where a hypothesis H is an uneliminated counterpossibility with respect to S_2 's knowing that p just in case (i) H implies that S does not know that p at t , and (ii) H

³⁶ Ram Neta, "Contextualism and the Problem of the External World," *Philosophy and Phenomenological Research* 66, 1 (2003): 1–31. Elsewhere, I discuss Neta's account in more detail. See Lee, "On the Standards-Variantist Solution to Skepticism," 13ff. Stanley also mentions but does not develop a similar view, but he finds this line of thought "overwhelmingly plausible," see Stanley, *Knowledge and Practical Interests*, 124.

³⁷ This just means that while Neta is happy to label himself as a 'contextualist,' his view (or at least an interpretation of it) is in fact radically different from the orthodox contextualism.

³⁸ Neta takes one's evidence to be one's mental states. In this paper, I will simply take one's evidence to be a set of propositions. I do not take a particular stance on the nature of evidence here.

and “S knows that p at t” are introspectively indistinguishable for S.³⁹ Crucial to (R) is the idea that what counts as one’s evidence for p varies with contexts, depending on the alternative H to p that is in question in the present context.

Evidence is a kind of truth-relevant factors. If Neta is right, S’s truth-relevant factors (i.e., evidence) with respect to p may vary with contexts in an epistemically interesting way, depending (in part) on whether an uneliminated counterpossibility is raised or not. This result supports (18), for Neta’s theory can be naturally extended to account for cases like *Airport*.

For instance, in HIGH, Mary has raised an uneliminated counterpossibility that the itinerary contains a misprint. According to (R), Smith’s evidence for *The flight stops in Chicago* would be restricted to those propositions that are introspectively available to Smith, whether or not the itinerary contains a misprint or not. Presumably, Smith’s evidence, so restricted, no longer contains propositions such as *The itinerary does not contain a misprint*. “Smith knows that the flight stops in Chicago” is true in LOW but not in HIGH since Smith’s truth-relevant factors (i.e., evidence) with respect to *The flight stops in Chicago* have varied across LOW and HIGH, i.e., not-(17).

I do not intend to argue for Neta’s theory. Nor do I think Neta’s account is the only plausible, or tenable, approach to position variantism. My point, rather, is to highlight the fact that it is not impossible, nor implausible, to argue for (a kind of) position variantism that leads to (18).

In other words, the fact that Neta’s theory can accommodate (18) should not be taken as a direct justification for (18) and/or position variantism. Rather, this fact suggests that one is in no position to simply assume that position variantism and intellectualism are incompatible. A lot more still needs to be said about whether position variantism is in fact a correct view; likewise for whether (18) is a correct characterization of stakes-shifting cases like *Airport*.

9. Conclusion

I have shown that (a) our initial intuitions regarding knowledge attributions in stakes-shifting cases, (b) position variantism, and (c) standards variantism are jointly inconsistent. I suggest that we should give up standards variantism. Focusing on *Airport*, I argue that all prominent accounts of the standards for knowledge have failed to deliver the correct verdicts on our intuitions regarding stakes-shifting cases.

³⁹ Neta, “Contextualism and the Problem of the External World,” 23–24.

The moral is this: the only way we can preserve our initial intuitions regarding knowledge attributions in stakes-shifting cases is to appeal to position variantism. This is an important lesson, since the position-variantist explanation has been largely overlooked in the contemporary literature. One reason that position variantism has slipped under the radar, perhaps, is that many have regarded it as incompatible with intellectualism.

I have shown that position variantism and intellectualism are compatible. The price for marrying position variantism with intellectualism is to endorse the view that *S*'s truth-relevant factors with respect to *p* may vary with contexts in an epistemically interesting way. Is this the right price to pay? I do not give an answer here. If what has been argued above is correct, perhaps this is the price we have to pay if we are going to respect our intuitions regarding stakes-shifting cases.⁴⁰

⁴⁰ I want to Matt McGrath for helpful comments on an earlier draft of this paper. This paper is funded by the Ministry of Science and Technology (MOST) of Taiwan (R.O.C.) (MOST 107-2410-H-194-MY2)

ON SOME ARGUMENTS FOR EPISTEMIC VALUE PLURALISM¹

Timothy PERRINE

ABSTRACT: Epistemic Value Monism is the view that there is only one kind of thing of basic, final epistemic value. Perhaps the most plausible version of Epistemic Value Monism is Truth Value Monism, the view that only true beliefs are of basic, final epistemic value. Several authors—notably Jonathan Kvanvig and Michael DePaul—have criticized Truth Value Monism by appealing to the epistemic value of things other than knowledge. Such arguments, if successful, would establish Epistemic Value Pluralism is true and Epistemic Value Monism is false. This paper critically examines those arguments, finding them wanting. However, I develop an argument for Epistemic Value Pluralism that succeeds which turns on general reflection on the nature of value.

KEYWORDS: epistemic value monism, epistemic value pluralism, epistemic value, Jonathan Kvanvig, Michael DePaul

On a widely held view, true beliefs are of final epistemic value. An interesting question is whether, and why, anything else is. Some authors hold that truth is the most basic thing of final epistemic value, embracing a version of “Epistemic Value Monism” that is sometimes called “Vertisim” or “Truth Value Monism.” Other authors demur, maintaining that the epistemic value of truth cannot explain the epistemic value of everything. Such authors embrace a kind of “Epistemic Value Pluralism.” The debate between Epistemic Value Monists and Pluralists is an important one. For instance, some philosophers might be inclined to understand other epistemic categories—e.g., epistemic obligations or epistemic virtues and vices—in terms of their relation to epistemic value. Clearly settling what is of epistemic value would be important for such projects.

Various arguments have been given against Truth Value Monism and in favor of Epistemic Value Pluralism. We can separate those arguments into two categories. *Knowledge based Arguments* argue that because the epistemic value of truth cannot explain the epistemic value of knowledge we must embrace Epistemic Value Pluralism to explain the epistemic value of knowledge. *Non-Knowledge based Arguments* argue that the epistemic value of truth cannot explain the

¹ For helpful feedback, I thank Dan Buckley, Jordi Cat, Dave Fisher, Adam Leite, Dan Linsenbardt, Mark Kaplan, Tim O'Connor, Andrew Smith, and Harrison Waldo.

epistemic value of things besides knowledge. In other work,² I have discussed Knowledge based Arguments and will not discuss them here.

Rather, the aim of this paper is to examine Non-Knowledge based Arguments for Epistemic Value Pluralism. I will argue several such arguments fail—they are implausible, obscure, actually consistent with Truth Value Monism, or neglect the relevant distinction between basic and non-basic final value (see below). Nonetheless, I will claim that there is one Non-Knowledge based Argument for Epistemic Value Pluralism that succeeds. That argument turns on plausible general claims about final value.

After setting the stage in section I, I examine an argument for Epistemic Value Pluralism due to Jonathan Kvanvig in section II. I show how his argument is too obscure to carry force. In section III, I focus on a more straightforward argument from Kvanvig on the nature of understanding. But I argue Kvanvig's view is actually consistent with Truth Value Monism. In section IV, I examine a sequence of arguments from Michael DePaul, including one about the appropriateness of responding to experience. I argue that DePaul's account is implausible and a more plausible one is consistent with Truth Value Monism. Finally, in section V, I argue that there is an argument for Epistemic Value Pluralism that is plausible that turns on the idea that it is of final value to value what is of final value. I defend this argument from an objection based on an alternative account of the nature of epistemic value.

I. Stage Setting

By “final epistemic value,” I have in mind the kind of value that, from the epistemic point of view, is valuable in and of itself or for its own sake. I will not offer any analysis of the epistemic point of view. I assume that final epistemic value is a kind or species of final value. I do not assume that final epistemic value exhausts all the kinds of final value.

It is widely thought that final value has some sort of important connection to valuing. There are different accounts of this connection. Some offer a deontic connection: when something is of final value, we are obligated to value it; others offer a rationalist connection: when something is of final value, it is rational to value it; some offer a reason based account: when something is of final value, there is reason to value it. My own preferred view is that something is of final value just

² Timothy Perrine, *Epistemic Value and Accurate Representation* (PhD Dissertation: Indiana University, 2017), 146-188.

when it is appropriate to value it. I will assume it in what follows. Little will hang on this internal dispute in what follows.

In evaluating the dispute between Epistemic Value Monism and Epistemic Value Pluralism it will be important to have an account of the distinction between the two. However, one natural way of formulating the distinction between them is problematic. Specifically, it is natural to suppose that proponents of Epistemic Value Monism hold:

For any x and some P , if x is of final epistemic value, then x is P .

Different proponents of Epistemic Value Monism may give different accounts of P . For instance, on this way of thinking, Truth Value Monists hold:

For any x , if x is of final epistemic value, then x is a true belief.

Epistemic Value Pluralists would then be people who reject this general characterization.

This way of formulating the dispute is problematic because it contains a problematic characterization of Value Monism. As I've argued elsewhere,³ value monists will want their view to have ontological flexibility. That is, they will want a wide range of things to be of final value—e.g., outcomes of actions, lives, entire possible worlds, etc. But this view will not have the desired flexibility. (Outcomes of actions, for instance, are not true beliefs.) Instead, we should think that what makes that position a monistic one is not that it maintains that only one kind of thing is of final value. What makes it monistic is that it maintains that any time something is of final value an explanation of its final value will refer to its connection to one kind of thing. For these reasons, proponents of Epistemic Value Monism should reject:

For any x and some P , if x is of final epistemic value, then x is P .

For this is too narrow a view of what is of value even by the monist's lights.

To understand the dispute between Value Monists and Pluralists, it will be helpful to introduce some terminology. Let us say something is of "basic final value" just when it is of final value but there is no explanation of the final value it has in terms of other things of final value.⁴ Something is of "non-basic final value" just when it is of final value but there is an explanation of the final value it has in

³ Timothy Perrine, "Basic Final Value and Zimmerman's *The Nature of Intrinsic Value*," *Ethical Theory and Moral Practice* 21, 4 (2018): 979-996

⁴ This leaves open that there is an explanation of why something is of basic final value that does not appeal to final value. In this way, the distinction does not assume that epistemic value is, in some important meta-normative sense, reducible or not.

terms of other things of final value. Disputes between Value Monists and Value Pluralists can then be understood as disputes about basic final value. Specifically, in this context, Epistemic Value Monism holds:

For any x and some P , if x is of basic final epistemic value, then x is P .

Epistemic Value Pluralists would reject this. But both Epistemic Value Monists and Pluralists can agree that a wide range of otherwise metaphysically distinct things are of final epistemic value. Their dispute is simply over whether, to explain such facts, we need to appeal to the final value of one thing or many.⁵

Understood in this way, Epistemic Value Monism says that there is only one kind of thing that is of basic final epistemic value. However, this leaves open what exactly is of basic final epistemic value. There could be different “versions” of Epistemic Value Monism that identify different kinds of things as being of basic final epistemic value. For purposes of this paper, I will focus on Truth Value Monism, understood here as the position that the only kinds of things that are of basic final epistemic value are true beliefs. I will focus on this version of Epistemic Value Monism because it is the usual foil to Epistemic Value Pluralism.

II. A Plurality of Evaluations—Kvanvig

One critic of Epistemic Value Monism is Jonathan Kvanvig. In an important paper defending Epistemic Value Pluralism, he urges that seeing truth as the sole or fundamental goal has a “strong reductionist flavor.”⁶ To get us to see this, he first characterizes epistemology as “the study of purely theoretical cognitive success,” and urges that we see value in each “independent kind of cognitive success” so that what is of final epistemic value would include a wide range of things including “knowledge, understanding, wisdom, rationality, justification, sense-making, and empirically adequate theories in addition to getting to the truth and avoiding error.”⁷ For ease of reference, let’s call this list ‘Kvanvig’s laundry list.’

However, it is not clear exactly what argument against Epistemic Value Monism Kvanvig intends to be defending. Perhaps Kvanvig’s argument is this:

⁵ For more on basic final value and issues involving Monism vs Pluralism, Fred Feldman, “Basic Intrinsic Value,” *Philosophical Studies* 99, 3 (2000): 319-46; Fred Feldman, *Pleasure and the Good Life* (Oxford: Clarendon Press, 2004); Michael Zimmerman, *The Nature of Intrinsic Value* (Landam: Rowman & Littlefield, 2001); Perrine, “Basic Final Value and Zimmerman’s *The Nature of Intrinsic Value*.”

