

Epistemic Relativism and Reasonable Disagreement

[Version of June, 2007]

Alvin I. Goldman
Rutgers University

[To appear in: R. Feldman and T. Warfield, eds., Disagreement; OUP]

1. Introduction

Two active topics in current epistemology are epistemic relativism and the reasonableness of disagreement between equivalently positioned agents. These topics are usually treated separately, but I'll discuss them in tandem because I wish to advance a new conception of relativism that bears on the issue of reasonable disagreement.

I begin with some familiar conceptions of epistemic relativism. One kind of epistemic relativism is descriptive pluralism. This is the simple, non-normative thesis that many different communities, cultures, social networks, etc. endorse different epistemic systems (E-systems), i.e., different sets of norms, standards, or principles for forming beliefs and other doxastic states. Communities try to guide or regulate their members' credence-forming habits in a variety of different, i.e., incompatible, ways. Although there may be considerable overlap across cultures in certain types of epistemic norms (e.g., norms for perceptual belief), there are sharp differences across groups in other types of epistemic norms.

What about the normative status of these different E-systems? Is one of them right and the rest of them wrong from an objective or absolute point of view? Are some "more right" than others? Descriptive pluralism takes no stand on this issue, but epistemologists generally want this normative issue resolved. A second brand of relativism, nihilistic relativism, does take a stand. It augments descriptive pluralism in holding that there is no objective right or wrong in this matter. As the sociologists of knowledge Barry Barnes and David Bloor (1982) express the matter:

For the relativist there is no sense attached to the idea that some standards or beliefs are really rational as distinct from merely locally accepted as such. [The relativist] thinks that there are no context-free or super-cultural norms of rationality ... (1982: 27)

Philosophers would express the view by saying that there is no fact of the matter about which community is ("objectively", or "absolutely") right.

Standing in contrast with nihilistic relativism is epistemic objectivism. This view holds that there is objective rightness in matters of epistemic norms, standards, or principles. Epistemic objectivists characteristically hold that there is a uniquely correct E-system and all systems incompatible with this one are wrong. Alternatively, an objectivist might hold that E-systems can be ordered by the (objective) binary relation of being at least as correct as. Such an ordering does not entail the existence of a uniquely

correct E-system, because two or more non-equivalent systems might tie for the most correct. Also, there might be infinitely many systems for each of which there is a more correct one. For purposes of the present paper, however, I'll make the simplifying assumption that if objectivism is true, there is a uniquely correct E-system.¹

I shall propose a new form of relativism that strikes a compromise between nihilistic relativism and objectivism. I call it objectivity-based relativism. As its name suggests, this form of relativism presupposes the truth of epistemic objectivism; there is nothing nihilistic about it. Nonetheless, it manages to preserve some of the pluralism associated with relativism. Moreover, I'll argue that objectivity-based relativism allows the possibility that two people can reasonably disagree about a given proposition even when they have equivalent evidence vis-à-vis that proposition.

I assume a close link between epistemic objectivism and the status of being (objectively) justified or unjustified with respect to beliefs and other doxastic states. The idea is that a belief or another doxastic state is justified or unjustified so long as it conforms or fails to conform to what is prescribed by the correct E-system, given the subject's evidence. Since mainstream epistemologists generally assume that beliefs are objectively justified or unjustified, the truth or falsity of epistemic objectivism is a critical issue. The truth of objectivism is also important to the prospects for reasonable disagreement, because such prospects depend on whether two evidentially equivalent people can each be objectively justified if they hold conflicting doxastic attitudes. Objectivism, understood as entailing a uniquely correct E-system, seems to imply the impossibility of reasonable (i.e., justified) disagreement.

To further situate the discussion, I move to two recent pieces of epistemology: Richard Feldman's paper "Reasonable Religious Disagreement" (forthcoming) and a chapter on epistemic relativism from Paul Boghossian's book, Fear of Knowledge (2006).

One of the two main questions Feldman poses is the following:

(Q1) Can epistemic peers who have shared their evidence have reasonable disagreements?

To say that two people have a disagreement, according to Feldman, is to say that one of them believes a certain proposition and the second disbelieves it. To say that two people have a reasonable disagreement is to say that each is justified in holding his or her belief (or disbelief). To say that people are epistemic peers is to say that they are roughly equal with respect to intelligence, reasoning powers, background information, etc. People have shared their evidence about a topic when they have had a full discussion of the topic and have not withheld relevant information.

Let us slightly amend Feldman's formulation. Instead of confining disagreement to cases of one person believing a proposition and another disbelieving it, let the term "disagreement" apply to any case of two people holding contrary, or incompatible, credal

attitudes toward the same proposition. This includes one person believing the proposition and the other suspending judgment. And instead of restricting the range of doxastic attitudes to the tripartite categories of belief, disbelief, and withholding, let us include graded beliefs or subjective probabilities among the set of categories, either point probabilities or interval probabilities (i.e., partial beliefs that are somewhat fuzzy). Finally, assume that not only belief but any doxastic attitude can exemplify the property of being justified or unjustified.

The core of Feldman's paper is his defense of a negative answer to (Q1). A crucial element in this defense is what he calls "The Uniqueness Thesis":

This is the idea that a body of evidence justifies ... at most one attitude toward any particular proposition. As I think of things, our options with respect to any proposition are believing, disbelieving, and suspending judgment. The Uniqueness Thesis says that, given a body of evidence, one of these attitudes is the rationally justified one (typescript, p. 10).

Feldman's appeal to the Uniqueness Thesis leads directly to the question of nihilistic relativism's viability, because if nihilistic relativism is true there is no uniquely correct system of epistemic norms. And if there is no uniquely correct system of norms, there is no guarantee that the Uniqueness Thesis is correct. Perhaps two or more different systems of epistemic norms are equally legitimate. One implies that a given body of evidence makes doxastic attitude D vis-à-vis proposition P rationally justified whereas another implies that the same body of evidence makes an incompatible attitude D* vis-à-vis P rationally justified. Thus, if the Uniqueness Thesis is false, Feldman's argument for the impossibility of reasonable disagreement falls through.

However, let us examine relativism more circumspectly, looking at Boghossian's treatment of the topic in chap. 5 of *Fear of Knowledge*. Boghossian formulates epistemic relativism as the conjunction of three theses (the second of which is abridged here):

- (R1) There are no absolute facts about what a particular item of information justifies. (Epistemic non-absolutism)
- (R2) Epistemic judgments of the form "E justifies belief B" express the claim: "According to the epistemic system C, that I, S, accept, information E justifies belief B." (Epistemic relationism)
- (R3) There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others. (Epistemic pluralism) (2006: 73)

Both R1 and R3 are good formulations of the standard version of relativism I shall consider. It is fundamentally the thesis that there are no objective or absolute facts that make an epistemic system right or correct. If the justifiedness or unjustifiedness of beliefs and other doxastic states is linked in the indicated way to the objective rightness or correctness of a unique system of epistemic norms (E-system), then if relativism is true, no objective status (e.g., truth or falsity) attaches to statements that a particular

doxastic state is justified or unjustified, reasonable or unreasonable. Thus, epistemic relativism seems to be equivalent to epistemic nihilism.

