

Some Problems With Steadfast Strategies for Rational Disagreement

Hamid Vahid

Abstract: Current responses to the question of how one should adjust one's beliefs in response to peer disagreement have, in general, formed a spectrum at one end of which sit the so-called 'conciliatory' views and whose other end is occupied by the 'steadfast' views. While the conciliatory views of disagreement maintain that one is required to make doxastic conciliation when faced with an epistemic peer who holds a different stance on a particular subject, the steadfast views allow us to maintain our confidence in our relevant beliefs. My aim in this paper is not to adjudicate between these views. Rather, I shall focus on a particular strategy, namely, denying the appearance of epistemic symmetry between peers, that the steadfast views standardly invoke in support of their position. Having closely examined certain representative examples of the steadfast approach, I will argue that this strategy is fundamentally flawed.

Keywords: rational disagreement, 'conciliatory' views, 'steadfast' views, epistemic norms, epistemic peer

It is common for equally competent epistemic agents to arrive at different conclusions in regard to a particular subject matter on the basis of roughly the same body of evidence. Although it is rarely possible for two people to share all their evidence, let us, idealizing away from the actual cases, call two agents, S_1 and S_2 , epistemic peers if they are equal with respect to the general intellectual virtues and are acquainted with the same evidence E in regard to a particular proposition p . Suppose further that, having reached different conclusions on the basis of E , they are then apprised of each other's competing views about p but continue to disagree. The question of what epistemic peers ought to do under such circumstances constitutes the main concern of the epistemology of disagreement.

There are roughly two fundamental positions on the proper epistemic response to a situation involving peer disagreement. According to the so-called 'conciliatory' views, S_1 and S_2 should revise their attitudes towards p , though the extent of such revisions varies depending on how conciliatory the views in question are.¹ In the opposite direction, one finds the so-called 'steadfast' views which require peers to maintain their confidence in their attitudes in the face of

¹ See, for example, Feldman 2006, Christensen 2007 and Elga 2007.

disagreement.² My aim in this paper is not to adjudicate between these views. Rather, I wish to focus on a strategy that the steadfast views standardly invoke in support of their position and against conciliationism. Having closely examined certain representative examples of the steadfast approach, I will try to show that this strategy is fundamentally flawed.

1. Reasonable Disagreement after Full Disclosure: The Questions

Consider again our agents S_1 and S_2 who hold different attitudes towards a proposition p on the basis of the evidence E . Suppose while S_1 believes that p is true, S_2 believes that it is false. The fundamental question is whether they can reasonably maintain their beliefs after becoming fully aware of each other's views. Following Feldman (2006; 2007), it is important to realize that the question of the possibility of reasonable disagreement actually boils down to two questions.

- (a) Can both S_1 and S_2 reasonably maintain their beliefs after learning of the other's opinion?
- (b) Can both S_1 and S_2 reasonably continue to disagree after disclosure while at the same time think that the other's belief is reasonable as well?

To say yes to the first question is to admit that reasonable disagreement after full disclosure is possible, while a positive answer to the second question is tantamount to admitting that it is also possible to have *mutually* recognized disagreement. It is arguable that intellectual humility requires that the steadfast views have the resources to answer the second question in the positive. It is, thus, the question (b) that is going to feature prominently in our discussion of the steadfast views.

1.1. Steadfast Views: The Standard Strategy

As noted earlier, there is a major divide separating positions on the issue of the possibility of reasonable disagreement. There are, on the one hand, the conciliatory views that require epistemic peers to temper their views in the face of disagreement and there are, on the other hand, the steadfast views that allow peers to maintain confidence in their beliefs after disclosure. The steadfast views are, however, initially puzzling for it is difficult to see how peers who are recognized to be roughly equal in terms of intellectual virtues can reach and maintain incompatible, and yet equally reasonable, beliefs on the basis of the same body of evidence. If the peer relationship involves evidential as well as general intellectual equality, how can S_1 and S_2 come to adopt incompatible attitudes towards a particular proposition p ? Moreover, how can they reasonably maintain their respective attitudes when, after full disclosure, they come to be apprised of all their evidence, background beliefs, etc.? What is it that

² See, for example, Kelly 2005, Bergmann 2009 and Goldman 2010.

gives, S_1 , for example, the right to maintain her belief and resist the epistemic pull of S_2 's contrary opinion? The same can be asked of S_2 and her belief.