⁶ Jonathan Kvanvig, “Truth Is Not the Primary Epistemic Goal,” in *Contemporary Debates in Epistemology*, eds. Matthias Steup and Ernest Sosa (Malden, MA: Blackwell Publishing, 2005), 287.

⁷ Kvanvig, “Truth Is Not the Primary Epistemic Goal,” 287.

On Some Arguments for Epistemic Value Pluralism

(P1) If something is a “purely theoretical cognitive success,” then it is of final epistemic value.

(P2) There are many independent kinds of purely theoretical cognitive successes.

(C1) So, Epistemic Value Monism is false.

In defense of the second premise, Kvanvig may point to his laundry list. However, once we’ve drawn the distinction between basic final epistemic value and non-basic final epistemic value, we can see that this argument is invalid. For even if there are many kinds of “purely theoretical cognitive successes” it may still be that their value is always explained by appealing to one kind of thing of basic final epistemic value. Once we recognize this distinction, pointing to a plurality of things of final epistemic value cannot, in and of itself, show Epistemic Value Monism false.

Kvanvig might shore up this argument by maintaining that:

(P1) If something is a “purely theoretical cognitive success,” then it is of final epistemic value.

(P2) There are many independent kinds of purely theoretical cognitive successes.

(P3) At least one of the independent kinds of purely theoretical cognitive successes have final epistemic value that cannot be explained by appealing to a single kind of basic, final epistemic value.

(C1) So, Epistemic Value Monism is false.

This argument avoids the problem of the previous one. But the inclusion of the terminology “purely theoretical cognitive success” is now unnecessary. If (P1)-(P3) is true, then a weaker set of premises will also produce a valid argument against Epistemic Value Monism:

(P4): There is at least one thing of final epistemic value whose final value cannot be explained by appealing to a single kind of basic final epistemic value.

(C1) So Epistemic Value Monism is false.

At this point, it appears that the terminology of “purely theoretical cognitive success” is, at best, doing no necessary work and, at worse, is unduly obscure.

The best way to defend (P4) would be through existential generalization—to give an example of something that is of final epistemic value whose value cannot be explained by the final epistemic value of one thing like truth. I will focus on whether Kvanvig has given us any promising examples of this. The most promising example would be understanding. I focus on it next.

III. Kvanvig on Understanding

Kvanvig argues that the value of understanding is not explained by the value of true belief. If Kvanvig is correct about this, then we have an argument for (P4). This section critically examines Kvanvig's argument.

Kvanvig focuses on the kind of understanding at issue when one understands that something is the case.⁸ To use his example, consider someone's understanding of the "Comanche dominance of the southern plains of North America from the late seventeenth until the late nineteenth centuries."⁹ Kvanvig makes three key claims about this kind of understanding. First, it is, for the most part, factive: such a person has a large number of true beliefs, and in so far as they have false beliefs on the subject matter, those false beliefs are peripheral.¹⁰ Second, these true beliefs need not amount to knowledge. A person whose true beliefs are "Getterized"—who, for instance, by pure coincidence picks up a book which contains true claims about the Comanche which were, nevertheless, shots in the dark by the author—can still possess understanding. In this way, understanding is not "a species" of knowledge.¹¹ Finally, understanding requires "grasping" the relations between the items of knowledge, specifically the way in which that information "coheres" with one another.¹²

The value of understanding, as Kvanvig sees it, derives from two places. First, it derives from the number of true beliefs that help make up understanding. But, secondly, it derives from the "grasping" that is required for understanding. Kvanvig writes:

[To account for the value of understanding] we need to return to the notion of subjective justification, the value of which was defended earlier. Subjective justification obtains when persons form or hold beliefs on the basis of their own

⁸ Jonathan Kvanvig, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003), 189-90.

⁹ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 197.

¹⁰ Kvanvig does not spend much time on what counts as peripheral; neither will I. For critical discussion, see Catherine Elgin, "Is Understanding Factive?," in *Epistemic Value*, eds. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford: Oxford University Press, 2009) and Wayne Riggs, "Understanding, Knowledge, and the *Meno* Requirement," in *Epistemic Value*, eds. Haddock, Millar, and Pritchard.

¹¹ Again, I neither endorse nor deny this claim. For skepticism regarding it, see Stephen Grimm, "Is Understanding a Species of Knowledge?," *British Journal for the Philosophy of Science* 52, 3 (2006): 515-535 and Duncan Pritchard, "Knowledge and Understanding," in *The Nature and Value of Knowledge*, eds. Duncan Pritchard, Alan Millar, and Adrian Haddock (Oxford: Oxford University Press, 2010), 77-80.

¹² Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 197, 202.

subjective standards for what is true or false.¹³

We thus get the following explanation of the value of understanding. The distinctive element involved in it, beyond truth, is best understood in terms of grasping of coherence relations. Such coherence relations in this context contribute to justification. Such justification is subjective, because the person in question must grasp the marks of truth within that body of information in order to grasp correctly the explanatory relationships within that body of information.¹⁴

So, on Kvanvig's view, the grasping of coherence relations helps lead to subjective justification, and because the latter is valuable, the former is as well.

Clearly, if Kvanvig's view on the value of understanding is inconsistent with Truth Value Monism it will be inconsistent because of his view on the valuing of graspings. On Kvanvig's view, graspings lead to subjective justification, which he claims is of value. He distinguishes between two kinds of extrinsic value.¹⁵ The first kind is the standard *instrumental* value, where something is of instrumental value when (roughly) it is an effective means to a valuable end, increasing the likelihood of securing that value.¹⁶ Kvanvig considers a second kind of extrinsic value, which need not be an effective means to a valuable end but is rather an "intentional means." An action (for instance) is an intentional means to a valuable end, when a person undertakes that action with the aim of achieving that valuable end.¹⁷ The distinction between an effective means and an intentional means are illustrated in cases where there is no action that I can perform that will make it more likely that I'll achieve a valuable end, but nevertheless there are actions I can undertake with the aim of achieving that valuable end. To use Kvanvig's own example, perhaps there is nothing I can do to sink a basketball shot from half court and win a million dollars—so that there are no effective means to that end—but there are actions I can perform with the aim of achieving the end—so there are intentional means.¹⁸

The notion of intentional means thus far developed only applies to actions. But, Kvanvig claims, it can also be extended to beliefs.¹⁹ Thus, consider a person who follows their own standard—whatever it is—for getting at the truth. Let's say that a person's belief is *subjectively justified* when it is held in accordance with their own standard.²⁰ Even if the person's own standards are woefully

¹³ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 200.

¹⁴ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 202.

¹⁵ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 60-5.

¹⁶ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 63.

¹⁷ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 60.

¹⁸ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 60-1.

¹⁹ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 65-75.

²⁰ Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, 56.

inadequate—so that it is not an effective means to get to the truth by following those standards—following those standards will be an intentional means and thus valuable. Thus, subjective justification is valuable as a kind of intentional means.

One might object to Kvanvig's argument at several places here. One might argue that understanding does not require subjective justification. Or one might argue that the notion of intentional means cannot apply to belief. Or one might argue that intentional means are not extrinsically valuable. But none of these objections are necessary to defend Truth Value Monism. For Truth Value Monism is a position about *final* epistemic value—it is a thesis about what is valuable for its own sake. But it is perfectly consistent to accept Truth Value Monism and hold there are many different kinds of things with *extrinsic* epistemic value. For instance, one might hold that reliable belief forming processes are valuable but only extrinsically, specifically, instrumentally because they are likely to lead to true beliefs. But even if we follow Kvanvig and “loosen up” extrinsic value to allow for another kind of extrinsic value distinct from instrumental value, this is perfectly consistent with Truth Value Monism. Consequently, Kvanvig's account of the value of understanding provides no problem for this version of Epistemic Value Monism.²¹

IV. DePaul against Epistemic Value Monism

Another proponent of Epistemic Value Pluralism is Michael DePaul. DePaul criticizes Truth Value Monism before offering up his own version of Epistemic Value Pluralism. In what follows, I'll briefly sketch and respond to his criticisms of Truth Value Monism before discussing his positive view.

A. DePaul's Argument against Truth Value Monism

DePaul's first criticism goes:

... I think deep down we all recognize that truth is not the only thing of epistemic value. Here is an easy demonstration. Take your favorite well-established empirical theory, a theory you believe that we know. Throw in all the evidence on the basis of which we accept that theory. Depending on the theory you selected, all this will likely add up to a substantial number of beliefs. Now compare this set of beliefs with an equal number of beliefs about relatively simple

²¹ These points do not require that final epistemic value always supervenes on the intrinsic features of something. (In fact, Truth Value Monism probably could not say that, since truth is not an intrinsic property.) They only require that the category of extrinsic value is distinct from the category of final value in that being of extrinsic value does not entail being of final value, which is clearly true.

arithmetic sums and about assorted elements of one's current stream of consciousness. I suspect that most of us would want to say that the first set of beliefs is better, epistemically better, than the second set. But the two sets contain the same number of true beliefs. And so, to the extent that we are inclined to say that these sets differ with respect to broad epistemic value, it would seem that we are committed to saying that truth is not the only thing has broad epistemic value.²²

The thrust of his criticism is clear: the only way to accommodate a difference in epistemic value between these two sets is to postulate something *else* of epistemic value and embrace a kind of Epistemic Value Pluralism.

However, DePaul is wrong that the *only* way to accommodate this difference in epistemic value is to postulate something else of basic final epistemic value besides truth. First, one can retain Truth Value Monism and account for the difference of value between these two sets by appealing to the conditions under which truths have any epistemic value whatsoever. Specifically, one might hold that whether a set of truths has any epistemic value depends partially upon extrinsic (and contingent) features of the set. For instance, Goldman explicitly holds that epistemic value depends partially upon whether or not a person is *interested* in whether the relevant proposition is true or false.²³ Thus, *contra* DePaul, even if one can provide two sets with the same number of true beliefs, it does not follow that they have the same epistemic value, given this view on the conditions under which something has epistemic value.

Second, one can retain Truth Value Monism and account for the difference of value between the two sets by appealing to the particular *contents* of the truths. Part of the intuitive motivation behind DePaul's criticism is that truths about (e.g.) organic chemistry are more *important* than truths about (e.g.) what's going on right now on the left side of my visual field. One might try to cash out this importance in terms of the interest of inquirers, which would lead us back to a response similar to the one given in the previous paragraph. But one might cash out this importance in terms of the contents of the propositions themselves. The idea that some propositions are more "natural" or "cut nature at its joints" or are otherwise descriptively superior to others has gained some currency recently. One can hold that while any true belief has some epistemic value, a true belief has *more* epistemic value if its contents are more "natural." In this way, the particular *objects* of belief can play a role in determining the overall value of a set of beliefs. While

²²Michael DePaul, "Value Monism in Epistemology," in *Knowledge, Truth, and Duty*, ed. Matthias Steup (New York: Oxford University Press, 2001), 173.

²³Alvin Goldman, *Knowledge in a Social World* (New York: Oxford University Press, 1999); see also William Alston, *Beyond "Justification"* (Ithaca: Cornell University Press, 2005).

few have fully developed such a position, and I won't do so here, I see no reason why it cannot be. So DePaul is wrong that the only way to account for the difference in epistemic value between those two sets is to abandon Truth Value Monism.²⁴

B. DePaul's Argument for Epistemic Value Pluralism

Turning to DePaul's positive proposal, he argues that there are two kinds of things that have final epistemic value: true belief and warrant.²⁵ He follows Plantinga as holding that warrant is "the epistemic feature which plays the preeminent role in distinguish mere true belief from knowledge."²⁶ Nevertheless, in contrast to Plantinga, DePaul does not take truth-conduciveness to be necessary for warrant; warrant, for DePaul, is not believing in a way that is likely to be true. But if warrant is decoupled from believing in a truth-conducive way, what is it? DePaul despairs of giving a particularly helpful, positive account of warrant. He holds that it is believing "appropriately" specifically believing appropriately "in the face of experience."²⁷

DePaul gives an argument that believing appropriately in the face of experience is of final epistemic value. The argument is a thought experiment.²⁸ Imagine a non-deceiving demon. The demon does not aim to make most of your beliefs about your immediate environment false; rather, the demon aims to disrupt the connection between your experiences on one hand and your beliefs and the world on the other. To this end, the demon gives you a visual field as if you were watching old Laurel and Hardy movies. Nevertheless, you continue to believe that you are (e.g.) currently sitting, reading a paper even as a slapstick gap unfolds before your eyes.