I am uncertain about this interpretation of the conjunction of R1, R2, and R3 because thesis R2 proffers a construal of justification statements that seems inconsistent with nihilism. Epistemic relationism says that ordinary justification statements covertly refer to the epistemic system that the speaker accepts. It offers a relational translation of justification statements that ostensibly promises an escape from nihilism. Of course, Boghossian raises serious problems for relationism (a critique I won't undertake to assess²). But the relationist component of relativism presented by R2 seems to be non-nihilistic. For this reason, R2 doesn't mesh so well with R1 and R3, in my view. So I am cautious about saying, unqualifiedly, that (epistemic) relativism is a form of nihilism. However, I don't myself wish to use relationism as a partial specification of E-relativism. Under my preferred construal, standard E-relativism is indeed a form of nihilism.

How does epistemic relativism, as defined by R1 and R3, relate to Feldman's Uniqueness Thesis and to the dispute over reasonable disagreement (among peers with shared evidence)? As defined by R1 and R3, E-relativism clashes with the Uniqueness Thesis. Moreover, as a species of nihilism, E-relativism undercuts the entire dispute about reasonable disagreement. Agents who disagree in their attitudes toward a given proposition aren't objectively unreasonable because, without an objectively correct E-system, their attitudes cannot be assessed as objectively unreasonable. At the same time, there can be no assessment of their attitudes as objectively reasonable or justified. So, as would be expected under nihilism, the entire issue simply melts away.

Is there any respectable form of non-nihilistic relativism, and what would such a form of relativism imply about the dispute over reasonable disagreement? A chief aim of this paper is to articulate a form of non-nihilistic relativism and explore its ramifications. First, however, I shall advance a brief defense of reasonable disagreement unconnected with relativism. This defense focuses on matters that need to be settled before discussing relativism, i.e., how to conceive of E-systems and their connection to justifiedness or reasonability.

2. Epistemic Systems, Doxastic Categories and Reasonable Disagreement

An epistemic (E-) system is a system of rules or norms directed at doxastic attitudes or choices. The norms in question presumably take roughly the following form: "If an agent has such-and-such evidence pertinent to proposition P, or possesses such-and-such prior beliefs, or undergoes such-and-such experiences or cognitive processes (perceptual, memorial, or reasoning processes), then doxastic attitude D is the appropriate attitude for the agent to hold vis-à-vis P." With respect to such systems of norms, we can formulate the following linkage principle L:

(L) Agent A is justified in holding doxastic attitude D vis-à-vis proposition P iff A's total evidence vis-à-vis P (e.g., antecedent beliefs, experiences, and/or cognitive processes relevant to P) is such that the objectively right epistemic

system implies that D is the appropriate attitude for A to adopt vis-à-vis P; in other words, iff A's holding D conforms to the right epistemic system.³

The formulation of principle L links justifiedness not to any random E-system but to a right E-system, because there are indefinitely many possible E-systems and conformity with an arbitrary system does not confer genuine, objective justifiedness. Only a right epistemic system has the appropriate connection with objective justifiedness or reasonability (see Goldman 1986, chaps. 4-5).

One question here is what “appropriate” should mean? Should it be construed as permission or prescription? A number of writers, myself included (Goldman 1986), opt for the permission construal. The prescription construal might seem to bias the landscape against the possibility of reasonable disagreement. For various theoretical reasons, however, I am going to adopt the prescription construal. As we shall see, this doesn't unduly prejudice the case against reasonable disagreement.

A crucial question is: What makes an E-system correct, right, or best? What is the ground, rationale, or criterion that confers such a status on an E-system? The question isn't whether or how we could we tell which E-system is right; this is an epistemological question.⁴ The question is a metaphysical one, about the constitution or ground of epistemic rightness. There are various possible approaches, and although I won't defend any such approach in detail, we should at least get a feel for some of the alternatives to persuade ourselves that the notion of such a ground or rationale isn't a mere chimera.

One family of approaches is externalist, roughly reliabilist. Here's a specimen of this approach, a reliabilist criterion of system superiority or comparative goodness that might induce a uniquely correct E-system.

(RCSS) Epistemic system E is better than epistemic system E* iff belief-forming practices that conform to E would produce a higher proportion of true beliefs than belief-forming practices that conform to E*.

One obvious worry here is whether RCSS would really induce a uniquely best system. A second problem is that RCSS takes account of only one type of credal state: belief. Since a general theory of justified credal states is desirable, not merely a theory of justified belief, shouldn't a truth-based criterion also make use of the truth-values of graded doxastic states in addition to flat-out belief?

The latter problem might be accommodated by moving from reliability to the related notion of degree-of-truth-possession, or veritistic value (Goldman and Shaked 1991; Goldman 1999a). Just as we say that someone “possesses” the truth categorically when she categorically believes something true, so we can associate with a graded belief a degree of truth possession (N.B., not a degree of truth) as a function of the degree of belief and the truth-value of its content. A graded belief of degree n ($0 \leq n \leq 1.0$) with respect to P is assigned a degree n of truth-possession if P is true, and a graded belief of

degree n is assigned a $1-n$ degree of truth-possession if P is false. Thus, having subjective probability or credence 0.70 with respect to P yields a 0.70 degree of truth-possession if P is true and a 0.30 degree of truth-possession if P is false. And so forth. We might then propose the following truth-possessional criterion of system superiority (which, under suitable assumptions, might induce a uniquely correct system):

(TPCSS) Epistemic system E is better than epistemic system E^* iff conformity to E would produce (in the long run) a higher total amount of degrees of truth-possession than conformity to E^* would produce.

Doubtless this criterion is also open to criticism. I merely offer it as an illustration.

What about internalist criteria of system superiority or system goodness? One possible internalist criterion of rightness is intuitive compellingness in reflective equilibrium. If a norm is intuitively compelling, after suitable reflection, this might make it right or correct. A system of all such norms would be a uniquely correct E -system. Again I don't mean to endorse this intuition-based criterion or ground of E -system rightness. But it illustrates a species of internalist approach that some may find appealing.⁵

I turn now to the problem of reasonable disagreement, and offer a first, very simple argument for the plausibility of reasonable disagreement. In considering the relationship between psychological attitudes and epistemic prescriptions for psychological attitudes, the following mismatch can in principle arise. The minimal "width" of doxastic attitudes might be narrower, at least in many cases, than the width of the categories employed by some (correct) prescriptions. Presumably, there are psychological limits on how narrow or wide a doxastic state can be. For example, the ordinary-language category of belief doesn't seem to designate a maximally narrow doxastic attitude. This is why many theorists prefer to talk about gradations of belief or degrees of confidence. On the other hand, it is questionable that we can have graded beliefs as fine as point probabilities -- that's too narrow in terms of psychological feasibility.⁶ But moderately fine-grained degrees of credence are certainly available.