A natural strategy (henceforth, 'the standard strategy'), for the steadfast view, is to deny the appearance of epistemic symmetry between peers. However, given our assumptions, any putative asymmetry must be characterized in terms other than the agents' general intellectual virtues or their normal body of evidence. It might, for example, be claimed³ that S_1 possesses insights about the disputed proposition that S_2 lacks or that S_1 and S_2 are employing different systems of epistemic norms (Goldman 2010).⁴ Whatever the merits of these suggestions, what is important to note is that no steadfast account of the possibility of reasonable disagreement can come for free. The initial epistemic symmetry between peers has to be broken at some point along the way. If the parties to a dispute are to maintain different attitudes in regard to a particular topic this should be traced back to some prior difference on their part. It is difficult to see how, despite the full epistemic symmetry, peers can rationally arrive at different conclusions in regard to a particular subject matter. To see this, let us consider a recent steadfast account of the possibility of reasonable disagreement due to Marc Moffett (2007).

While agreeing that the most plausible line of defense for a steadfast response to our central question is to appeal to some sort of epistemic difference on the part of the peers, Moffett claims that it is possible to avoid downgrading their competence while maintaining that they can reasonably hold on to their views. The trick is "to take advantage of the transient underdetermination of theory by evidence" (Moffett, 2007, 360). It is not, however, clear why the postulation of an epistemic asymmetry between peers is tantamount to downgrading their competence. To assume that S_1 and S_2 are equal with respect to general intellectual virtues is consistent with the possibility of one of them being more competent or reliable on a particular topic. More importantly, however, the appeal to the underdetermination thesis hardly explains the claim that the same body of evidence (E) can rationalize two contradictory claims p and not- p . If E supports p (to a sufficient degree), then it would be reasonable to believe p on that basis. But this would, in turn, entail that it is not reasonable to believe not- p . It is therefore not clear how the same body of evidence can rationalize two contradictory beliefs. To say that the beliefs (theories) are underdetermined by E is merely to repeat the claim in question. The underdetermination thesis is, in other words, semantically too close to the claim that reasonable disagreement is possible to explain it.

Moreover, assuming that the underdetermination thesis is true, it can, at most, explain how reasonable disagreement is possible before disclosure (i.e., in isolation). It does not tell us why S_1 and S_2 can maintain confidence in their

³ For example, van Inwagen 1996 and Bergmann 2009.

⁴ For a different account of asymmetry in terms of the information one has about oneself and one's opponent see Sosa 2010.

beliefs *after* disclosure. After all, the conciliationists claim that when S_1 and S_2 learn that the other holds a different attitude, on the basis of the same evidence, the awareness of this fact would provide a defeater for their beliefs.

In response, Moffett appeals to Sklar's principle of methodological conservatism (MC) according to which it would be reasonable to hold on to a hypothesis despite coming to know of evidentially equivalent alternatives (Sklar 1975). But (MC) is precisely what the conciliationists are challenging. Unless Moffett can provide some independent justification for (MC), he cannot, on pain of begging the question against the conciliationists, appeal to it to defend his steadfast view. He does, however, make some remarks in defense of (MC) by referring to Harman's (1986) observation that some such conservative principle is indispensable to our epistemic well-being for

if we are not able to rely on our standing beliefs without re-verifying them or – what is cognitively impossible for creatures like us – carrying along our justifications, we will be unable to make justifiable progress in our theoretical undertakings (Moffett 2007, 361).