²⁴ A similar kind of move to appeal to the particular *objects* of attitudes has been made by Fred Feldman in defending a form of hedonism; see his *Pleasure and the Good Life*. For a different kind of response to DePaul on this issue see Nick Treanor, "Trivial Truths and the Aim of Inquiry," *Philosophy and Phenomenological Research* 89, 3 (2014): 552-9.

²⁵ Michael DePaul, *Balance and Refinement* (London: Routledge Press, 1993), 77.

²⁶ DePaul, *Balance and Refinement*, 67.

²⁷ DePaul, *Balance and Refinement*, 82-3. There are problems with DePaul's view when applied to a wide range of cases of knowledge. Perhaps responding appropriately to one's experiences is important for distinguishing between true belief and knowledge for certain *kinds* of knowledge like perceptual or even testimonial knowledge. But it is not clear how it will apply to other cases, including not only *moral* knowledge (as DePaul is aware) but logical, mathematical, or inductive knowledge. I'll set aside these worries in what follows, though.

²⁸ DePaul, *Balance and Refinement*, 80-1, 191-2.

This case illustrates a breakdown of warrant, according to DePaul. In it, while many of one's beliefs may be true, they don't *appropriately fit* one's experience as of old movies. But it is not just that this case illustrates how warrant can breakdown; according to DePaul, it reveals an overly narrow conception of the epistemic value of *experience*. Truth Value Monists are concerned to evaluate the truth of *beliefs* and insofar as experiences are mentioned it is as instrumental to forming true beliefs. But experiences should play a more important role:

When one recognizes the possibility of correspondences among experiences, belief, and reality, it is easy to see that such a person's cognitive state may fall short of epistemic excellence. For it might be that there is the same sort of incoherence between the person's experience and his belief as epistemologists fear to find between belief and reality. And, I maintain, where there is such an experiential incoherence, we fall short of warrant and knowledge, no matter what the connection between our beliefs and truths.²⁹

It is not obvious how best to regiment DePaul's argument. I think the following captures it fairly. First, in the non-deceiving case, there is an "incoherence" between the experiences of the subject and the way the subject is forming beliefs that is disvaluable. Second, that disvalue cannot be understood in terms of the instrumental disvalue of forming false beliefs because the subject is forming true beliefs. Therefore, we must think that the disvalue is a kind of final epistemic disvalue. Thus, there is something of final epistemic value in forming beliefs that "appropriately fit" one's experiences and something of final epistemic disvalue in forming beliefs that do not "appropriately fit" one's experiences. Thus, Truth Value Monism is false.

Why think that, in the non-deceiving case, there is an incoherence between the experiences the subject is having and the way the subject is forming beliefs? DePaul assumes that we can tell there is something defective here by simply comparing experiences and beliefs. In other words, the appropriate belief response to one's experience supervenes just on those qualitative experiences, or sensations, themselves. Other facts—about whether, e.g., one is being messed with by an evil demon—are irrelevant. And it is natural for him to think this. After all, incoherence is an *internal* relation. So if there is an incoherence here it should be determined *solely* by the beliefs and experiences.

However, this view is implausible. For this view ignores the general or specific cognitive abilities of the cognizer having the experience. The relevance of a cognizer's cognitive abilities becomes clear when we consider less extreme examples. For instance, when I was seven and I had a certain olfactory experience,

²⁹ DePaul, *Balance and Refinement*, 86.

I did not form any beliefs about what caused it; now when I have the very same olfactory experience, I form the belief that someone is brewing coffee. It was inappropriate for me to form the belief that someone was brewing coffee then, it is not so now. Or, consider a novice bird watcher. Upon seeing a bird initially, it will be inappropriate for the bird watcher to believe it's a woodpecker (he's only started watching birds yesterday). But after a decade of watching birds, if the now expert bird watcher has the exact same visual experience, it would be appropriate to form a belief that it's a woodpecker. These examples show that what constitutes an appropriate response to experience doesn't supervene on just the experience the person has.

Indeed, there is a more principled reason for denying that appropriate responses supervene on just the experiences a cognizer has. Recall that, for DePaul, warrant is both responding appropriately to one's experience and the property that plays the chief role in distinguishing knowledge from mere true belief. There are (and could be) many different kinds of cognizers that know things, and even among cognizers of the same type or kind (such as human beings), there are many different kinds of things they know—different “sources of knowledge” as it is sometimes put. Consequently, if warrant is that property which helps account for the difference between knowledge and mere true belief in all (or even most) of these cases, warrant (or the degree of warrant) will presumably supervene partially on the different cognitive facilities of the different cognizers. But if, as DePaul claims, warrant is also responding appropriately to one's experience, then it follows that responding appropriately to one's experience will supervene partially on the different cognitive faculties of different cognizers.

In response, DePaul might press that even if it's not true, generally speaking, that what is appropriate to believe should supervene solely upon our experiences, surely in the cases provided above it is clear that *those* cognizers aren't responding appropriately to their experiences. But even this is doubtful. After all, in those cases, the non-deceiving demon has radically altered their cognitive faculties so that, really, the experiences they have are playing no role in how they are forming beliefs. But given how radically different that way of forming beliefs is from how we form beliefs, we should not be very confident that *not* responding to their experiences is the right way of “responding” to their experiences. So, I claim, it is not clear that in DePaul's non-deceiving demon case the subjects are forming beliefs in an “incoherent” way or a way that is inappropriate.

Additionally, when we think more about the role that cognitive abilities play in determining appropriate responses to experience, we are led back to the instrumentally valuable picture of experience. Specifically, it is natural to think

that the appropriateness of certain beliefs vary with the *reliability* or *truth-conduciveness* of a cognizer's cognitive abilities. The reason why it is appropriate for my current self to believe that someone is brewing coffee on the basis of a particular olfactory experience, but not my seven year old self, is that the former can very reliably pick out coffee by scent whereas the latter could not. Similarly, the expert bird watcher is much more reliable when it comes to identifying woodpeckers. This explains why it is appropriate for the expert, but not the novice, to believe a certain bird is a woodpecker on the basis of a certain visual experience.³⁰ (It's worth noting that when DePaul goes into detail about beliefs that are appropriate for him they are all cases of beliefs that were arrived at reliably.³¹) Thus, when we reflect on *how* cognitive abilities are relevant to the appropriateness of beliefs, we are most naturally pushed back to understanding warrant, i.e. responding to one's experiences appropriately, as having a close connection to truth-conduciveness and the instrumental model DePaul criticizes.

C. More on Epistemically Appropriate Responses to Experiences

In discussing DePaul's argument, I briefly argued that appropriately responding to one's experience required forming beliefs in a reliable or otherwise truth-conducive way. My argument for this turned on a discussion of how agents can learn to acquire beliefs on the basis of sensations. To be sure, I have not offered a full defense or development of these ideas. But I will briefly consider some alternative accounts of responding appropriately to one's experience. To be clear, even if these other accounts are right, it would still not yet show that responding to one's experiences appropriately or properly is of basic final epistemic value. We would still need an argument for that. Rather, they would at best undermine my positive proposal for the instrumental value of responding appropriately or properly to one's experience.

One account is Markie's.³² Broadly speaking, on Markie's account, when a response to an experience is "epistemically appropriate" it is because we have learned or otherwise know how to identify objects and their features on the basis of those experiences.³³ Markie then teases out three different "ways" a belief might

³⁰ Compare Alvin Goldman, "Towards a Synthesis of Reliabilism and Evidentialism," reprinted in *Reliabilism and Contemporary Epistemology*, ed. Alvin Goldman (New York: Oxford University Press, 2012).

³¹ DePaul, *Balance and Refinement*, 82-3.

³² Peter Markie, "Epistemically Appropriate Perceptual Belief," *Nous* 40, 1 (2006): 118-42.

³³ Markie, "Epistemically Appropriate Perceptual Belief," 123, 130, 139.

be epistemically appropriate.³⁴ However, a full review of Markie's account is unnecessary. For Markie thinks that a belief is "most fully" appropriate when it satisfies all three of his ways.³⁵ Additionally, one of those ways requires that the way the belief is formed is authorized by a reliabilist norm. So I doubt Markie's account is in deep tension with what I say here.

A different proposal would be to appeal to seemings. It is unclear what a seeming is, though most authors think they are *sui generis* mental states wherein a proposition is presented "as true" or "forcefully." So understood, seemings are not beliefs, inclinations to beliefs, or sensations.³⁶ The proposal would then be that while (e.g.) my 7 year old self and my current have the same sensations, I have a seeming that coffee is being brewed while my 7 year old self does not. Further, it is this difference of seemings that explains why it is epistemically appropriate for me to believe that someone is making coffee but it is not epistemically appropriate for my 7 year old. (Though he is speaking of justified beliefs, and not appropriate responses to experiences, Tucker offers essentially this view.³⁷)

Underlying this response is the view that, absent reasons for doubt, it is epistemically appropriate to believe that *p* if it seems to one that *p*. But such a view is very implausible. One problem is that there are a number of counterexample to it. For instance, Peter Markie gives the following example.³⁸ I have a sensation as of a walnut tree. I have two seemings. First, that there is a walnut tree. Second, that the walnut tree was planted in April 24th, 1914. I form both beliefs. But clearly the second one of these beliefs is not an epistemically appropriate response, setting aside whatever reasons for doubt I might have. But there is also a deep theoretical problem. Seemings can be caused in all sorts of epistemically problematic ways. But this view ignores that fact. Thus, this view will have the result that, so long as one lacks a relevant reason to doubt, it is appropriate to form a belief as a result of a seeming even if that belief was formed by biases, wishful thinking, poor reasoning,

³⁴ Markie, "Epistemically Appropriate Perceptual Belief," 130-4.

³⁵ Markie, "Epistemically Appropriate Perceptual Belief," 134.

³⁶ William Tolhurst, "Seemings," *American Philosophical Quarterly* 35, 3 (1998): 293-302; Michael Huemer, "Compassionate Phenomenal Conservatism," *Philosophy and Phenomenological Research* 74, 1 (2007): 30-55; Andrew Cullison, "What Are Seemings?" *Ratio* 23, 3 (2010): 260-75; Blake McAllister, "Seemings as Sui Generis," *Synthese* 195, 7 (2018): 3079-96.

³⁷ Chris Tucker, "Why Open-minded People Should Endorse Dogmatism," *Philosophical Perspectives* 24 (2010) 529-545, 537-8.

³⁸ Peter Markie, "The Mystery of Direct Perceptual Justification," *Philosophical Studies* 126, 3 (2005): 347-73, 357.

poor education—not to mention brain lesions, evil geniuses and clairvoyant powers. And that is very implausible.³⁹

V. Valuing the Valuable

This final section presents an argument for Epistemic Value Pluralism that I believe succeeds. It does not turn on the particularities of epistemological theories but plausible general claims about value.

A. An Argument for Epistemic Value Pluralism

The argument contains the following premises. The first is this:

Iterated Appropriateness. If someone bears an appropriate attitude towards something of final value, then it is appropriate to bear a pro-attitude towards the fact that they bore an appropriate attitude towards something of final value.