Now it seems unlikely that correct E -norms will make doxastic prescriptions only in categories as narrow as the narrowest graded beliefs. On the contrary, for many evidential situations, correct E -norms will probably issue prescriptions in doxastic categories substantially wider than the narrowest graded beliefs. For example, the Intergovernmental Panel on Climate Change, on February 2, 2007, made projections that, they said, were "very likely," translated as "better than 90 percent." In its previous report, in 2001, the panel of scientists said that the confidence level for its projections was merely "likely," translated as "66 to 90 percent." Presumably, associated with each of these confidence levels was a (tacitly) prescribed doxastic attitude interval, an interval within which a correct doxastic attitude should fall. Such prescriptions, however, leave considerable leeway. If this is the right mold for correct epistemic norms (especially where the evidence is far from probative), different choices of doxastic states will each comply with the norms. Two people can have different (i.e., contrary) graded beliefs

within this interval – for example, one around 70 percent and one around 85 percent – yet each would conform to the norm. The difference in graded belief would constitute genuine (albeit mild) disagreement. The disagreement would be reasonable because each of the graded beliefs would conform to the norm.⁷

Roger White (2005) offers several intriguing arguments against such permissiveness. I don't find these arguments entirely compelling, but there isn't space to examine them here. Later I'll offer another, quite different argument for reasonable disagreement, an argument more intimately related to the distinctive themes of this paper. So I won't probe any further into this initial argument for reasonable disagreement.

3. A Different Conception of Relativism: Objectivity-Based Relativism

As previously indicated, I want to define a species of relativism that can co-exist with epistemic objectivism. Objectivism says that there is a uniquely correct E-system such that, for any proposition P and set of evidential circumstances, it prescribes to anyone with that evidence a doxastic attitude toward P within some interval. Such a prescription holds universally for all agents, whatever their community, culture, context, historical niche, etc. This is because the system is presumed to be objectively -- hence universally -- right. Let's assume the truth of objectivism and call the objectively right system "SYS". SYS's being right doesn't entail that anybody in any culture or context is justified in believing that SYS is right. In general it doesn't follow from the truth of an arbitrary proposition P that everybody or anybody is justified in believing P. Some truths are hidden; they don't automatically generate evidence of their truthfulness to all populations, or any populations. Gaining epistemic access to them may be difficult and problematic. Truths concerning E-systems are likely to be in this boat. Philosophers like to think of themselves as enlightened, but it's distinctly possible that even members of the philosophico-methodological subculture fail to be justified in believing, either with respect to the complete E-system SYS, or with respect to some of its individual norms, that it is the correct E-system or correct individual norm.

Failure to be justified in believing a correct norm to be correct is not the only possibility; people might be positively justified in believing some incorrect E-system or E-norm to be correct. In both cases, the justifiedness in question could be objective justifiedness. Thus, people could hold mistaken but objectively justified beliefs (or weaker doxastic attitudes) about E-norms. How could this occur? Wouldn't it require a right E-system, together with suitable evidence, to undermine itself? Is this possible? Yes. Let's elaborate a few plausible examples of such a scenario. But first let's back up and say more about the contents of plausible E-systems.

We can characterize all E-norms as source authorizations. Vision is one possible source, and a vision-based norm might be, "If it looks to an agent as if P, then (in the absence of defeating conditions) the agent should believe that P." Another possible source is memory, for which a related norm might be, "If an agent seems to remember that Q, then (in the absence of defeating conditions) the agent should believe that Q." Many sources will be psychological sources, like vision and memory, but some sources

probably won't be, for example, testimony. A testimonial norm might be, "If a random speaker or writer testifies that P, then (in the absence of defeating conditions) the agent should believe that P." To be sure, one couldn't apply a testimonial norm without relying on psychological sources to decide whether a speaker has delivered testimony that P. But that doesn't undercut a testimonial source as an epistemic source. The first group of examples I shall give of justifiably believing an incorrect norm is drawn from the domain of testimony. In advancing these examples, I shall presuppose the correctness of a certain genre of testimonial norm that is widely accepted by epistemologists. However, I won't presuppose any highly specific testimonial norm as the correct one in its territory.

It is common in many cultures for children to be told by their elders that specific sources should be trusted as guides to belief.⁸ In religious communities, young children are taught that a certain scripture should be trusted as a guide to the truth about religious matters and historical events, possibly including such things as the age of the Earth and when various species came into existence. The same scripture might be cited as the supreme source on moral matters. Children are in effect given E-norms with the content, "If the scripture says P, you should believe P." In scientific educational contexts, students might be given E-norms with the content, "If scientific researchers agree on P, you should assign a high credence to P."

Are children in such instructional contexts justified – objectively justified– in believing that such norms are correct? Whether they are so justified depends on the contents of genuinely correct E-norms. Although epistemologists do not speak with one voice about these contents, almost all believe that generic testimony-based norms -- norms concerning testimony from arbitrary speakers -- are among the right norms. Whatever the exact contents of generic testimonial norms, it's plausible that when children receive religious or scientific instruction -- especially early instruction, when their ability to engage in autonomous criticism is relatively weak -- the instruction received from their teachers or parents renders them objectively (O-) justified in believing that the norms so transmitted belong to a correct E-system. The children are O-justified in accepting such norms. This seems especially clear if the children hear roughly the same testimony from numerous elders and no conflicting testimony, a likely scenario in many communities both historical and contemporary.

This paper doesn't aim to resolve which of the specific norms just illustrated are correct and which aren't. But many of the norms conflict with one another in such a way that not all could belong to a uniquely right E-system. For example, students in different contemporary American educational systems are exposed to different teachings about the epistemic force of evolutionary science. Those instructed in a modern biology curriculum are taught to assign high credence in whatever evolutionary science says about the world. Although this material might not be explicitly formulated in the form of E-rules, such an implication would be present. By contrast, students taught in fundamentalist schools (especially private ones, with an Intelligent Design mission) are taught to be skeptical about whatever evolutionary science says. They are encouraged to accept E-norms urging low levels of credence in evidence of that kind. Given their respective exposures to the testimony of their teachers and the presumed correctness of

generic testimonial norms, all students would be justified in believing the recommended E-norms to be correct. But the two E-norms concerning evolutionary science clearly conflict with one another, so they cannot both belong to the uniquely correct E-system. So we have at least one type of case in which an E-norm is justifiably believed to be right but doesn't in fact belong to a right E-system.⁹

The foregoing examples, however, feature derivative norms as opposed to fundamental ones. Even if one grants that generic testimonial norms are fundamental, surely testimonial norms prescribing trust in particular texts or authorities must be derivative norms. Readers might concede that false though justified beliefs can be held with respect to derivative norms but resist the idea that the same holds for fundamental ones.

But there is no difference here. Even a fundamental norm can have its apparent authorizing credentials amended by experience. I don't mean that the norm ceases to be correct, only that some cognizer ceases to be justified in believing it to be correct. A standard example of a fundamental norm is "If it looks to you as if P, then (in the absence of defeaters) you should believe that P." Now imagine a scenario in which someone hears credible testimony to the effect that the visible world is a sham or delusion, so that vision shouldn't be trusted (cf. the Matrix). Isn't this a case in which the epistemic status of a fundamental norm for an individual is affected by his experience? (This assumes a reliabilist or veritistic grounding of E-norm correctness.) The same thing can transpire for a norm of reasoning, which might seem to be a paradigmatic fundamental norm. Choose your favorite methodological norm: a statistical norm like the Neyman-Pearson method, the Chi-square method, or even a Bayesian norm. Such norms are often the subjects of serious debate and critique -- in statistical-theory circles or philosophical circles. If a novice hears a lecture from a well-certified theoretical statistician or philosopher of statistics that mounts a compelling critique of such a norm, the hearer could well be justified in reducing or moderating her credence in the correctness of the norm. So the justificational status of even a fundamental norm can be amended by application of other norms, ultimately by appeal to the ground of E-system rightness.¹⁰

We have been considering cases in which a person is objectively justified (O-justified) in believing of a certain norm that it is correct or incorrect. But we might also be interested in cases in which someone is O-justified in having a graded belief rather than a full belief in a norm's correctness or incorrectness. Although her evidence might not support full belief in norm N's correctness, it might support a credence of, say, 0.60. Our framework permits justifiedness not only for full beliefs but for all grades of credence.