But Harman's observation actually concerns a different version of the thesis of epistemic conservatism according to which one is justified to continue to hold a belief as long as there are no good reasons against it. What underpins Harman's principle of conservatism is the phenomenon that people do not usually keep track of the justification relations among their beliefs which is presumably why they tend to preserve their beliefs in the face of evidential discrediting.⁵ These considerations hardly lend any support to (MC).

Moffett seems to realize that Harman's principle is not sufficient to establish (MC) and, thus, his version of the steadfast view. He, thus, chooses to highlight, what he takes to be, the implausible skeptical consequences of the conciliatory views (like Feldman's), namely, that they would be 'epistemically devastating' if consistently applied across the board. But this remark is hardly relevant to our concerns here, not least because not all versions of conciliationism have such implications. Moffett's animadversions, thus, seem to lead nowhere. However, he makes one final observation in defense of his view claiming that holding to one's beliefs has certain cognitive advantages in that it helps us develop and refine theoretical frameworks and worldviews.

But are such alleged advantages to be counted as *epistemic* reasons for beliefs? Yes, he thinks, unless one

[construes] the epistemic in an inappropriately narrow manner, as reasons which the belief itself is more likely to be true (Moffett 2007, 363).

But this revisionist step threatens to change the rules of the game, for it is no longer clear that he can continue to characterize an individual x as an

⁵ For an analysis of Harman's version of epistemic conservatism see Vahid 2004.

Some Problems With Steadfast Strategies for Rational Disagreement

epistemic peer with respect to a given domain (if and) only if we regard x's judgments concerning that domain to be as likely as our own to be correct given the same evidence (Moffett 2007, 356).

More importantly, for our purposes here, the appeal to the pragmatic virtues of belief perseverance as reasons for why epistemic peers can hold on to their beliefs in the face of disagreement is to postulate an asymmetry, though not in the space of reasons but, in the space of pragmatic considerations. Just as the appeal to insights might help justify why epistemic peers can reasonably hold on to their beliefs after disclosure, so does the appeal to the purported pragmatic consequences of maintaining such beliefs. Either way, no steadfast view comes for free as, one way or another, it is forced to appeal to factors regarding which epistemic peers diverge. Let us then return to our main question and examine some steadfast views that employ the standard strategy to show how reasonable disagreement is possible. I begin with a recent account (due to Goldman 2010) that locates the required asymmetry in the systems of epistemic norms that are said to license the agents' beliefs.

2. Objectivity-based Relativism as the Basis of Reasonable Disagreement

To show that reasonable disagreement is possible, Goldman appeals to two fundamental theses of epistemic objectivism and epistemic relativism. He takes 'objectivism' to mean that there is a uniquely right epistemic system (E-system) comprising of rules and norms directed at doxastic attitudes. Thus, according to objectivism, for any proposition p , a right E-system determines what attitude someone should adopt towards p on the basis of some evidence e . The thesis of epistemic relativism, on the other hand, construes justification statements as covertly referring to the E-system the speaker accepts. The combination of these elements, Goldman thinks, would provide a solid basis for the possibility of reasonable disagreement. To see this, consider our two agents S_1 and S_2 believing p and not- p respectively on the basis of the same evidence E . Suppose, in accordance with Goldman's objectivity-based-relativism (OBR-) account of justification, S_1 and S_2 operate with two different epistemic systems, E_1 -system and E_2 -system respectively that authorize different attitudes towards p . Since, by hypothesis, there is a uniquely correct E-system, this means that S_1 and S_2 cannot be both objectively (O-)justified in their beliefs. So, at the level of the target proposition p , at least one of those agents is unreasonable or unjustified. Does this mean that the story for the possibility of reasonable disagreement is over? No, for Goldman thinks that his OBR-account of epistemic justification has the resources to show how S_1 and S_2 can reasonably maintain their beliefs on the basis of the same evidence E .