This assumption is plausible on the face of it. Here is an additional reason for thinking it is true. Suppose a person bore a pro-attitude towards something of value—e.g., a friend is pleased that her son is happy. Now suppose a further person was aware of this pleasure but was either indifferent towards her attitude or even adopted a con-attitude towards it. We would normally think that such a person is behaving in way that is at least insensitive if not inappropriate. A natural explanation for this is that it is appropriate to value the fact that a person is adopting an appropriate attitude towards something of value.⁴⁰ Of course, in this situation, the

³⁹ The issues mentioned here mirror issues about the cognitive penetration objection to Phenomenal Conservatism. (I criticize Phenomenal Conservatism at greater length in Timothy Perrine, “Strong Internalism, Doxastic Involuntarism, and the Costs of Compatibilism,” *Synthese*, forthcoming.) But there are some differences. First, we are here considered with epistemically appropriate responses to experiences, not necessarily justified beliefs. Second, cognitive penetration occurs when a cognitive state directly impacts a perceptual state (Jack Lyons, “*Seemings and Justification*,” *Analysis Reviews* 75, 1 (2015): 153–64, 154) and my objection is not of that form. For additional critical discussion of the view in the text, as well as Phenomenal Conservatism, see Markie, “The Mystery of Direct Perceptual Justification” and “Epistemically Appropriate Perceptual Belief;” Jackson Alexander, “Appearances, Rationality, and Justified Belief,” *Philosophy and Phenomenological Research* 82, 3 (2011): 564–593; Susanna Siegel, “Cognitive Penetrability and Perceptual Justification,” *Nous* 46, 2 (2012): 201–22; Susanna Siegel, “The Epistemic Impact of the Etiology of Experience,” *Philosophical Studies* 162, 3 (2013): 697–722; Berit Brogaard, “Phenomenal Seemings and Sensible Dogmatism,” in *Seemings and Justification*, ed. Christ Tucker (New York: Oxford University Press, 2013); Matthew McGrath, “Phenomenal Conservatism and Cognitive Penetration: the “Bad Basis” Counterexamples,” in *Seemings and Justification*, ed. Tucker; and Lyons, “*Seemings and Justification*.”

⁴⁰ Of course, sometimes people have *excuses* for not valuing things—they are too busy, their

person is *aware* of the appropriate attitude. But even if someone is not aware of an appropriate attitude, it can still be appropriate for someone to bear a pro-attitude towards it. To use an analogy, it might be appropriate to praise a person for a very difficult basketball shot. (After all, the shot was difficult.) This might be appropriate *even if* no one is aware of it (besides, of course, the person who made the shot).

Earlier I assumed that when it is appropriate to adopt a pro-attitude towards something, then that thing is of final value. Given that assumption, *Iterated Appropriateness* implies:

Iterated Value: If someone bears an appropriate attitude towards something of final value, then the fact that they bore an appropriate attitude towards something of final value is, itself, of final value.

Like *Iterated Appropriateness*, *Iterated Value* is quite plausible. Several contemporary philosophers have adopted something close to it, though they usually add some qualifications and make additional claims about such a principle that are independent to our discussion.⁴¹ Now *Iterated Value* is formulated simply in terms of final value. But my immediate concern is final epistemic value. Since final epistemic value is a kind of final value, it is plausible that *Iterated Value* implies:

Iterated Epistemic Value: If someone bears an appropriate attitude towards something of final epistemic value, then the fact that they bore an appropriate attitude towards something of final epistemic value is, itself, of final epistemic value.

Here's a simple argument for Epistemic Value Pluralism. Suppose, as is very plausible, some true belief is of final epistemic value. Now suppose an agent bears a positive attitude towards the fact that someone has a true belief. By *Iterated Epistemic Value*, it follows that such a pro-attitude is of final epistemic value. Therefore, Epistemic Value Pluralism is true.

This simple argument contains one lacuna. Recall that a proper formulation of Epistemic Value Pluralism must hold that there are several things of *basic* final epistemic value. Merely maintaining that there are a number of ontologically distinct things of final epistemic value is not enough. So this simple argument needs to be shored up by maintaining that adopting a pro-attitude towards something of final epistemic value is, itself, of *basic* final epistemic value.

minds are elsewhere, etc. The existence of such excuses does not undermine the point.

⁴¹ See, e.g., Robert Nozick, *Philosophical Explanations* (Cambridge: Harvard University Press, 1981), 428ff.; Thomas Hurka, *Virtue, Vice, and Value* (New York: Oxford University Press, 1992), chapters 1&2; Zimmerman, *The Nature of Intrinsic Value*, chapter 6; Robert Adams, *A Theory of Virtue* (Oxford: Clarendon Press, 2007), chapter 2.

Here is a reason for thinking that adopting a pro-attitude towards something of final epistemic value is of basic final epistemic value. If it were merely of non-basic final value, then all of the value it has could be explained by appealing to the thing of final epistemic value that the attitude is directed to. But it cannot. Suppose some belief that *p* is of final epistemic value. Now suppose one agent adopts a pro-attitude towards it while another agent adopts a neutral attitude towards it. Given *Iterated Epistemic Value*, one of those attitudes is of final epistemic value, while (plausibly) the other is not. But both are about the same thing of final epistemic value. Thus, appealing to just what the attitude is about—its object—cannot explain the final epistemic value of adopting the pro-attitude. So adopting a pro-attitude towards something of final epistemic value is of basic final epistemic value. But some true beliefs are also of basic, final epistemic value. Thus, Epistemic Value Monism—and all versions of it, e.g., Truth Value Monism—are false. Epistemic Value Pluralism is true.

B. Epistemic Value and Value Simpliciter

In the remainder, I want to focus on what will be a surprising inference to some authors working on epistemic value: my inference of *Iterated Epistemic Value* from *Iterated Value*. That inference assumed that final epistemic value is a kind of final value. However, that assumption has been questioned by some philosophers. They deny that epistemic value is a kind of final value, or more weakly, that if something is of final epistemic value, then it is of value *simpliciter*. For instance, Ernest Sosa claims that there are various “domains” of evaluation, with the epistemic domain being just one among many. These domains admit of “value.” And, for each domain, some of that value is “fundamental” and others “derived” from the fundamental value of that domain. But none of this indicates that the fundamental value of a given domain is of final value *simpliciter*. Perhaps it is of some domain independent value, but it is not final value but (e.g.) instrumental value to some domain independent value. As Sosa once wrote, “Truth may or may not be intrinsically valuable absolutely, who knows? Our worry requires only that we consider truth the *epistemically fundamental* value, the ultimate explainer of other distinctively epistemic values.”⁴² Similar kinds of views have been endorsed by others. Duncan Pritchard likewise allows that something might be of “fundamental epistemic good” without that good being “finally valuable *simpliciter*.”⁴³ Pritchard even suggests that

⁴² Ernest Sosa, *A Virtue Epistemology* (Oxford: Clarendon Press, 2007), 72.

⁴³ Pritchard, “Knowledge and Understanding,” 12.

from the fact that truth is of epistemic value it need not follow that it has any value *simpliciter at all*⁴⁴

These kinds of views are inconsistent with my claim that final epistemic value is a kind of final value, or at least that when something is of final epistemic value that implies it is of some final value. However, this disagreement would not simply undermine my argument from *Iterated Value* to *Iterated Epistemic Value*; it is inconsistent with the basic way that I have setup the issues of this paper. For this reason, evaluating this kind of position is a large task that I cannot complete here. With that in mind, I raise two issues.

First, I assume that when something is of value it is valuable—that is, is *worthy* of value or it would otherwise be *appropriate* or *fitting* to value it. This view says that there is a kind of “value”—epistemic value—on which that is false. From the fact that something is of epistemic value it does *not* follow that it is worthy of value or that it would be appropriate to value it. (Maybe it is; maybe it isn’t.) To be sure, this view has a fallback position. If something is epistemically valuable, it may be *epistemically* appropriate to value it; or, from the epistemic point of view, it is worthy of valuing.⁴⁵ But this view denies that it follows from the fact that something is epistemically appropriate to value that it is also appropriate to value it *simpliciter*. At this point, I have begun to loose touch with what these words are supposed to mean. The problem is not the lack of a formal semantic device for this view.⁴⁶ The problem is more simply to understand what kind of thing deserves the title of value if it is not valuable!

Second, even if this way of thinking about epistemic value is inconsistent with mine, we might want to know why we ought adopt it. So far as I can see, the main argument is this. The inference from ‘x is a value in domain D’ so ‘x is a value *simpliciter*’ is invalid. (After all, there might be values in the “coffee domain” or even the “torture domain” that are fundamental to those domains. But we would not normally claim that their values are values *simpliciter*.⁴⁷ Thus, the inference from ‘x

⁴⁴ Duncan Pritchard, “The Ethics of Beliefs,” in *The Ethics of Belief*, eds. Jonathan Matheson and Rico Vitz (Oxford: Oxford University Press, 2014), 113. Something like this picture is also implicit in Luis R. G. Oliveira “Deontological Evidentialism, Wide-Scope, and Privileged Values,” *Philosophical Studies* 174, 2: 485-506, though generalized beyond issues of epistemic value.

⁴⁵ Cf. Pritchard, “The Ethics of Belief,” 113.

⁴⁶ For instance, one could utilize Geach’s distinction between “attributive” and “predicative” adjectives (or an analogous version for adverbs)—see P. T. Geach, “Good and Evil,” *Analysis* 17, 2 (1956): 33-42. Michael Ridge, “Getting Lost on the Road to Larissa,” *Nous* 47, 1 (2013): 181-201 is relevant here.

⁴⁷ See, e.g., Sosa, *A Virtue Epistemology*, 73-4; Ernest Sosa, *Knowing Full Well* (Princeton:

is a value in the epistemic domain' to 'x is a value *simpliciter*' is likewise invalid. To mimic the wording of Pritchard, it is a "further step" to say that if something is of final value in some domain that it is of final value *simpliciter*.

However, this argument is itself invalid. The following inference rule is certainly invalid: ϕ or ψ ; therefore, ϕ . After all, there are many instances of this inference pattern that do not preserve truth. But some instances of it *do* preserve truth. The inference 'A or A; therefore, A' is perfectly valid and is an instance of that inference. Similarly, there may be many domains that do not track value *simpliciter*. Those domains are merely ways of evaluating things. But from the fact that some domains do not track value *simpliciter* it does not follow that some *particular* domain also does not track value *simpliciter*. So this argument fails.

Of course, some might want to know why we should think that if something is of fundamental value in the epistemic domain that it follows that it is valuable *simpliciter*. A number of arguments could be developed here. One promising argument is through the similarity of epistemic evaluation and moral evaluation. Some authors have noticed that the kind of normativity, broadly construed, involved in moral and epistemic evaluations are very similar. But moral evaluations are widely thought to involve value *simpliciter*. If something is valuable in the moral domain it is valuable *simpliciter*. So too if something is valuable in the epistemic domain it is valuable *simpliciter*.⁴⁸

Additionally, there is a problem for epistemologists who deny that values in the epistemic domain are not necessarily values *simpliciter*.⁴⁹ Let us take seriously Sosa's suggestion that human beings are "zestfully judgmental across the gamut of

Princeton University Press, 2011) 63.

⁴⁸ For discussion, positive and critical, of this kind of reasoning, see Spencer Case, "From Epistemic to Moral Realism," *Journal of Moral Philosophy* 16, 5 (2019): 541-562; Terence Cuneo, *The Normative Web* (Oxford: Clarendon Press, 2007); Christopher Cowie, "Why Companions in Guilt Arguments Won't Work," *Philosophical Quarterly* 64, 256 (2014): 407-22; Christopher Cowie, "Good News for Moral Error Theorists," *Australasian Journal of Philosophy* 94, 1 (2016): 115-30; Ramon Das, "Why Companions in Guilt Arguments Still Work: Reply to Cowie," *Philosophical Quarterly* 66, 262 (2016): 152-160; Ramon Das, "Bad News for Moral Error Theorists," *Australasian Journal of Philosophy* 95, 1 (2017): 58-60; Richard Rowland, "Moral Error Theory and the Argument from Epistemic Reasons," *Journal of Ethics and Social Philosophy* 7, 1 (2013): 1-24; Richard Rowland, "Rescuing Companions in Guilt Arguments," *Philosophical Quarterly* 66, 262 (2016): 161-171.

⁴⁹ This problem is inspired by Stich's arguments against the value of true belief—see Stephen Stich, *The Fragmentation of Reason* (Cambridge: A Bradford book, 1990). But it is not quite the same. I discuss Stich's argument in Perrine, *Accurate Representation and Epistemic Value*, 246-253.

our experience.”⁵⁰ Of course, different “domains” of evaluations can issue different judgments of the same thing. Inside the “aesthetic” domain, a particular gourd might be bad because it is ugly; but inside the “culinary” point of view it might be excellent, ready for one’s fall soup. But there might be an “epistemic*” domain and in it true beliefs are not a fundamental epistemic* value; perhaps nothing is, or maybe only reasonable attitudes are. And perhaps there is an “epistemic**” domain where true beliefs are a fundamental epistemic** value, but knowledge is not more epistemically** valuable than true belief. Given such domains,⁵¹ a natural question is why should we give more attention to the epistemic domain than the epistemic* or epistemic** domain? This question is pressing for those who do not see the epistemic domain as tracking value *simpliciter*. But for those of us who see the epistemic domain as tracking value *simpliciter* there is a straightforward response. We should focus on the epistemic domain, instead of these competitors, because the epistemic domain tracks value *simpliciter* in a way that those other domains need not.