We can now state the central theses of the new form of relativism I wish to consider: objectivity-based relativism. These theses can be formulated as follows:

(OBR) There is a uniquely correct E-system that governs the objective justifiedness and unjustifiedness of people's doxastic attitudes. However, people occupy different evidential positions vis-à-vis this system and other candidate E-

systems. Hence, the objective justificational status of different people vis-à-vis different E-systems is varied rather than uniform. Some people are objectively justified in believing certain E-norms and E-systems to be correct; others are objectively justified in believing other E-norms and E-systems to be correct. Similarly for attitudes other than full belief toward E-norm-related propositions.

Objectivity-based relativism is very different from nihilistic relativism, and also -- by my lights -- fairly attractive.¹¹ It has the virtue of accommodating an important intuition that actuates many proponents of E-relativism, the intuition that differences in intellectual procedure found in diverse cultures, communities, and historical periods do not reflect wholesale irrationality or epistemic depravity. There is something epistemically legitimate about divergent choices of procedures. Objectivity-based relativism captures this intuition by allowing members of epistemically diverse cultures to have objective justification (O-justification) for different beliefs about intellectual norms. In virtue of this norm-justification, they may also enjoy a distinct but significant justificational status for their garden-variety beliefs (beliefs about ordinary matters rather than E-norms). This is the status of being O-justified in believing that they are O-justified in believing P. When they use their adopted E-norms to form beliefs in garden-variety propositions, these beliefs will often fail to be O-justified. Nonetheless, they may be iteratively O-justified: the people are justified in believing that their beliefs are justified.

Suppose Amanda is O-justified in believing norm X to be a correct E-norm. Furthermore, given Amanda's evidential circumstances, norm X authorizes her to believe proposition P. Then she is O-justified in believing that she is O-justified in believing P.¹² However, second-order O-justifiedness doesn't entail first-order O-justifiedness.

(Non-entailment) $J_o[J_o(P)] \not\Rightarrow J_o(P)$

Perhaps norm X is actually incorrect although Amanda is O-justified in believing it to be correct. Objectivity-based relativism does not imply that all norm-systems are equally right in the sense of being equally capable of conferring first-order justifiedness.

It may be helpful here to flag the distinction between propositional and doxastic justifiedness. A person is doxastically justified in having attitude D toward P if she actually has D and it's justified. A person is propositionally justified in having attitude D toward P if her epistemic position is such that D is the proper attitude to adopt toward P -- whether or not she actually adopts it. Arguably, iterative justifiedness makes better sense when interpreted in propositional rather than doxastic justificational terms. This is because comparatively few individuals form explicit beliefs about the justificational status of their own (first-order) attitudes. Only fairly reflective minds contemplate this sort of thing. Nonetheless, even if Jerome doesn't actively wonder whether he is justified in believing P, and hence doesn't come to any belief (or other opinion) on the subject, his evidential condition might entitle him to believe that he is so justified. Hence, he is propositionally justified in believing that he is justified even if he isn't doxastically justified. The truth of this iterative-justificational proposition might well interest epistemologists.

4. Objectivity-Based Relativism and Reasonable Disagreement

I turn now to OBR's implications for the reasonable disagreement controversy. The implications I'll extract are fairly limited in scope, because objectivity-based relativism bears on the reasonable disagreement issue only from a single restricted angle, that of iterative justifiedness. The analytical framework presented here doesn't provide the resources for a full-bore attack on the reasonable disagreement problem, because it takes no stance on the contents of a right E-system. Without specifying such contents, it is hard to draw firm conclusions about the doxastic moves an agent should make if she learned various things about her peers, such as the fact that some of them disagree with her. Should she stick to her guns in believing P? Should she "split the difference" with them? These questions cannot be adequately answered without identifying the right E-system – at least the general contours of such a system. Although I won't tackle these central issues, I'll use our broader analytical framework (mainly, the linkage principle) to reveal a connection between iterative justifiedness and reasonable disagreement.

An objectivist framework for E-system rightness requires any two people who have the same total evidence vis-a-vis P to take the same attitude toward P – at least if we ignore permissible differences within the prescribed attitude, as discussed in section 2. (Henceforth I ignore such differences.) For both individuals to have objectively justified attitudes toward P, their attitudes must be the same. If they differ, at most one attitude can be justified. Hence, reasonable disagreement is precluded at the first-order level of justifiedness. Even if their attitudes differ, however, each might be objectively justified in believing that her attitude is (objectively) justified. In other words, disagreement among evidentially equal agents is compatible with each agent possessing second-order justifiedness.

Here's a scenario by which this can transpire. Amanda and Jerome have the same evidence with respect to P but different evidence about E-system correctness. In virtue of this evidence, Amanda is O-justified in believing system E to be correct whereas Jerome is O-justified in believing E* to be correct. Finally, the attitude required by E toward P (given the specified evidence) is incompatible with the attitude required by E*. Thus, Amanda is justified in believing that she is justified in adopting attitude D toward P whereas Jerome is justified in thinking that he is justified in adopting attitude D* toward P, where D and D* are incompatible. At the first-order level of justification such a difference in attitude implies that at least one of them is unreasonable, but at the second-order level of justification both can be reasonable -- i.e., iteratively justified.

The preceding sentence incorporates a crucial step in the argument: the proposal that higher-order justifiedness can ensure, or at least make a positive contribution toward, the reasonability of a first-order belief. This is despite the fact that higher-order justifiedness doesn't entail first-order justifiedness. Thus, the reasonability of an agent's attitude toward P isn't fixed by its first-order justificational status. This point is worth marking with a new principle:

($J^1 \not\Rightarrow R$) The first-order justificational status of an attitude does not fix its (overall) reasonability; reasonability can also be influenced by higher-order justificational status.¹³

What considerations might support this principle? And if second- (or higher-) order justifiedness counts in fixing reasonability, how much does it count?

A first point to make is that second-order justifiedness has some epistemic value, indeed, substantial value. Consider an agent who (i) forms a justified but mistaken belief that system E is correct, (ii) correctly applies E's requirements to her own evidential state, and therefore (iii) selects attitude D toward P. How well does this agent proceed in epistemic terms? She clearly proceeds well at stage (i). She justifiably forms a belief that E is correct. She cannot be faulted there in epistemic terms. Similarly, how can she be faulted for the procedures she executes at stages (ii) and (iii)? In these stages, the norms she justifiably believes to be correct are applied to her evidence, and her attitude toward P is formed on the basis of this evidence. Perhaps she can be faulted for failing to obtain a true belief at any of these stages. But if we assume fallibility even for objectively right E-norms -- and I do assume such fallibility throughout -- then truth attainment is never guaranteed by first-order, second-order, or any order of justifiedness. So why should failure to obtain the truth imply culpability? In short, when a person's belief enjoys second-order justifiedness, there is much to be said for her epistemic conduct. If epistemic conduct can be characterized as "culpable" or "non-culpable," a belief's second-order justifiedness entitles an agent to a respectable level of non-culpability. At a minimum it makes a contribution toward attainment of an overall level of positive non-culpability or reasonableness. Furthermore, it's a contribution that might trump the epistemic culpability associated with holding a belief (or other attitude) that is first-order unjustified.