To show this, Goldman starts by claiming that S_1 and S_2 are both O-justified in believing their respective E-systems or individual E-norms. To defend this claim, he asks us to consider communities and cultures where young children and people are instructed by their elders to believe and follow certain

E-norms like "If the Scripture says p , you should believe that p " or "If scientists agree on p , you should assign a high credence to p ." He further assumes that these "children hear roughly the same testimony from numerous elders and no conflicting testimony" (Goldman 2010, 198). Such scenarios, he thinks, provide examples in which an incorrect E-norm is O-justifiably believed to be correct. These norms would, in turn, sanction incompatible attitudes in regard to particular propositions. How does this bear on the question of whether reasonable disagreement is possible? According to Goldman, since S_1 and S_2 are O-justified in believing the E_1 -norm and E_2 -norm respectively, they may also enjoy a distinct but significant justificational status for the target beliefs those norms authorize. So while they may not be first-order O-justified in their beliefs, they are, nonetheless, iteratively O-justified in holding those beliefs, that is, we can say of, say, S_1 that " S_1 is O-justified in believing that she is O-justified in believing p ."

Likewise for S_2 . Although second-order O-justifiedness does not entail first-order O-justifiedness, it does, says Goldman, make a significant contribution towards the reasonability of a first-order belief as he denies that the first-order justificational status of an attitude fixes its overall reasonability. By bringing epistemic norms to bear on the reasonability of the target beliefs, Goldman admits that he is, in effect, claiming that S_1 's and S_2 's total evidence differ. For although they have, what he calls, the same 'material' evidence (E), their norm evidence (E_1/E_2 -norms) differ which is why their attitudes towards p can legitimately diverge. Thus, an OBR-account of epistemic justification, he concludes, can explain how divergent attitudes based on the same body of evidence can be both reasonable.

There are, however, a number of places in Goldman's argument for the possibility of reasonable disagreement where one can take issue with it. By appealing to an OBR-account of justification, Goldman is obviously implementing the standard strategy to defend his version of the steadfast view. He only chooses to construe the required epistemic asymmetry in terms of different epistemic norms that guide the agents' doxastic behavior. Thus, while S_1 adopts the E_1 -norm to arrive at p , S_2 adopts the E_2 -norm to arrive at not- p . To assess Goldman's account of the possibility of reasonable disagreement, I propose to examine it in view of the following questions: (1) Is Goldman's account internally coherent? And (2) does it successfully explain how reasonable disagreement is possible? I begin with the first question.

What makes Goldman say that the epistemic agents (like the children of his example) are O-justified in believing incompatible E-systems or E-norms? Presumably, this is because they adopt their E-norms on the basis of different testimonial evidence coming from their elders in different communities. Consonant with his norm-based account of epistemic justification, Goldman thinks that S_1 's and S_2 's O-justified beliefs in their pertinent E-systems themselves also involve norms (though of a more fundamental kind). These

fundamental norms are source authorizations like the following 'generic testimony-based norm' pertinent to his example involving children: "If a random speaker or writer testifies that p , then (in the absence of defeating conditions) the agent should believe that p ." So there are actually two kinds of evidence involved here that ground S_1 's and S_2 's beliefs in their incompatible E-norms. There is the norm evidence, like the generic testimony-based norm above, and there is the material evidence consisting of the different testimonies that the children obtain in the form of instructions from their elders.