For these reasons, I think, this kind of response to my argument for Epistemic Value Pluralism is unpromising.⁵²

⁵⁰ Sosa, *A Virtue Epistemology*, 70.

⁵¹ If some domains do not already exist, we could easily introduce them. After all, prior to the cultivation and invention of coffee, there was not the “coffee domain” (or, there was, and it just wasn’t used).

⁵² After writing this paper, I stumbled upon Kurt Sylvan, “Veritism Unswamped,” *Mind* 127, 506 (2018): 381-435. In it, Sylvan argues that we can retain the idea that “accurate belief is the sole fundamental epistemic value” (Sylvan, “Veritism Unswamped,” 382) so long as we reject the view that “accurate belief is the sole non-instrumental epistemic value” (Sylvan, “Veritism Unswamped,” 382). (By ‘non-instrumental value’ Sylvan just means ‘final value’ (cf. Sylvan, “Veritism Unswamped,” 431.) Sylvan’s idea is that some things have a value that is derived from accurate belief in a way other than being instrumentally valuable. In defending his ideas, Sylvan even uses a principle of Hurka’s reformulated in terms of “derivative non-instrumental value” (Sylvan, “Veritism Unswamped,” 383). Comparing Sylvan’s argument to mine is difficult because he does not setup his discussion in terms of basic vs. non-basic final epistemic value. But there is a close affinity of our ideas. With that in mind, I’ll make three brief comments. First, if by ‘fundamental epistemic value’ Sylvan means “basic final epistemic value” then we disagree. For I’ve argued that there are more things of basic final epistemic value than true beliefs. Second, if Sylvan thinks things of “derivative non-instrumental value” are things whose value can be entirely explained by appealing to the value of true belief, then we disagree. For I’ve argued there are some things of final epistemic value whose value cannot be fully explained by appealing to true beliefs. Finally, though Sylvan and I both sympathetically cite authors like Hurka, Sylvan offers a thinner reading of them. Specifically, Sylvan maintains that there can be responses to final value that are themselves of final value while those responses do not require forming any pro-attitudes. I think that is implausible, but will not argue that here.

LUCK, KNOWLEDGE, AND EPISTEMIC PROBABILITY

Gregory STOUTENBURG

ABSTRACT: Epistemic probability theories of luck come in two versions. They are easiest to distinguish by the epistemic property they claim eliminates luck. One view says that the property is knowledge. The other view says that the property is being guaranteed by a subject's evidence. Asbjørn Steglich-Petersen defends the Knowledge Account (KA). He has recently argued that his view is preferable to my Epistemic Analysis of Luck (EAL), which defines luck in terms of evidential probability. In this paper, I defend EAL against Steglich-Petersen's arguments, clarify the view, and argue for the explanatory significance of EAL with respect to some core epistemological issues. My overall goal is to show that an epistemic probability account of luck rooted in the concepts of evidence and evidential support remains a viable and fruitful overall account of luck.

KEYWORDS: luck, epistemic probability account of luck, knowledge

1. Two Epistemic Probability Accounts of Luck

You stop for lunch at your favorite hot dog stand near your office. The owner says, "You come here all the time! Here's one for free."

By luck, you got a free hot dog. It is natural, as a first pass, to think of yourself as lucky to receive your free lunch because you had no idea that you would. You are surprised, you had no reason to expect this, you did not know it would happen. These general thoughts support an *epistemic* theory of luck: lucky events are those significant events that you lack good epistemic reasons to expect.

Here are two different ways of being slightly more precise about this case. One way: you did not *know* that by going to the stand today, you would get a free hot dog (and per hypothesis receiving the hot dog is significant to you). Another way: the *evidence* in your possession did not guarantee that you would receive a hot dog (and per hypothesis receiving the hot dog is significant to you). What both explanations have in common is that *epistemic* factors are used to explain why this event is a lucky one. What sets them apart?

Distinguishing those two superficially similar explanations requires taking positions on several further issues concerning the nature of knowledge, the nature of evidence, what it is to possess evidence, and how evidence can support a proposition sufficiently to 'guarantee' it. Both, though, are probability theories of

luck that identify *epistemic* probability as the kind of probability that matters for luck. There is no need to rehearse different categories of views of luck and pointing to the important similarities between the knowledge account and the evidential probability account, or in rehearsing how these two epistemic accounts of luck compare with non-epistemic accounts.¹ This paper is about how exactly to characterize the epistemic properties that factor into luck. This is a place for highlighting *differences*.

Steglich-Petersen defends this view, which he calls the “knowledge account.”²

KA: S is lucky with respect to E at t only if, just before t, S was not in a position to know that E would occur at t.

I have defended an alternative view, and called it the “epistemic analysis of luck.”³

EAL: S is lucky with respect to E only if the evidence S had immediately prior to E occurring did not guarantee that E would occur.⁴

The fundamental concept in KA is the concept of *being in a position to know* that an event will occur. In EAL, it is the concept of *a subject's evidence guaranteeing* that an event will occur. The difference between those two concepts is more important than one may think, even while the views have much in common. Steglich-Petersen has argued that KA is superior to EAL, as only KA can handle a certain class of cases. The chief failing of EAL, according to Steglich-Petersen, is that there are conditions on knowledge that make a difference for luck, and these conditions are neglected by EAL, with its focus on evidential probability as the sole epistemic condition relevant to luck.

This paper proceeds as follows. Section 2 evaluates an argument that EAL cannot account for lucky necessities and solves the problem by appealing to intuitive conditions for evidence possession that are not met in cases of lucky

¹ Cf. Fernando Broncano-Berrocal, “Luck,” *The Internet Encyclopedia of Philosophy*, <https://www.iep.utm.edu/>, 2020; Steven D. Hales, “Why Every Theory of Luck is Wrong,” *Nous* 50, no. 3 (2016): 490-508; Gregory Stoutenburg, “In Defense of an Epistemic Probability Account of Luck,” *Synthese* 196, 12 (2019): 5099-5113.

² Asbjørn Steglich-Petersen, “Luck as an Epistemic Notion,” *Synthese* 176, 3 (2010): 361-377, Asbjørn Steglich-Petersen, “Does Luck Exclude Knowledge or Certainty?” *Synthese* (2018), <https://doi.org/10.1007/s11229-018-1790-z>.

³ Gregory Stoutenburg, “The Epistemic Analysis of Luck,” *Episteme* 12, 3 (2015): 319-334, Stoutenburg, “In Defense.”

⁴ Steglich-Petersen calls this the ‘certainty’ account, but that label unfortunately leaves aside the concepts of evidence and evidential probability that are central to EAL. EAL also requires that the event is significant to the subject.

necessities. Section 3 argues that a higher-order evidence version of the lucky necessities case does not undermine EAL because higher-order evidence at least sometimes has lower-order import. Section 4 argues that even apparently non-probabilistic evidence can be understood probabilistically. Section 5 clarifies the role of evidence in EAL. Section 6 argues that EAL, and not KA, has the resources to provide an interesting, unified account of epistemic luck that explains many common intuitions about knowledge, including intuitions about closure principles, Gettier cases, and lotteries.

2. Logical Necessities, Knowledge, and Evidence

Steven Hales maintains that every theory of luck is open to counterexample.⁵ One of his counterexamples is this:

Logical Bandit: The logical bandit points a gun at you and tells you that unless you correctly answer a logic puzzle, he's going to steal your wallet. He gives you this poser:

Suppose you go to a diner where the cook is famous for pancakes. Actually, he is famous for burning 50% of the pancake-sides he cooks, and cooking the other 50% perfectly. The statistics: One third of his pancakes are golden on both sides; one third are black on both sides; and the remaining third are golden on one side and black on the other. You order a pancake. When it comes, the side you can see is golden. What is the chance that the other side is golden?

You are horrible at this sort of thing, and are completely flummoxed by the gun, the puzzle, and whole situation. You make a wild guess and say "it's 2/3." The logical bandit, who could tell you were just guessing, smiles ruefully and replies, "you're lucky the correct answer is indeed 2/3," and vanishes into the night.⁶

The point of the counterexample is that the answer 2/3 is necessary, but all modal and probability theories of luck say that an event is lucky only if it could have failed to occur (or could have failed to be true) or has a probability below 1.0.⁷ So, concludes Hales, all theories of luck, including probability theories, are false.

I have argued that EAL can handle the Logical Bandit:

You do not understand the puzzle, so the details of the puzzle that entail the correct answer are not included in your evidence. If the probability of the answer being 2/3 had been 1.0 on your evidence, you would have been in a position to know that 2/3 was the answer, and consequently not lucky to answer correctly.

⁵ Hales, "Why."

⁶ Hales, "Why," 495.

⁷ Intuitively, an event counts as *improbable* only if its probability is below 0.5. The requisite notion of improbability for luck is technical.

You clearly were not in a position to know the answer, since you had to guess. So, the probability condition is satisfied.⁸

Since the subject does not know that the answer is $2/3$, that implies that the evidential probability for the subject is below 1.0. Steglich-Petersen objects: “On the face of it, [the probability that $2/3$ is the answer] is 1, since it is entailed by the details of the puzzle.”⁹

It is important to separate the question of how probable it is that the answer to the puzzle is $2/3$ from the question of how probable it is *on the subject's evidence* that the answer to the puzzle is $2/3$. The questions are different. EAL says that the probability that determines luck is *evidential* probability, so the probability of the answer to the puzzle being $2/3$ does not matter for luck. What matters is whether the probability that the answer is $2/3$ is 1.0 *on the subject's evidence*. If and only if it is, then Steglich-Petersen is right that the evidential probability account fails against the case.

Steglich-Petersen says, “On the face of it, [the probability that $2/3$ is the answer] is 1, since it is entailed by the details of the puzzle.”¹⁰ The reasoning seems to be this: The Logical Bandit tells S the details of the puzzle, so S has the details of the puzzle as evidence, and those details yield a probability of 1.0 that the answer is $2/3$. For that reasoning to be sound requires the truth of this principle, or at least something very close to it:

Hearing Implies Having (Hearing): If S hears and understands that p, then p is included in S's total evidence.

Without (Hearing) or something much like it, there is no clear reason why we should think that the Bandit stating the details of the puzzle puts the details of the puzzle into the subject's evidence.¹¹

There are examples that support (Hearing). I ask the gate agent if it will be a full flight. The agent says “no.” It is reasonable to include *it is not a full flight* in my total evidence, in the set of propositions and experiences that I rely on. Ordinary cases of testimony are like this. But (Hearing) is not true. If the agent answers “no”

⁸ Stoutenburg, “In Defense,” 10.

⁹ “Knowledge or Certainty,” 6.

¹⁰ “Knowledge or Certainty,” 6.

¹¹ Similar principles could be invented for other sensory modalities, and similar concerns would arise. Does a brief memorial flash of an image in my mind put into my evidence facts about a distant family event, if the event was the original cause of the image? Does the faint smell of burning metal put into my evidence that I am in a motorcycle factory, if I am in a motorcycle factory and should be able to figure that out? The suggestion in each parallel case is that the threshold for what makes something in experience count as evidence is beyond mere exposure.

and I am immediately distracted by something else, to the point that the interaction is forever beyond retrieval for me, then the proposition *it is not a full flight* will *not* be included in my evidence. Hearing and understanding that *p* is insufficient for having *p* as evidence.

The relevant notion of “understanding” is critical here, too. Even without a full account of understanding, we need a relatively clear idea of what role understanding plays in turning an auditory sensation into evidence. A person who can repeat spoken sentences from another language does not count as understanding the language unless that person can also perform certain inferences like correctly processing alterations in word suffixes by number or gender and is able to answer at least some questions about what was spoken. Someone who can repeat the sentences but cannot do these other things is merely parroting words, without understanding.

The question now is, is the Logical Bandit case more like a straightforward instance of receiving testimony, or is it more like an instance of hearing something and immediately forgetting it, or hearing sentences from a foreign language? If the details of the puzzle are included in the subject’s evidence, and they are all the relevant evidence the subject has concerning the answer to the puzzle, then EAL would indeed imply that the probability for the subject that $2/3$ is the answer is 1.0.