Some might complain that epistemic non-culpability is a rather weak status, not strong enough to imply justifiedness or reasonability. But the argument can be re-phrased in terms of "propriety" of epistemic conduct. Doesn't an agent engage in proper epistemic conduct if she applies the correct norm-system to her evidence at stage (i) to form a belief that system E is correct? Assume that no later evidence mandates a change in this belief. Given her proper choice of system E, doesn't she engage in proper epistemic conduct at stages (ii) and (iii) in applying E's requirements to her P-relevant evidence and selects attitude D? By similar steps, a different agent might properly form a belief that system E* is correct and properly select a different attitude D* with respect to P, despite having the same P-relevant evidence as the first agent. Thus, epistemic peers who share the same P-relevant evidence can reasonably disagree about P even when this involves first-order unjustifiedness on the part of at least one of them.

Here's an additional consideration to support the significance of second-order justifiedness. When considering the reasonability of someone's belief, its truth-value doesn't settle the issue. A false proposition can be reasonably believed. What determines a belief's reasonability is the agent's evidence (or belief-forming methods), not the belief's truth-value. The same point holds on the topic of norm correctness. The

actual rightness of an E-system doesn't determine the reasonability of an agent's conforming to it. What is critical is the agent's evidence about its rightness. If an agent conforms her attitude to the prescriptions of a properly chosen E-system, this should be an important -- perhaps decisive -- element in assessing the attitude's reasonability, even if the evidence supporting that E-system's rightness happens to be misleading.

This proposal poses a problem, however. If second-order justifiedness is relevant to reasonability, why isn't every order of iterative justifiedness relevant? Indeed, given what we have said, shouldn't each higher order of iterative justifiedness be more relevant to reasonability than its predecessor? Won't this generate a vicious infinite regress, which threatens to scotch the entire enterprise of assigning determinate justificational statuses to doxastic attitudes? Each higher-order status will trump the immediately lower-order status, and as the orders increase they will tend to swamp first-order justifiedness entirely. Doesn't it become radically unclear what overall reasonability consists in, or whether it can be determinate?

The problems in this territory aren't as devastating as the foregoing portents suggest. As one ascends the hierarchy, the evidence an agent possesses vis-à-vis the preceding level of iterative justifiedness rapidly becomes negligible. In fact, it may quickly become null. If the right E-system is anything like what epistemologists suppose, an agent will typically be instructed, at the $n+1^{\text{st}}$ level, to suspend judgment about the n^{th} level of justifiedness. While justified belief about a lower level of iterative justifiedness can trump lower-level justifiedness, justified agnosticism, or suspension of judgment, should not have comparable trumping power. Even an infinite series of judgment suspensions will be in the same boat. So the threat of higher order justifiedness totally swamping first-order justifiedness is not so severe.

It would be helpful, no doubt, to quantify the appropriate weightings for lower-order and higher-order justificational status. How exactly do they influence overall, or ultima facie reasonability? Unfortunately, I do not know how to address this issue in adequate generality. Two points should suffice for present purposes. First, it is not proposed that first-order justifiedness gets "washed out" entirely by second-order justifiedness. Forming opinions in accord with an objectively right E-system is surely worth something, if not everything, in terms of justification and reasonability. The suggestion is only that first-order justifiedness can be outweighed or superseded by higher-order justifiedness. Second, for present purposes, we don't have to specify the precise circumstances in which higher-order justifiedness trumps first-order justifiedness. As long as this can sometimes happen, it falsifies the general principle that two people with the same (first-order) evidence vis-à-vis P cannot reasonably adopt different attitudes toward P. The foregoing considerations adequately establish that this can sometimes happen.¹⁴

5. Evidence and the Peer Disagreement Controversy

Participants in the peer disagreement controversy are likely to complain that I am ignoring their controversy because their controversy centers on cases in which people have the same evidence vis-à-vis target proposition P. It concerns cases involving epistemic peers, where the peer relationship typically includes “evidential equality” (as Christensen 2007 calls it). By contrast, my cases are ones in which people have evidential differences, i.e., differences concerning the correct E-norms. So, critics might mutter, how am I contributing to the debate?

My contribution might be viewed from the following perspective. It contributes to the debate by identifying a category of evidence that bears on the reasonableness of peer disagreement but is generally ignored in the literature. Contributors to the debate typically divide the determinants of reasonability into two sectors. The first sector consists of the agents’ evidence relevant to the target proposition. This evidence is usually divided into three categories: (a) evidence “directly” concerning the target proposition, (b) evidence concerning one’s own epistemic competence, and (c) evidence concerning the peer’s epistemic competence. The second sector consists of rules or norms that should govern their epistemic conduct. Such rules are prescriptions or permissions, which are not, strictly speaking, propositions. Hence they are not the sorts of things for which there can be evidence; the three types of evidence in the first sector exhaust the evidence relevant to peer disagreement. However, I am pointing out an additional type of proposition with respect to which evidence might diverge. This is a proposition of the form, “Norm X is a correct norm (and applies to the present doxastic choice).” Two agents can have different bodies of evidence that bear on norm correctness and are relevant to the reasonability of their respective attitudes.¹⁵ So here we highlight a species of evidence – norm evidence, we might call it, as contrasted with material evidence – that is generally ignored in the literature.

The peer disagreement literature tends to miss this point because it presumes that justifiedness or reasonability is conferred by de facto norm correctness. Correct norms, not an agent’s evidence about the correct norms, set the standard for epistemic conduct. The issue of norm evidence simply isn’t raised. I am arguing that norm evidence is among the determinants of reasonability. Where two agents are equal with respect to material evidence but differ with respect to norm evidence -- though the correct norm-system stays fixed – it is legitimate for their attitudes toward a given proposition to diverge.

If this is my view, don’t I side with those espousing the maxim “No reasonable disagreement without evidential difference”? Yes, that’s a fair characterization of my position -- if all categories of evidence are included. But if attention is restricted to material evidence (“sector 1” evidence), as it usually is, this maxim doesn’t characterize my position. In cases where evidential equality extends only to material evidential equality, there is room for reasonable disagreement.¹⁶

As acknowledged at the beginning of section 4, this paper doesn’t develop a comprehensive approach to peer disagreement. However, let me identify some other

contours of the topic and explain why I remain silent on many of them. I'll also identify some problems with existing treatments, especially concerning the nature of evidence.

My approach to peer disagreement, it might be said, embraces a synchronic perspective. If two agents are evidential equals with respect P at time t, can they reasonably differ in their attitudes toward P? A more standard perspective is a diachronic one. The diachronic question is how an agent should change her opinion vis-à-vis P over time. It focuses on the following problem: At time t an agent forms an opinion vis-à-vis P in ignorance of a certain peer's opinion. At a later time t*, the agent learns that the peer, despite being an evidential equal, holds a different opinion. How (if at all) should the agent revise her opinion? Here is David Christensen's diachronic formulation of the problem, which is fairly representative: "How should I react when I discover that my friend and I have very different beliefs on some topic? ... Should my discovery of her differing degree of belief in P lead me to revise my own confidence in P?" (Christensen 2007: 188) We might call this the peer responsiveness formulation of the problem.