Suppose we agree that, on the basis of their total evidence, S_1 and S_2 are justified in believing that their respective E-norms are correct. On the basis of these norms, S_1 and S_2 then go on to hold different attitudes towards p even though they happen to share the same material evidence E . As we have seen, however, Goldman denies that S_1 and S_2 can be both O-justified in holding different attitudes towards the target proposition p . But it is not clear why S_1 and S_2 can be O-justified at the norm level but not O-justified at the target-belief level. That is, it is not clear why S_1 and S_2 can be O-justified in holding different attitudes about the correct E-system but fail to be O-justified in holding incompatible attitudes towards p . For, on both levels, they have the same type of total evidence consisting of appropriate norm evidence as well material evidence. The only difference seems that, at the norm level, it is the norm evidence that is shared by S_1 and S_2 with their material evidence varying while, at the target-proposition level, it is the material evidence E that is fixed with the norm evidence being different (E_1/E_2 -norms). So if S_1 and S_2 can be both O-justified at the norm level because their total evidence differs, there is no reason why they cannot also be O-justified at the target-proposition level. Indeed, if, as Goldman claims, S_1 and S_2 are both O-justified in believing the correctness of the E_1 -norm and the E_2 -norm respectively, then, obviously enough, they should also be O-justified in believing what these norms authorize. A belief that is grounded in another belief that is O-justified is itself O-justified.

It might be argued that what grounds the distinctness of the epistemic status of S_1 and S_2 at these two levels has to do with the fact that the E_1 -norm and the E_2 -norm cannot both belong to the uniquely correct E-system. At least one of these norms must be incorrect and so must the belief it sanctions. But this point is irrelevant to what is at issue here which is justification not truth. A belief can be false and yet O-justified. What counts for justification at both the norm and the target-proposition levels is one's total evidence which might include misleading evidence. As Goldman himself emphasizes, truth is not necessary for justification:

what determines a belief's reasonability is the agent's evidence...The same point holds on the topic of norm correctness. The actual rightness of an E-system does not 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 (Goldman 2010, 206).

Accordingly, since both S_1 and S_2 have evidence for their adopted E-norms (which is why they are said to be both O-justified in believing them), they are also O-justified in holding the attitudes that those norms license about p . So, they can also be O-justified at the level of the target proposition despite having misleading evidence.

I think Goldman's inconsistent treatment of the agents' beliefs at the norm and the target-proposition levels arises from his expressed wish to combine elements from both objectivism and relativism in his theory of epistemic justification. On the one hand, the objectivist ingredient prompts him to link

justifiedness not to any random E-system but to a *right* E-system, because...[o]nly a right epistemic system has the appropriate connection with objective justifiedness (Goldman 2010, 193)

or reasonability. This, in turn, leads him to deny that S_1 and S_2 can both be O-justified in holding different attitudes at the target-proposition level. The relativist ingredient, on the other hand, impels him to say that the

actual rightness of an E-system does not determine the reasonability of an agent's conforming to it (Goldman 2010, 206)

which is presumably why he claims that S_1 and S_2 are both O-justified in adopting different E-systems. Let us now turn to our main question of whether Goldman's OBR-account of epistemic justification can successfully explain the possibility of reasonable disagreement. We need to address this question both before and after the agents come to know of their respective views. I start with the possibility of reasonable disagreement before full disclosure.

According to Goldman, reasonable disagreement before disclosure is possible because S_1 and S_2 are iteratively, though not first-order, justified in holding different attitudes towards p . This contention is, as we have seen, in turn, grounded in Goldman's claim that the agents are O-justified in believing that their incompatible E-norms are correct. There are, however, a number of ways that this latter claim can be challenged. Recall Goldman's example of the children in isolated communities who adopt incompatible E-norms on the basis of their elders' advice. Now, given Goldman's 'reliabilist criterion of [E-]system superiority' (Goldman 2010, 194), shouldn't those testimonial sources themselves be reliable if the children are to be *objectively* (O-)justified in believing their E-norms? One may concede that the children enjoy some sort of justification for their beliefs but objective justification seems to require that their evidence be, at least, reliably adequate.