I claimed that the hearer in the Bandit case does not understand the details of the puzzle, which implies that the subject does not have the details of the puzzle as evidence.¹² Steglich-Petersen objects that “there is no reason why he should not [understand the details of the puzzle]” and in support of the claim points out that the details of the puzzle are simple, before claiming that the subject merely struggles to draw the correct inference from the puzzle. Here, though, the subject makes a *complete* guess. It is not just that the subject fails to draw the correct inference, but that the subject fails to draw any inference *at all* and instead must resort to picking a number at random. That is a good reason to think that the subject does not understand the puzzle because the subject is in a position like that of a person who is hearing sounds from another language. There is of course a distinction between failing to infer and failing to understand, but in this case being able to competently think about how one might devise an answer to the question is constitutive of understanding the puzzle, just as having competence with a language is needed if hearing spoken words in that language puts the words spoken into one’s evidence.

¹² “In Defense,” 10.

So, it is reasonable to say that the subject indeed does not understand the puzzle. And thus, the details of the puzzle are not included in the subject's body of evidence from which the probability of the answer being $2/3$ can be derived.

3. Evidence and Higher-Order Defeat

Steglich-Petersen offers a variation on the Logical Bandit intended to show that if the subject *does* understand the puzzle, EAL gets the case wrong for another reason. The idea is that evidential certainty is incapable of eliminating luck, because it is possible for p to be certain on one's evidence while p is nevertheless lucky. The new case is a modification of the Logical Bandit, with a twist concerning higher-order defeat. The case starts the same as before, then continues:

Logical Bandit with Higher-Order Defeat: You are usually pretty good at logic puzzles, and after a bit of thinking, you reach the correct answer of $2/3$. However, the bandit then tells you that he slipped a powerful reason-distorting drug into your coffee just before pointing his gun at you. Reasonably, this convinces you that you have most likely made a mistake, even though in this instance you did not in fact make a mistake.¹³

Yet the subject guesses $2/3$ and is intuitively lucky that the answer is so. Steglich-Petersen says that the luckiness here cannot be explained as a matter of epistemic probability, because the first-order evidence that $2/3$ is the answer is the same now as it was in the first version of the case.

In arguing for his verdict, Steglich-Petersen says of the subject's guess: "*In his evidential situation*, this answer isn't much better than any other..."¹⁴ But that explanation plainly *supports* the EAL explanation of the case. On the subject's *total evidence* the probability that the answer is $2/3$ is indeed *under* 1.0, because in the new version of the case the subject's evidence includes the details of the puzzle, which taken alone would make the probability of $2/3$ equal to 1.0, *and* the additional evidence that the subject's reasoning is faulty, which *lowers* the probability that the answer is $2/3$. Higher-order evidence is evidence, so it affects epistemic probability.

An intuitive way of reaching the same conclusion is this. Take the subject immediately after the subject has arrived at the answer of $2/3$, and ask: "How likely is it that the answer is $2/3$?" Presumably, the subject will be very highly confident that the answer $2/3$ is true, and will perhaps say: "It is definitely $2/3$." Now, let the Bandit make the claim about the coffee, and ask again: "Now that you have heard

¹³ "Knowledge or Certainty," 7-8.

¹⁴ "Knowledge or Certainty," 8.

that the Bandit drugged you, how likely is it that the answer is 2/3?" Here, presumably the subject will give some answer that is lower than the initial answer. If the higher-order evidence is irrelevant to epistemic probability, the probability of the 2/3 answer has not gotten *even a little* lower than it was prior to being told that you have been drugged. It has stayed *exactly* the same.

There is precedent for thinking that higher-order beliefs and evidence at least sometimes defeat first-order reasons.¹⁵ When we have evidence concerning the quality of our evidence, we thereby have more evidence that bears on the probability of the proposition in question. There are a few ways that evidence concerning the quality of one's evidence may impact the probability of a proposition for a subject. One way is this: whenever a person, considering whether *p*, obtains evidence *that there is* evidence that bears on the truth of *p*, that indirect evidence is evidence concerning *p* for that subject.¹⁶ Anyone who thinks that how probable a proposition is depends upon the evidence for and against that proposition is thus committed to accepting that higher-order evidence bears on the *overall* probability of a proposition, and is not bracketed away as probability-neutral 'evidence'.¹⁷

Unless it is evidence that mysteriously has *no* bearing *at all* on the probability of the proposition in question, then higher-order concerns plainly *do* have an effect on the probability of the proposition for the subject. Surely we can, for theoretical purposes, distinguish different relationships that subjects' evidence may bear to propositions, and that is precisely what we see in the literature on higher-order evidence and defeaters. But bracketing for theoretical purposes does not imply commitment to the substantive claim that higher-order evidence and defeaters make absolutely *no* difference to epistemic probability.

Thus the Bandit's plausible claim to have distorted the subject's reasoning constitutes evidence for the subject that 2/3 is the wrong answer. The subject,

¹⁵ David Alexander, "The Problem of Respecting Higher-Order Doubt," *Philosophers' Imprint* 13, 18 (2013): 1-12; Richard Feldman, "Respecting the Evidence," *Philosophical Perspectives* 19, 1 (2005): 95-119; Sophie Horowitz, "Epistemic Akrasia," *Nous* 48, 4 (2014): 718-744; Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, eds. Richard Feldman and Ted A. Warfield (New York: Oxford University Press, 2010), 111-174.

¹⁶ Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers without God: Meditations on Atheism and the Secular Life*, ed. Louise Antony (New York: Oxford University Press, 2007), 194-214.

¹⁷ Arguably, all evidentialists (broadly speaking) should think of epistemic probability as 'flat': evidence for *p* is evidence for *p*, whether that evidence comes in directly for *p*, or in the form of evidence about one's evidence for *p*, or evidence about one's evidence for one's evidence for *p*, and so on. Thanks to Lisa Miracchi for discussion on this point.

Gregory Stoutenburg

believing in accordance with the evidence, revises the probability of 2/3 downward from 1.0. Consequently, the subject's belief that 2/3 is the answer is lucky.

4. Inference and Probability

Steglich-Petersen's final argument that evidential certainty is insufficient for luck begins with this example from Martin Smith:¹⁸

Background Color: Martin has set up his computer such that, whenever he turns it on, a random number generator determines the background color on his display. For one out of a million possible values, the background will be red. For the remaining 999,999, it will be blue. Martin turns on his computer, and then leaves the room before seeing the background color. A few minutes later, Martin's housemate Bruce enters the room and sees that the background color is blue.

Steglich-Petersen then applies the case to luck:

Suppose that both Martin and Bruce, in light of their respective evidence, form the belief that the background is blue. And indeed it is blue. Are any of them lucky to have formed a true belief? Bruce's true belief does not seem lucky. He is looking at the display, and sees that the background is blue. If it hadn't been blue, he wouldn't have believed it. Very plausibly, Bruce thus knows that it is blue... Martin, on the other hand, seems to enjoy at least some degree of luck in his true belief, even if it was statistically very unlikely to be false. Martin would not be justified in believing outright that the screen is blue.¹⁹ He does not know that it is blue, only that it is very likely to be blue. And if it hadn't been blue, he would still have believed it to be blue. So it does seem at least slightly lucky that he has ended up believing the truth. At the very least, the following seems clear: Martin is luckier than Bruce is, in ending up with a true belief.²⁰

Steglich-Petersen's point is that knowledge requires more than evidential probability of 1.0. Bruce's belief has these additional knowledge-making properties, whatever they are, while Martin's belief is based merely on probability. So, Bruce knows and Martin does not.²¹

¹⁸ Martin Smith, "What Else Justification Could Be," *Nous* 44, 1 (2010): 10-31; Martin Smith, *Between Probability and Certainty: What Justifies Belief* (New York: Oxford University Press, 2016).

¹⁹ Steglich-Petersen continues after the passage: "As Smith notes, in cases like this, it seems that while Bruce would be justified in believing that the background color is blue, Martin would not. Martin would of course be justified in believing it to be *very likely* that the color is blue, but not justified in outright believing this" (p. 9). Note that this diagnosis is correct only if one is justified in believing only what one has *infallible* reasons to believe.

²⁰ Steglich-Petersen, "Knowledge or Certainty," 10.

²¹ I am willing to entertain the example for the sake of argument, but it should not be forgotten that luck requires significance. It is rather implausible that the screen color is significant to

The same verdict can be explained probabilistically. It is stipulated in the case that Bruce *sees that* the screen is blue, while Martin infers the color from his knowledge of the color-generating algorithm. So, we can describe the conditional probability that the screen is blue for each subject this way:

For Bruce, $\Pr(\text{screen is blue} \mid \text{Bruce sees that it is blue}) = 1.0$

For Martin, $\Pr(\text{screen is blue} \mid \text{there is a } 999,999/1,000,000 \text{ chance that it is blue})$
 $= 0.999999$

Surely if one *sees that* p then the probability for that person that p is 1.0.²² If that is correct, then this case provides no reason to think that KA explains what EAL cannot, because *both* views secure the intuitive verdict: Martin is luckier than Bruce.

To press the point, one must supply a reason as to why we should think that the probability of Bruce's belief is lower than Martin's. Steglich-Petersen suggests that Bruce is unable to eliminate various possibilities of deception: that he is hallucinating, that the lighting is tricky, etc., and that such uneliminated possibilities affect the probability of Bruce's belief. But surely these possibilities *are* eliminated if *seeing-that* is factive. Surely if Bruce sees that the screen is blue then he knows that the screen is blue, and if he knows that the screen is blue, then it is impossible relative to everything Bruce knows that the screen is *not* blue (or that there is no screen, or that he is dreaming, and so forth). The objection still fails because the probability of Bruce's belief is higher than Martin's.

5. Luck and Evidence

Even if the previous objection fails—the objection was that more than evidential probability matters for luck—the strength of Steglich-Petersen's argument forces me to admit that my description of the EAL as claiming that luck is determined by significance and a subject's evidence *alone* is slightly misleading, as there really are factors that contribute to justification (and knowledge) that are in a way different

either of these subjects. One might say with Pritchard (according to his updated view that luck does not require significance) that *risk* is present even if luck is not: Duncan Pritchard, "The Modal Account of Luck," *Metaphilosophy* 45, 4-5 (2014): 594-619. Also see Fernando Broncano-Berrocal, "Luck as Risk," in *The Routledge Handbook of the Philosophy and Psychology of Luck*, ed. Ian M. Church and Robert J. Hartman (New York: Routledge, 2019).

²² In his defense of epistemological disjunctivism, Pritchard denies that seeing that p entails knowing that p . If that were correct (it is not), it would apply equally to Steglich-Petersen's diagnosis of this case and to the argument I just gave. Cf. Duncan Pritchard, *Epistemological Disjunctivism* (Oxford: Oxford University Press, 2012).

from just first-order *evidence*.²³ We can see this by comparing the role of additional factors in justification with the extra premise used to generate a vicious regress in “What the Tortoise Said to Achilles.”²⁴ Achilles attempts to infer C from

1. If A and B, then C
2. A and B

using modus ponens. The inference is plainly valid, but the Tortoise challenges him to *prove* that the inference is legitimate by adding a new premise that states the modus ponens rule:

3. ‘If [A and B, then C] and [A and B] then C’

But by granting the Tortoise’s request, Achilles has conceded enough for the Tortoise to show that Achilles is committed to a vicious regress. The requirement that from any set of premises that entail a conclusion one *also* needs a *premise* that describes the *rule* that permits the inference begins an infinite series that prevents the conclusion from ever being legitimately inferred.

We could, using the language of ‘evidence,’ say that Achilles’s evidence for C consisted of just two things: premises 1 and 2 of the argument. But then the Tortoise would rightly point out that Achilles has no business inferring C without at least some recognition of the *connection* between 1 and 2, and C. If awareness of that connection is counted as ‘evidence’ *in the same way* as belief in the premises 1 and 2, then the Tortoise’s vicious regress begins. The premises and the awareness of the relation between premises and conclusion are two different things, and it is misleading to call both of them “evidence” as though there is no difference between them.