The framework I employ here lacks sufficient tools for a detailed analysis of the peer responsiveness problem. Ignoring the arguments of Section 2, the framework implies that under complete evidential equality (including norm-evidence equality) two people cannot reasonably have differing opinions. But suppose evidential equality isn't complete. We can still ask what doxastic choices are expected of peers when they discover their disagreement? Must their degrees of belief converge? Must the mode of convergence involve "splitting the difference"? To tackle these issues we need more than framework principle L. We need a correct and detailed E-system. However, a formulation, defense, and application of such an E-system is beyond the scope of this paper. Nonetheless, we can further contribute to the debate in two ways: first, by identifying weaknesses in certain treatments of evidence and evidential equivalence, and second, by pinpointing the impact of alternative theories of evidence.

Contributors to the peer disagreement topic tend to assume that evidential equality is a readily producible scenario, that two people can make themselves evidential equals by simply "sharing" relevant evidence with one another. Feldman (forthcoming) makes much of this procedure. The assumption seems to be that if one person verbally communicates what he regards as his (relevant) evidence, a hearer acquires the same evidence as the speaker. If the second person reciprocates, both will have shared their total information with their opposite number and they will be evidential equals. This assumption, however, involves an unnoticed pun on the word "share." Sharing evidence in the sense of communicating the content of an evidential state does not necessarily imply that the hearer shares – in the sense of possesses -- the same evidence as the speaker. This is well illustrated by the following example due to Apolonio Latar (2007).

Suppose Billy is accused of committing a certain crime and his friends have weighty evidence that he did it. They know, for example, that he threatened to commit a crime of that very description the day before it happened. In fact, Billy did not commit the crime. However, he was alone in his room when somebody else was committing it, so he cannot prove to his friends that he didn't do it. He clearly remembers not having

done it (he recalls not being near the crime scene, etc.), and this vivid memory is excellent evidence for him. But Billy cannot literally transmit this memory to his friends (e.g., by duplicating this portion of his brain state in them). All he can do is verbally report its content. Even if his friends believe his report, this doesn't make them evidentially equal to Billy vis-à-vis the criminal accusation. They believe he is innocent, but they don't have personal recall to support that belief. Moreover, as Latar points out, even if the friends accept Billy's report, they may not accord it as high a degree of confidence as he does. So the proposition will have less evidential power for them than it does for Billy.

Feldman rightly supposes that if one person verbally shares his evidence with another, the latter acquires evidence of the speaker's evidence. He therefore articulates the principle, "evidence of evidence is evidence." What he seems to mean by this is that if Smith truthfully reports certain evidence concerning P (call it Q), this evidence is also acquired by the hearer. This is not quite right. Hearing such testimony may give the hearer default justification for believing Q, but such default justification can be defeated by other information the hearer possesses. In that case, Q doesn't qualify as an item of evidence for the hearer. Furthermore, even if there is no such defeat, the hearer does not necessarily acquire the same evidence possessed by Smith. Smith's saying that he had a certain visual experience, for example, doesn't reproduce in the hearer the same visual experience, with its full evidential load. So the hearer doesn't acquire the same evidence for P as Smith has.

Furthermore, it's impossible to convey to others all the subtle strands of evidence one harbors, or has harbored, for one's opinions. For example, failing to observe any counterexamples to a certain hypothesis may justify one's acceptance of it -- at any rate, if there is a high likelihood that one would observe such counterexamples if the hypothesis were false. But the evidential "omissions" that collectively constitute this (past) support tend not to be stored in memory and aren't readily retrieved if one is asked to defend one's belief.¹⁷ Finally, one often forgets even past observations that play a lively causal role in belief acquisition. These now-forgotten observations are relevant to the current justificational status of a belief that has been preserved over time, but they aren't available for "sharing" when asked for one's evidence. Thus, what speakers manage to communicate when asked for their reasons rarely approximates the whole of their relevant evidence.

The peer disagreement literature assesses the extent to which people should defer to others with similar evidence and cognitive competence but differing opinions. It asks what it would be rational or reasonable for people to do in such cases. I have offered some conclusions about what reasonability requires based on a very general conception of epistemic justifiedness. It must be emphasized, however, that such general conclusions cannot be applied to concrete cases without first settling the question of what counts as evidence. This in turn depends on the contents of the correct norms.

Let me illustrate this point with the help of Kelly's example of a mathematician who is initially confident he has a proof of a certain theorem but whose colleagues deny

that it's a genuine proof (see note 16). Kelly characterizes the case as one of evidential equality, because each mathematician has surveyed exactly the same evidence, presumably, the steps of the proof as written on paper. Another epistemologist, however, might hold that the parties are unlikely to have the same evidence in Kelly's example. Suppose that although all the steps in the proof are correct (as Kelly supposes), each colleague, while examining the proof, seems to detect a mistaken step. The proof creator, while reviewing his proof, has no such experience of seeming to spot an error. Then under some conceptions of evidence the parties do not have the same evidence. Consider, for instance, Michael Huemer's principle of "phenomenal conservatism," which may be classified as a principle of evidence: "If it seems to S as if P, then S thereby has at least prima facie justification for believing that P" (2001: 99). If this evidence principle is correct, the mathematicians don't all have the same evidence. Each colleague has evidence of the proof's containing a mistaken step whereas the proof creator has no such evidence. In general, comparatively "subjective" conceptions of evidence will tend to produce fewer cases of exact evidential equality than comparatively "objective" conceptions.¹⁸

Conceptions of evidence will not coincide across all E-systems. A conception of evidence is implicitly specified by the set of antecedents of an E-system's (conditional) prescriptions. These antecedents fix what the system considers to be evidence, and they won't generally be the same across E-systems. An upshot of this is that we cannot settle questions about evidential equality without settling questions about the properties of a right E-system. This task outstrips the compass of the present essay.

6. Conclusion

The paper has undertaken two projects. The first is to advance a non-standard conception of epistemic relativism under which relativism is compatible with objectivism or absolutism. Even if there is a uniquely right system of E-norms, people in different communities can justifiably (though not correctly) accept different E-systems as right. The second project is to draw a moral for the problem of reasonable disagreement. Once we distinguish first-order and second-order justifiedness, we find that two people with the same ("material") evidence for proposition P can have contrary attitudes toward P both of which are second-order justified. Since second-order justifiedness is (or can be) as important a determinant of reasonability as first-order justifiedness, these divergent attitudes can both be reasonable.¹⁹