Indeed, Goldman's examples resemble the so-called 'epistemic poverty cases' in the philosophical literature which are often invoked to show the inadequacy of deontological justification. To say that a belief p is deontologically justified is to say that in holding that belief, an agent has flouted no (subjective)

epistemic obligations, and is, thus, subject to no blame or disapproval. A standard way to express the thought that an agent is deontologically justified in believing p is to say that he is iteratively justified (in a truth-conducive sense) in holding that belief.⁶ The deontological conception of justification has been criticized for, among other things, failing to give us what we expect of epistemic justification. The thought behind this criticism is often formulated in terms of the so-called 'epistemic poverty cases' where, despite doing all that can be reasonably expected of agents, they form their beliefs on less than adequate grounds. One such case, 'the cultural isolation case' (Alston 1988), involves an agent, S , growing up in an isolated community in which everyone unquestionably accepts the traditions of the tribe as authoritative. Having had no opportunity to come across circumstances in which this authority is challenged, S can hardly be blamed for holding beliefs that are grounded in the traditions. He is thus deontologically justified in holding those beliefs despite the beliefs being unjustified (in a truth-conducive sense).

The epistemic circumstances of S are similar to those of the children in Goldman's scenario who are said to be exposed to "no conflicting testimon[ies]." It, thus, seems unreasonable to describe the children in Goldman's example as having objectively (O-)justified belief in their epistemic norms. Indeed, Goldman himself, when challenged on this issue, goes on to concede that he can be taken to be addressing a conception of "justification [that is] grounded entirely in one's subjective perspective" (Goldman 2010, 199, fn.9). Of course, once we grant that the agents are only iteratively justified in believing their adopted E-norms, it would be quite plausible to think (with Goldman) that they are also iteratively justified in holding beliefs that those E-norms sanction. In this case there would be no inconsistency in the treatment of beliefs at the norm and the target-proposition levels. However, as we can see, this requires giving up on the notion of objectivity that Goldman's account involves.

Another reason for denying objective justifiedness at the norm level is this: Any explanation of the possibility of reasonable disagreement between epistemic peers that helps itself with the assumption that they are O-justified in believing their respective E-norms would be question-begging. For suppose we assume, with Goldman, that S_1 and S_2 are O-justified in believing incompatible E-norms. To say this, however, is to admit that reasonable disagreement is possible. So, on pain of circularity, this assumption cannot be used as a premise (as in Goldman's account) to justify the conclusion that reasonable disagreement is possible. Differently put, Goldman manages to show that reasonable disagreement is possible at the level of the target proposition only by assuming it at another level, namely, the norm level.

Finally, suppose we agree, with Goldman, that S_1 and S_2 are only iteratively justified in holding different beliefs about p . Does this settle the issue of

⁶ See, for example, Alston 1988. Goldman himself uses the term 'weak justification' to describe this conception of justification. See, for example Goldman 1988.

reasonable disagreement before closure? No, because the notion of iterative justifiedness is too weak. Indeed, one could already acknowledge this fact about the agents without buying into Goldman's conceptual framework. To say, for example, that S_1 is iteratively justified in believing p is, as we have seen, to say that she is deontologically justified in holding that belief and that she has flouted none of her epistemic obligations in forming that belief. In other words, it is to say that her belief is blameless and responsibly formed. But this is something one could already acknowledge about such agents having assumed that they are rational peers.

Let us now see how Goldman's observations bear on the question of the possibility of disagreement *after* full disclosure. Hardly at all, it seems. Indeed, by the end of his discussion, Goldman confesses that his approach is only intended to address the issue of reasonable disagreement in isolation (or 'synchronic disagreement' as he calls it). It is not concerned with the 'diachronic' question, namely, with how an agent should epistemically react when, having formed her belief at t , she comes to know of her peer's contrary opinion at a later time t' . I end my discussion of Goldman, however, by arguing that not only his account is silent about whether one can reasonably maintain her belief after disclosure, it *cannot* recognize such a possibility.