With this in mind, we should revisit Steglich-Petersen’s argument that the Logical Bandit case cannot be shown to be lucky because the probability of the 2/3 answer is below 1.0 given the subject’s evidence. (Let us also grant what I have argued above is impossible: that the subject understands the details of the case while being unable to do more than guess blindly at the answer.) We *could* categorize the factors that contribute to the subject’s justification for believing that 2/3 is the answer this way:

The subject’s evidence = the details of the puzzle

The subject’s non-evidential justification-contributor = awareness that the

²³ I should have seen this problem coming, since I have used the same distinction to argue against a different thesis, in Bryan C. Appley and Gregory Stoutenburg, “Two New Objections to Explanationism,” *Synthese* 191, 7 (2016): 1391-1407.

²⁴ Lewis Carroll, “What the Tortoise Said to Achilles,” *Mind* IV, 14: 278-280.

evidence entails 2/3

If we describe the factors that contribute to the subject's justification in this way, then appeal to probability on the subject's evidence does not solve the Bandit case. We would then be conceding that the subject's awareness of the connection between the evidence and the proposition has no bearing on epistemic probability.

However, while there is an important difference between evidential and non-evidential justifiers, there is a much more substantial difference between contributors to a subject's justification that have to do with the *truth* of the proposition believed, and those that do not, like a subject's level of confidence or the stakes at play in a context. Indeed, I introduced talk of evidential probability only when distinguishing my epistemic probability account from the *interest-relative* version of Steglich-Petersen's view.²⁵ If we draw the distinction between what counts as evidence and what does not count as evidence by factors that contribute to the truth of a proposition and those that do not, then for the purpose of talking about how probable a proposition is for a subject, it is reasonable to collapse the evidential/non-evidential justifier distinction as it was drawn above. Reasonable, but still misleading.

So, if we use 'evidence' broadly to include all and only those factors that bear on the truth of a proposition, factors like a subject's experiential states, knowledge, justified beliefs, and awareness of evidential connections, we can say that on this *evidential* probability theory of luck the probability that the answer is 2/3 is indeed below 1.0, since the subject fails to appreciate the connection between those factors that support the 2/3 answer and that answer being correct (even granting, once again, that the subject understands the puzzle).²⁶ Since the subject's awareness of the connection between the answer and the other factors that support that answer is one of the factors that supports the truth of the 2/3 answer, the overall probability for the subject is below 1.0.

6. Epistemic Accounts of Luck and Epistemic Luck

A long-standing hope of the literature on luck is that an analysis of luck will provide insight into some lasting philosophical problems, like how we can have

²⁵ Stoutenburg, "In Defense," 6.

²⁶ Eventually it will be harder to be this inclusive about evidence, as philosophers disagree on what constitutes knowledge and justification. Safety theorists, for instance, have a modal condition on knowledge. So, if 'S knows that p' entails that p is safe, and p is a part of S's evidence, then a modal fact constitutes some of S's evidence. This is an example of how specific accounts have implications for more general ones. Philosophers interested in an epistemic probability account of luck can take sides on this issue as they wish.

moral responsibility when so much of what we are and do is lucky, and why we cannot have knowledge when the truth of a belief is due to luck.²⁷ An epistemic account of luck is at least initially appealing in this regard: we have the intuition that a belief that is true by luck is not knowledge, and then we add that by ‘true by luck’ we mean that the subject’s belief suffers a distinctively *epistemic* failing.

The initial appearance is deceiving, however, if by “B is true by luck” we just mean “B is not knowledge,” as KA would have it. Then the account is saying what we knew all along: that when a belief is not knowledge, it is not knowledge. This trivial equivalence is all the insight that a knowledge account of luck can give us. Whatever other virtues such an account could have, this consequence is disappointing, and prevents KA from having much theoretical significance.²⁸

In contrast, an epistemic account of luck understood in terms of epistemic probability can identify an interesting and principled connection between luck and knowledge. We can say that a belief is true by luck when the epistemic probability of the belief falls below a specific threshold, most plausibly (in my view) the limit of epistemic probability: 1.0. Any significant true belief is at least *very slightly* lucky if the probability of the belief on the subject’s evidence is below 1.0.

This way of thinking about *epistemic* luck illuminates some connections between epistemic probability and knowledge. If the view is true, then underdetermination arguments for skepticism, lottery problems, closure-based skeptical arguments, and Gettier scenarios all have in common that in each, the truth of a subject’s belief is to a slight degree due to luck, because the subject’s basis for belief does not guarantee the truth of the proposition believed. It is significant that an epistemic probability theory of luck can unite a number of interestingly similar issues about knowledge.

One might worry that this account of epistemic luck immediately leads to skepticism. Here is one defective argument for that conclusion:

1. S cannot know that p if the truth of p is just a matter of luck.
2. The truth of p is just a matter of luck when $P(p) < 1.0$.
3. So, we cannot know that p.

For the argument to be at all forceful would also require the stipulation that most ordinary beliefs are like p.

²⁷ See Ballantyne for arguments that this goal of the luck project is misplaced: Nathan Ballantyne, “Does Luck Have a Place in Epistemology?” *Synthese* 191, 7 (2014): 1391–1407.

²⁸ Steglich-Petersen is aware of this kind of limitation of his view (“Luck,” 376). Here, I am arguing that the limitation has implications for the significance of the view.

The argument is defective because premise 2 is false. It is highly misleading to say that the probability of p is “just a matter of luck” when p is *highly* probable but still below 1.0. To say an event is “just a matter of luck” implies that the occurrence of the event is due *only* to luck and to nothing else. But that is not what EAL says, even if combined with the idea that a belief is true by luck when it is not guaranteed by the subject’s evidence: it may be that the subject’s belief is 0.99 probable, and true. That would not make the truth of the belief “just a matter of luck.” It would make the belief .01 a matter of luck, and .99 a matter of evidential support.

Premise 1 is not implied by the epistemic probability account of luck. Nothing in that account of luck implies that we must accept that knowledge is incompatible with even the slightest trace of luck. Even those who think that luck and knowledge are incompatible, broadly speaking, should still clarify *how much* luck is tolerable. Fallibilists might accept the epistemic probability account of luck and say that knowledge is compatible with some small degree of luck, and infallibilists can deny this. Either way, the truth of premise 1 is required for the argument to be successful, and nothing about the epistemic probability theory of luck itself implies that premise 1 is true. But if one thinks that absolutely no luck is compatible with knowledge, then the skeptical conclusion follows.

7. Conclusion

An epistemic probability account of luck successfully explains how beliefs about necessary truths can be lucky and delivers the correct verdict about luck when considering the epistemic status of beliefs formed by vision and through inference. While arguing for these claims in the course of responding to objections from Steglich-Petersen, I also further developed my epistemic probability theory of luck. I argued that higher-order evidence affects epistemic probability, clarified the conditions under which an experience is taken up as evidence, distinguished ways that factors relevant to justification affect epistemic probability, argued that the epistemic analysis of luck does not by itself imply skepticism, and underscored the impressive explanatory power of the epistemic analysis of luck. EAL remains a fruitful and viable option for a unified theory of luck.

DISCUSSION NOTES/ DEBATE

GETTIER BELIEFS AND SERIOUS BELIEFS: A REPLY TO BIRO AND FORRAI

James SIMPSON

ABSTRACT: In a recent exchange in the pages of this journal, John Biro responds to Gabor Forrai's argument against Biro's argument that in most, if not all, Gettier cases the belief condition, *contra* popular opinion, isn't satisfied. In this note, I'll argue that Biro's response to Forrai satisfactorily resolves the first of Forrai's two central objections to Biro's argument that the belief condition isn't satisfied in most, if not all, Gettier cases. But Biro's response leaves mostly unaddressed the most plausible way of construing Forrai's second objection. I'll take up the mantle of successfully defending Biro's argument from this more plausible construal of Forrai's second objection. However, even though I'll argue that Biro's argument is in good shape with respect to Forrai's objections, I'll show that the definition of serious belief that Biro offers us is mistaken.

KEYWORDS: John Biro, Gabor Forrai, Gettier case, serious belief, assert, action

In a recent exchange in the pages of this journal, John Biro¹ responds to Gabor Forrai's² argument against Biro's argument³ that in most, if not all, Gettier cases the belief condition, *contra* popular opinion, isn't satisfied.

To, briefly, recapitulate this exchange. In his paper "Non-Pickwickian Belief and 'the Gettier Problem'," Biro argues, at least in part, that the belief that subjects allegedly have in most, if not all, Gettier cases is a belief in merely a pickwickian sense. In other words, it's not a serious belief, where, for Biro, a serious belief is one that, *inter alia*, guides action⁴ and adequately reflects preparedness or willingness to assert the proposition believed.⁵ And, Biro thinks, it's a serious, non-

¹ John Biro, "No Reprieve for Gettier "Beliefs": A Reply to Forrai," *Logos & Episteme* X, 3 (2019): 327-331.

² Gabor Forrai, "Gettiered Beliefs are Genuine Beliefs: A Reply to Gaultier and Biro," *Logos & Episteme* X, 2 (2019): 217-224.

³ In John Biro, "Non-Pickwickian Belief and 'the Gettier Problem'," *Logos & Episteme* VIII, 3 (2017): 47-69.

⁴ In this paper, by "action" I mean "action or omission."

⁵ Biro, "Non-Pickwickian Belief and 'the Gettier Problem'," 53, 68.

pickwickian belief that the Gettier case subject must have in order to count as satisfying the belief condition for knowing.⁶

Forrai, on other hand, resists Biro's argument on dual grounds. First, Forrai argues—somewhat awkwardly—that there could be, *contra* Biro, a circumstance where some epistemic subject, S, counts as seriously believing that p, even though, S isn't prepared to assert that p.⁷ Second, Forrai argues that some actions aren't guided by single beliefs, but rather by “constellations of beliefs.”⁸ Take, for instance, the belief that Rod's couch is comfy. Even though I believe this, I might not sit on Rod's comfy couch, since I also believe that Rod's overly sensitive about people sitting on his comfy couch. Forrai's point is that what guides my omission of not sitting on Rod's couch is not my belief that Rod's couch is comfy, but my belief that Rod's overly sensitive about people sitting on his comfy couch. Hence, Forrai argues, there will be at least some serious beliefs that can't be said to guide action, at least not in a certain sense. Even still, Forrai thinks, those beliefs are serious beliefs.

The above reconstruction of Forrai's general argument represents some reading between the lines on my part. This is mostly the result of some unclarity on Forrai's part on how his criticisms of Biro actually connect up with Biro's view in a genuinely problematic way. Perhaps, this unclarity still persists. To see both of the above lines of argument more clearly, then, let's consider the following two cases adapted from Forrai:⁹

NoSay. Suppose Jim wants to buy a used Ford and he believes that Havit's Ford is up for sale. Jim, however, was told that if he asserted that Havit's Ford is up for sale, then he wouldn't be able to buy it.

NoSale. Suppose Greg wants to buy a used Ford and he believes that Havit's Ford is up for sale. It would then be perfectly rational for Greg to talk to Havit about buying it. However, if Greg also believes that Havit would not sell him his car because he hates Greg's guts, Greg would not talk to Havit about buying his Ford.

In NoSay, Jim wouldn't be willing to assert that Havit's Ford is up for sale, although, intuitively, Jim seriously believes that Havit's Ford is up for sale. If being willing to assert that p is a necessary condition for seriously believing that p, as Biro alleges, then Jim doesn't count as seriously believing that Havit's Ford is up for sale.¹⁰ But, of course, that's the intuitively incorrect result.

⁶ Biro, “Non-Pickwickian Belief and ‘the Gettier Problem’,” 59.

⁷ Forrai, “Gettiered Beliefs are Genuine Beliefs: A Reply to Gaultier and Biro,” 221.

⁸ *Ibid.*, 222.

⁹ *Ibid.*

¹⁰ Biro, “No Reprieve for Gettier “Beliefs”: A Reply to Forrai,” 328.

In NoSale, Greg believes that Havit's Ford is up for sale and he wants to buy a used Ford, but Greg actively avoids talking with Havit about buying his Ford. If guiding action, in Forrai's sense, is a necessary condition for seriously believing that *p*, as one might read Biro as suggesting, then Greg doesn't count as seriously believing that Havit's Ford is up for sale.¹¹ The reason is that the belief that Havit's Ford is up for sale isn't what guides Greg's omission of not talking with Havit about purchasing the Ford. That would be Greg's belief that Havit won't sell him the Ford because he hates his guts. Yet, intuitively, Greg seriously believes that Havit's Ford is up for sale, even though, that belief doesn't guide Greg's omission of not talking with Havit about purchasing his Ford.