References

- Barnes, B. and Bloor, D. (1982). "Relativism, Rationalism and the Sociology of Knowledge," in M. Hollis and S. Lukes, eds., Rationality and Relativism, pp. 21-47. Cambridge, MA: MIT Press.
- Boghossian, P. (2006). Fear of Knowledge. New York: Oxford University Press.
- Boghossian, P. (2007a). "The Case Against Epistemic Relativism: Replies to Rosen and Neta," Episteme: A Journal of Social Epistemology 4(1): xxx.
- Boghossian, P. (2007b). "Epistemic Rules." Presented at the Rutgers Epistemology Conference, Rutgers University, May 5, 2007. Forthcoming in Philosophy and Phenomenological Research.
- Broad, C. D. (1985). Ethics (edited by C. Lewy). Dordrecht: Martinus Nijhoff.
- Christensen, D. (2007). "Epistemology of Disagreement: The Good News," Philosophical Review 116: 187-218.
- Feldman, R. (forthcoming). "Reasonable Religious Disagreements," in L. Antony, ed., Philosophers without Gods: Meditations on Atheism and the Secular Life. New York: Oxford University Press.
- Goldman, A. (1992). "What Is Justified Belief?" Reprinted in Goldman, Liaisons: Philosophy Meets the Cognitive and Social Sciences. Cambridge, MA: MIT Press.
- Goldman, A. (1986). Epistemology and Cognition. Cambridge, MA: Harvard University Press.
- Goldman, A. (1999a). Knowledge in a Social World. Oxford: Oxford University Press.
- Goldman, A. (1999b). "Internalism Exposed," Journal of Philosophy.
- Goldman, A. and Shaked, M. (1991). "An Economic Model of Scientific Inquiry and Truth Acquisition," Philosophical Studies. Reprinted in Goldman, Liaisons: Philosophy Meets the Cognitive and Social Sciences (1992).
- Huemer, M. (2001). Skepticism and the Veil of Perception. Lanham, MD: Rowman and Littlefield.
- Kelly, T. (2005). "The Epistemic Significance of Disagreement," in T. Gendler and J. Hawthorne, eds., Oxford Studies in Epistemology, vol. 1. Oxford: Oxford University Press.

Kelly, T. (this volume). "Peer Disagreement and Higher Order Evidence." In R. Feldman and T. Warfield, eds., Disagreement. New York: Oxford University Press.

Lackey, J. (2006). "It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony," in J. Lackey and E. Sosa, eds., The Epistemology of Testimony, pp. 160-189. Oxford: Oxford University Press.

Latar, A. (2007). "Shared Evidence and Reasonable Disagreement." Department of Philosophy, Rutgers University, undergraduate essay.

Neta, R. (2007). "In Defense of Epistemic Relativism," Episteme: A Journal of Social Epistemology 4(1): xxx.

Pettit, P. (2006). "When to Defer to Majority Testimony – and When Not." Analysis 66: 179-187.

Rosen, G. (2007). "The Case Against Epistemic Relativism: Reflections on Chapter 6 of Fear of Knowledge," Episteme: A Journal of Social Epistemology 4(1): xxx.

Tversky, A. and Kahneman, D. (1983). "Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment," Psychological Review 90: 293-315.

White, R. (2005). "Epistemic Permissiveness," in J. Hawthorne, ed., Philosophical Perspectives, vol. 19, Epistemology, pp. 445-459. Malden, MA: Blackwell Publishing.

¹ A plausible terminology to adopt here would distinguish weak and strong objectivism. Weak objectivism would be the view that all pairs of E-systems have an objective ranking in terms of comparative (E-) goodness, and strong objectivism would be the view that some E-system is uniquely best in terms of such a ranking. In this terminology, our focus here is strong objectivism. However, the phrase ‘strong objectivism’ might have unintended connotations. It might suggest a highly rigid or constraining brand of objectivism, which issues extremely fine-grained prescriptions for all evidential scenarios. This would be the opposite of what Roger White (2005) calls “epistemic permissivism.” My conception of strong objectivism, however, does not have this implication. A uniquely best E-system might be a fairly loose or permissive system. See section 2 below.

² For additional discussion of possible solutions to the problems of relationism, see a mini-symposium on Boghossian (2006) consisting of Rosen (2007), Neta (2007), and Boghossian (2007a).

³ Notice that the linkage principle only requires doxastic attitude D to conform to the right norm system in order to be justified. It does not require the agent to follow the system’s rules in arriving at D. In particular, it doesn’t require the agent to mentally represent the relevant norms or be mentally guided by them. For example, assume that the correct E-system includes perceptual and memory norms, which “approve” of an agent’s holding perceptual or memory beliefs under specified circumstances. A non-reflective agent, who doesn’t mentally represent these norms, can nonetheless justifiably hold such beliefs as long as she conforms with the norms. It would be an excessive demand to place on justifiedness to require rule-following “all the way down.” See Boghossian (2007b), who argues that systematic rule-following involves a vicious regress.

⁴ Boghossian (2006) examines an argument in support of nihilistic relativism based on the premise that if there is an objectively right E-system, it is possible to be justified in believing that it is right. The argument then proceeds to deny that this is possible on the grounds that such justification would involve norm circularity, because it would have to presuppose the rightness of the system. Obviously, the argument also needs the further premise that norm circularity vitiates justifiedness. A worry I would pose for this argument (different from Boghossian’s criticisms) is where any proponent of such an argument would get the last premise. What justifies the premise that norm-circular arguments are justificationaly impotent? Is this premise embedded in the right E-system? Is it embedded in every E-system? Is it a higher-level constraint on any E-system? Each of these assumptions is problematic. So it is hard to see how to construct a successful epistemological argument of this sort against the existence of a uniquely correct E-system. Such reflections tilt against the cogency of any putative epistemological constraint on the existence of an objectively right E-system comparable to the constraint that it must be possible to have non-circular justification for such a system.

⁵ For an illustration of one possible ramification of an intuition-based criterion, see note 10 below.

⁶ Of course, the propositional content of a belief might include point probabilities. But that’s not pertinent to the present discussion.

⁷ Even if we adhere to the tripartite scheme of belief, suspension, and disbelief on the assumption that they are the “thinnest” doxastic states available (a very implausible assumption), a right E-system could still issue prescriptions for disjunctive categories like “belief or suspension” or “disbelief or suspension”. Single-word labels could be invented for such doxastic intervals (e.g., “belension” and “disbelension” respectively). Thus, the same argument for reasonable disagreement can be presented within the tripartite taxonomy.

⁸ I focus my examples on children’s receipt of testimony because children have smaller stores of real-world knowledge or belief as compared with adults. This implies, among other things, that they possess fewer evidential resources to challenge the testimony of their elders, hence fewer evidential resources to defeat the prima justifiedness that arises from receiving such testimony.

⁹ Some writers on testimony might resist my conclusion on the grounds that a hearer's justifiedness in accepting a piece of testimony from a speaker depends not only on the hearer's evidence but also on the reliability of the speaker. In the cases before us, therefore, children might not be justified in believing what their elders' say about the trustworthiness of a specified source, because the elders in these communications are not reliable sources. Jennifer Lackey (2006) gives an example called "NESTED SPEAKER" in which Fred has reasons to believe that Pauline is a reliable testifier about wild birds, but in fact she isn't reliable on this subject. Lackey contends that when Fred forms a belief about albatrosses based on Pauline's testimony, his belief isn't justified. That isn't the result of any flaw in Fred's reasons, but rather a result of Pauline's unreliability. Similarly, as both Lackey and Baron Reed have argued to me in conversation, if the elders in my example are unreliable speakers, the children aren't justified in believing in the correctness of the elders-commended norm. One response to these arguments is to question the judgment that Fred's belief is unjustified. My own intuition about this case is murky, by no means clear-cut in Lackey's direction. However, let's concede the classification of the case for purposes of further argument. It's clear in our examples that where the norm endorsed by the speakers is incorrect (in virtue of the norm's unreliability), the speakers are unreliable on this topic. But that doesn't mean that they are unreliable speakers in general. In fact, they may well be reliable with respect to all the mundane matters on which they also testify to the children (the locations of specific rooms in the school, chalk in the cupboard, etc.). Does the NESTED SPEAKER case show that topic-specific reliability is necessary for hearer justification? No, because Pauline is unreliable in general, not just in testimony about wild birds, and it could be her general unreliability that (partly) undercuts Fred's justifiedness. Moreover, it is noteworthy that Lackey herself concedes that there are some concepts of justifiedness that escape some of the arguments in her paper. In particular, Lackey concedes that the kind of justification she calls "justification grounded entirely in one's subjective perspective" escape these arguments (2006: 182, note 1). We can take ourselves here to be addressing such a conception of justification.