In the most recent paper in this exchange, Biro appears to rather neatly navigate both of Forrai's worries by adopting a *ceteris paribus* clause. Roughly, S counts as seriously believing that *p* only if, *ceteris paribus*, both S is prepared to assert that *p* and S's belief that *p* guides S's action. Let's call this Biro's Principle.

Applying Biro's Principle to NoSay, we see that, all things considered, Jim wouldn't be prepared to assert that Havit's Ford is up for sale, but, other things equal, he would be. For Biro, that's all being prepared or willing to assert that *p* comes to. This, then, swiftly resolves Forrai's first worry.

Applying Biro's Principle to NoSale, we see that, all things considered, Greg wouldn't talk to Havit about purchasing his Ford, but, other things equal, Greg would. Again, for Biro, all belief guiding action comes to is that, *ceteris paribus*, S's belief that *p* guide action. This, then, apparently nicely resolves Forrai's second worry.

Yet such resolution of Forrai's second worry is only apparent. The reason is that Biro overlooks how Forrai intends to characterize what it is for belief to guide action. Forrai appears to think that S's belief that *p* guides some action, *A*, only if the reason why S *A*-ed was her belief that *p*.¹² In which case, adding a *ceteris paribus* clause is of little help with NoSale, since, other things being equal, Greg's belief that Havit's Ford is up for sale won't guide, in Forrai's sense, Greg's omission of not talking to Havit about purchasing his Ford. Thus, on Forrai's characterization of belief guiding action, Biro's Principle doesn't yield the correct result in NoSale that Greg's belief that Havit's Ford is up for sale is a serious belief.

However, I think there's fairly simple response available to Biro here. To see this, consider the following case:

Book. Suppose Tim believes that a copy of *Waverly* is in his office and he wants it. There are three routes Tim can take turning out of his driveway. One route goes

¹¹ Biro, "Non-Pickwickian Belief and 'the Gettier Problem'," 68.

¹² Forrai, "Gettiered Beliefs are Genuine Beliefs: A Reply to Gaultier and Biro," 222.

to Tim's office and passes a Publix. One route doesn't pass a Publix, but it goes to Tim's office. One route goes the opposite direction of Tim's office, but it passes a Publix. Both Publix's are the same distance from Tim's driveway. Tim believes that he needs some milk for the house, so he decides to take the route that goes to his office and passes a Publix. On the way to his office, Tim stops at Publix and picks up some milk.

Now, Forrai's conception of belief guiding action yields the result that Tim's belief that a copy of *Waverly* is in his office doesn't guide his action of picking up milk from Publix, since Tim's belief that a copy of *Waverly* is in his office isn't the reason why Tim picks up milk from Publix. Indeed, the reason why Tim picks up milk from Publix is that he believes that he needs milk at the house.

Yet, intuitively, it is wrong to think that Tim's belief that a copy of *Waverly* is in his office doesn't guide his action of picking up milk from Publix. One way of sustaining this intuition is by noting that Tim chooses the route to his office that he does, at least to some obvious extent, on the basis of his belief that a copy of *Waverly* is in his office. If he didn't have that belief, then he might just as well rationally choose the route that passes Publix but doesn't go to his office. In Book, however, if he did that, he would be acting irrationally. But, on Forrai's view, Tim picking up milk from Publix by taking either route that passes Publix would be equally rational, since, on Forrai's view, the only belief that guides Tim's action is his belief that he needs milk at the house. Yet, certainly, it wouldn't be rational for Tim to take the route that passes Publix, but goes in the opposite direction of his office, given that he believes a copy of *Waverly* is in his office. This signals to me, then, that, quite plausibly, Forrai's conception of what it is for belief to guide action is mistaken.

But, of course, this invites a question: What is it for a belief to guide action? In this connection, I propose the following view of what it is for belief to guide action:

For S's belief that p to guide S's action in circumstance, C, S's belief that p must inform whatever S, in fact, does in C.¹³

To see the view more clearly, let's consider an example. Suppose I want a beer and I believe, seriously, that my fridge is empty of beer. That belief is rationally consistent with my doing all sorts of things, like going to the store to get beer, calling my wife to pick some beer up from the store, not getting any beer at all, and so on. But what the belief that my fridge is empty of beer is not rationally

¹³ Note, on such a view, for S to have a serious belief that p in C, S's belief that p needn't be the reason why S performs whatever action he performs in C.

consistent with is my going to the fridge to get a beer or telling my pal, Julie, that there's a beer in the fridge. In other words, my belief that there aren't any beers in the fridge guides the action that I do perform insofar as it "tells" me that some doings *are* and some doings *aren't* available to me as a rational agent.

Now let's consider NoSale once more. Greg's belief that Havit's Ford is up for sale can be said to guide Greg's omission of not talking to Havit about purchasing his Ford only if it "communicates" the space of doings that are and aren't available to him as a rational agent. And, to my mind, that's exactly one way, *inter alia*, that Greg's belief that Havit's Ford is up for sale functions in NoSale.

Interestingly, this line of reasoning appears to be consonant with the way that Biro thinks about how serious belief guides action. As Biro comments, a mark of serious belief "is that it guides action. If I am in the market for a used Ford and believe that Havit owns the one in the parking lot, it would *not* be *rational* for me to go around asking who owns it" (emphasis mine).¹⁴ In other words, a serious belief guides action only if it makes clear what doings fall inside and outside the scope of rational doings. Thus, as far as I can see then, Biro's view of serious belief is in good shape with respect to both of Forrai's worries.

However, I'll close by showing that Biro's view isn't free from danger completely. Biro¹⁵ offers us the following definition of serious, non-pickwickian belief:

Biro's Definition. For any set of propositions such that one knows that one of them follows from the others but could be true even if those others were not, one believes the entailed proposition if and only if one would believe it even if one did not believe (all) the entailing ones.

Biro's Definition is intended to give us a general account for what makes a belief serious. But does Biro's Definition apply to all beliefs or only Gettier beliefs? Realistically, there are two answers available to Biro here, but neither answer seems very satisfying. One is that for *any* belief to be serious it must satisfy Biro's Definition. The other is that *only* Gettier beliefs—the beliefs at issue in Gettier cases—must satisfy Biro's Definition in order to count as serious beliefs.¹⁶

The former answer appears to render Biro's Definition false. To see this, consider the following scenario:

Raven. Suppose I believe, falsely, that every bird is a raven. I know that the false proposition, every bird is a raven, entails the true proposition that the bird on my front porch is a raven, and so, on this basis alone, I come to believe, truly, that the

¹⁴ Biro, "Non-Pickwickian Belief and 'the Gettier Problem'," 68.

¹⁵ In Biro, "No Reprieve for Gettier 'Beliefs': A Reply to Forrai," 330.

¹⁶ It's worth noting that this is Biro's (conversation) preferred answer.

James Simpson

bird on my front porch is a raven. Let's call the bird on my front porch, Bird₁. Suppose I think to myself, "Wow. In seeing Bird₁, I've finally seen a raven." Further suppose that I tell my wife and kids that Bird₁ is a raven. I call the local aviary and tell them that I've found a raven, namely, Bird₁. I even try, unsuccessfully, to have Bird₁ become a member of an unkindness (flock) of ravens.

Now, Biro's Definition yields the result that I don't believe, at least seriously, that Bird₁ is a raven, since if I didn't believe that every bird is a raven, then I wouldn't believe that Bird₁ is a raven. Yet, intuitively, I seriously believe that Bird₁ is a raven. After all, I clearly think of myself as believing that Bird₁ is a raven, I tell my wife and kids that Bird₁ is a raven, I call the local aviary to tell them that Bird₁ is a raven, and I even try to have Bird₁ become a member of an unkindness of ravens. Quite plausibly, then, I seriously believe that Bird₁ is a raven. If this interpretation of Raven is correct, as seems eminently reasonable, then, on the former answer, Biro's Definition must be mistaken.

The latter answer, while it avoids the problem above, strikes me as *ad hoc*. On its face, it appears that placing conditions of seriousness of belief on Gettier beliefs, but not ordinary beliefs, would simply be a way for Biro to get his desired result that the belief condition isn't satisfied in most Gettier cases and, yet, avoid the problem above. Beyond this, though, it's not clear what could motivate placing conditions of seriousness of belief on just the beliefs at issue in Gettier cases.

NOTES ON THE CONTRIBUTORS

Frederik J. Andersen is a PhD student at the Arché Philosophical Research Centre for Logic, Language, Metaphysics, and Epistemology, University of St Andrews. His main research interests are in epistemology and logic. Recently, he published the paper “Moral Disagreement and Higher-Order Evidence” with Klemens Kappel in *Ethical Theory and Moral Practice*. Contact: frederikjandersen@gmail.com.

Daniele Bertini is adjunct professor of Philosophy at the University of Rome 2, Tor Vergata. While he specialized in analytic philosophy of religion, he is competent in overlapping disciplines as history of religious doctrines (especially Christian Theology), philosophy of language, and epistemology. His works focus mainly on the epistemology of religious belief and the challenge of religious diversity. He authored four books (on George Berkeley, the History of Christian Theology, and the problem of Religious Diversity), three collections (on the role of experience in Theology, on moral heteronomy, and on relations), and published in journals as *Dialegethai*, *International Journal of Philosophy and Theology*, *Religions*, *European Journal for Philosophy of Religion*, and *Philosophia*. Contact: daniele.bertini@uniroma2.it.

Nicole Dular is Assistant Professor of Philosophy at Franklin College. She received her PhD from Syracuse University in 2017. Her research interests are broadly on issues at the intersection of ethics, epistemology, and feminist philosophy. The focus of her most current research is on the ethics and normativity of social and epistemic practices, from a non-ideal approach. Website: nicoledular.weebly.com. Contact: nicoledular@gmail.com.

Kok Yong Lee is an associate professor in the department of philosophy at National Chung Cheng University, Minhsiung, Taiwan (R.O.C.). His research focuses on epistemology and philosophy of language. Contact: kokyonglee.mu@gmail.com.

Timothy Perrine is currently a Post-Doctoral Researcher at Wuhan University in Wuhan China. He has published over a dozen and a half papers in various areas including epistemology, metaphysics, ethics, and philosophy of religion. Contact: timperrine@gmail.com.

James Simpson is a Ph.D. student at the University of Florida. His interests are in epistemology, ethics, metaphysics, the philosophy of action, and the philosophy of religion. His most pressing current research project involves a dissertation-length defense of a novel analysis of knowledge. Contact: simpson.james@ufl.edu.

Logos & Episteme

Gregory Stoutenburg is a Visiting Assistant Professor of Philosophy at York College of Pennsylvania. His research focuses on knowledge and justified belief, with a particular interest in demanding standards for knowledge. He has published on the nature of luck in defense of the view that luck is partly an epistemic matter. Contact: gstoutenburg@ycp.edu.

***LOGOS & EPISTEME*: AIMS & SCOPE**

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

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

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

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

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

For further information, please see the Notes to Contributors.

Contact: logosandepisteme@yahoo.com.

NOTES TO CONTRIBUTORS

1. Accepted Submissions

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

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

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

2. Publication Ethics

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

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

3. Paper Size

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

4. Manuscript Format

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

5. Official Languages

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

6. Abstract

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

7. Author's CV

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

8. Review Process

The reason for these requests is that all articles which pass the editorial review, with the exception of articles from the invited contributors, will be subject to a strict double anonymous-review process. Therefore the authors should avoid in

their manuscripts any mention to their previous work or use an impersonal or neutral form when referring to it.

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

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

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

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

9. Acceptance of the Papers

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

10. Responsibilities

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

11. Checking Proofs

Authors should retain a copy of their paper against which to check proofs. The final proofs will be sent to the corresponding author in PDF format. The author must send an answer within 3 days. Only minor corrections are accepted and should be sent in a separate file as an e-mail attachment.

12. Reviews

Authors who wish to have their books reviewed in the journal should send them at the following address: Institutul de Cercetări Economice și Sociale „Gh. Zane” Academia Română, Filiala Iași, Str. Teodor Codrescu, Nr. 2, 700481, Iași, România.

The authors of the books are asked to give a valid e-mail address where they will be notified concerning the publishing of a review of their book in our journal. The editors do not guarantee that all the books sent will be reviewed in the journal. The books sent for reviews will not be returned.

13. Property & Royalties

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

14. Permissions

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

15. Electronic Archives

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