¹⁰ Let's see how this might work for probabilistic norms even if the posited ground of rightness is intuitive appeal in reflective equilibrium. Assume that correct norms for probabilistic reasoning are associated with the standard probability calculus. One such norm is the prescription not to assign a higher probability to a conjunctive event than to one of that event's conjuncts. Thus, in Tversky and Kahneman's (1983) well-known Linda example, the norm would imply that one shouldn't believe that "Linda is a feminist and a bank teller (F & T)" has a higher probability than "Linda is a bank teller (T)." Could anybody be justified in believing (mistakenly) that a different, incompatible norm is the correct one? Yes. Tversky and Kahneman's results showed that naïve intuitions tend to be driven by the "representativeness," or resemblance, heuristic. This heuristic leads subjects to judge the F & T event more probable than the T event, because Linda more closely resembles a prototypical feminist bank teller than a prototypical bank teller. When subjects were explicitly presented with two arguments, one using the conjunction rule and the other using the resemblance rule, 65% of the subjects found the resemblance argument more convincing than the conjunction-rule argument. Apparently, resemblance considerations are intuitively more compelling for naïve subjects than the conjunction rule. If we now assume that undefeated intuitions provide justification, then naïve subjects who haven't been tutored in probability theory may well be justified in believing that the resemblance norm is correct, even though it conflicts with the probability calculus. The ground of rightness we are considering, however, makes norm rightness a function of intuitive appeal in reflective equilibrium. In the present case, this ground would be applied as follows. Once people reflect systematically on matters of probability, they will come to appreciate – and find intuitively compelling – the appropriateness of norms based on the probability calculus. So we have a case in which some people -- the wholly untutored ones -- are justified in accepting norms that aren't genuinely correct under the posited ground of correctness. .

¹¹ The uniqueness requirement for a right E-system is admittedly not so attractive. It might be possible to replace this stringent condition by a weaker one.

¹² Perhaps a further condition should be added here, namely, that the agent must be justified in believing that she satisfies the evidential circumstances specified in the norm. Such a condition could easily raise

questions about the nature of evidential circumstances: whether to construe them “internalistically” or “externalistically.” That’s a topic for a different occasion.

¹³ At the outset of the paper, I used the term ‘reasonable’ interchangeably with ‘justified,’ as do many epistemologists. With the present principle, however, these terms acquire slightly distinct uses or meanings. When speaking strictly, talk of a doxastic attitude’s justifiedness should henceforth be qualified by reference to the order of justifiedness in question: first-order justifiedness, second-order justifiedness, etc. An attitude’s reasonability arises from one or more of its various justificational statuses. It remains to be explored just which justificational statuses are most determinative of reasonableness and under what conditions.

¹⁴ The issues in play here have obvious analogues in moral theory, where the operative terms of appraisal are “right,” “obligated,” etc. rather than “justified” or “reasonable.” Some moral theorists who probe analogous issues in moral theory like to distinguish different senses of “right” or “obligated.” For instance, Broad (1985: 128) discusses the question of whether a person is morally obligated to render military service if he is a citizen of a country that is at war, if he is of military age, and if his services are legally demanded of him. Broad assumes, for the sake of argument, that the situation in fact makes a moral demand on him. Still, there is a question of whether his obligation is to do what the situation in fact morally demands or whether he is obligated to do only what he recognizes the situation to demand. Broad says that one can go either way here. One can say that he is obligated by what the situation in fact demands or one can say that he is only obligated to do what he recognizes to be morally demanded of him. (A better analogue of what is under discussion here would be the claim that a person is only obligated by what he is justified in believing morality to require, not by what he does believe it to require.) This prompts Broad to speak of obligation or rightness in two different senses: an objective sense and a subjective sense. He writes: “[I]t is futile to pretend that there is just one right sense of ‘right’ and one sense in which we ought to use ‘ought’” (1985:127). Obviously, a similar strategy of distinguishing senses of ‘justified’ or ‘reasonable’ can be adopted in the epistemological case. In effect, this is part of what I am doing. (Thanks to Holly M. Smith for the reference to Broad.)

¹⁵ This evidence could have been acquired in the past, of course. And it may not be readily retrievable, as discussed in the text below. It is also important in this context to note the path-dependence properties of evidence acquisition (see, for example, Pettit 2006).

¹⁶ Does the disagreement literature already recognize the variety of evidence I am highlighting? Thomas Kelly (2005, this volume) speaks of “higher order evidence”, and sometimes it sounds as if he is concerned with evidence of E-norm correctness. On balance, however, Kelly’s discussion of higher-order evidence has a rather different thrust. He argues (this volume, p. xxx) that if you have higher-order evidence to the effect that you probably made a mistake in responding to your first-order evidence, you should temper that initial confidence in the conclusion. For example, suppose you are a professional mathematician who thinks you have proved a certain theorem, but each of several colleagues claims to find a mistake in the proof. Your colleagues’ dissent is higher-order evidence about your original performance in examining the proof. This case illustrates that what Kelly means by “higher-order” evidence isn’t evidence about the content of any (correct) E-norm but evidence about your competence in arriving at your initial credence -- hence evidence about how heavily to weight this initial credence when revising your opinion.

¹⁷ Another important category of evidence that doesn’t get encoded in memory and therefore isn’t available for subsequent report is observed evidence whose significance isn’t appreciated at the time of observation. Such evidence is unlikely to be recalled later. Nonetheless, on a plausible approach to evidence, it is part of the total (cumulative) evidence that bears on a belief’s justificational status. I assume here a historical approach to justification (cf. the historical reliabilist approach of Goldman 1979). Defenders of other approaches to justification might dissent at this juncture, but the problem of forgotten or neglected evidence is important (cf. Goldman 1999b).

¹⁸ A reader may be surprised to find me giving respectful treatment of a strongly subjective conception of evidence, because this seems at variance with justificational externalism that I have favored in the past.

Two clarifications are in order. First, I am not endorsing Huemer's phenomenal conservatism, merely citing it as an extant position. Second, in the present paper's framework the principal externalist dimension of justification or reasonability would arise from the ground or criterion of E-system rightness. That's where reliability would enter the picture. An architecture that introduces external factors at this level is entirely compatible with evidential states being highly subjective or "internal".

¹⁹ Earlier versions of this paper were presented at Northern Illinois University, Princeton University, and the Bled (Slovenia) epistemology conference of 2007. I am grateful to members of these audiences for very helpful discussions. A non-exhaustive list of those whose comments influenced the final product would include Elizabeth Harman, Gil Harman, Tom Kelly, Jennifer Lackey, Ram Neta, Baron Reed, Gideon Rosen, Bruce Russell, and Holly Smith,