To see this, we may recall that, to argue for the possibility of reasonable disagreement at the level of the target proposition, Goldman needed the premise that it is possible for agents to be O-justified at the norm level. The argument for *this* premise, on the other hand, consisted of Goldman's presenting a case, "a likely scenario in many communities both historical and contemporary" (Goldman 2010, 198) where, he claims, it is easy to see how agents can be O-justified in believing incompatible E-norms. The case involved children living in communities where they follow their elders and teachers about certain matters without much critical discussion and without being exposed to "conflicting testimon[ies]" (Goldman 2010, 198). What this implies is that intellectual or epistemic isolation is necessary if Goldman's scenario is to give him the premise he wants. However, with such a requirement in force, Goldman's agents can hardly proceed to a state of full disclosure. For, after disclosure, when their intellectual isolation ends, they will no longer be entitled to having O-justified beliefs in their E-norms, and, thus, there would be nothing that would authorize them to have different attitudes towards the target proposition p . After full disclosure, there will be no opportunity to maintain reasonable disagreement. In view of its structure, Goldman's account of reasonable disagreement is bound to crumble in the state of full disclosure. I shall now turn to another defense of the steadfast view (due to Bergmann 2009) that also appeals to the standard strategy.

3. Insights and Error Theories as the Bases of Reasonable Disagreement

Right from the start, Bergmann states that his target is to explain how reasonable disagreement is possible after full disclosure. Following Plantinga, he distinguishes between two kinds of rationality, internal and external rationality. A belief is internally rational if and only if it is an epistemically appropriate response to the agent's evidence (consisting of her mental states). Internal rationality, thus, concerns the epistemic standing of a belief downstream from experience. One might, thus, say of the beliefs of a brain a in a vat (BIV) or a victim of the Cartesian demon that they are internally rational or justified. External rationality, on the other hand, requires that, in addition, the agent's cognitive mechanism be functioning properly. This means that since the BIV's experiential evidence is due to malfunction, her beliefs are not externally rational.

Bergmann then appeals to the standard strategy to show that, whichever conception of rationality is in place, mutually recognized reasonable disagreement is possible. Unlike the previous steadfast views, he does not construe the required epistemic asymmetry in terms of pragmatic considerations or epistemic norms. Rather, he unpacks it in terms of the insights that agents might possess. Thus, we might say of two agents, S_1 and S_2 , that they differ not only with respect to their attitudes towards p but also in regard to, what Bergmann calls, their broader outlooks O_1 and O_2 which he takes to include the following ingredients.

O_1 • p

- a theory of error (applied to epistemic peers who believe the key ingredients of O_2) according to which the apparent insight that underpins the key ingredients of O_2 is not a genuine insight

O_2 • not- p

- a theory of error (applied to epistemic peers who believe the key ingredients of O_1) according to which the apparent insight that underpins the key ingredients of O_1 is not a genuine insight

So before S_1 learns of S_2 's contrary view, she has apparent insights that both p and the error theory in O_1 are true. The right response to these insights is to believe p and the associated error theory. Thus, both these beliefs are, by definition, internally rational. We can further stipulate, says Bergmann, that, given the strength of these apparent insights, the beliefs in question are very strongly justified. This takes care of the question whether the parties to the dispute can rationally disagree before disclosure. But what happens after S_1 learns of her disagreement with S_2 (whom S_1 takes to be an epistemic peer)? Let us say that the only piece of evidence that is added to S_1 's body of evidence after disclosure is S_2 's report of her apparent insight in support of her belief that not- p . Does S_1 's recognition of this fact count as a defeater for S_1 's own belief that p ?

No, says Bergmann, because, given S_1 's rational high confidence in her error theory for those peers (like S_2) who reject p , S_1 can rationally hold on to her belief: "As a result S_1 remains internally rational in her belief that p despite her recognition that S_2 ...believes not- p " (Bergmann 2009, 340). The same thing can, of course, be said of S_2 's belief that not- p . Bergmann further claims that these remarks also explain why S_1 and S_2 can both recognize that the other is internally rational in holding her respective belief. So mutually recognized reasonable disagreement is also possible.

Let us now examine to what extent these observations succeed in explaining the possibility of reasonable disagreement after disclosure. I shall focus on the notion of internal rationality that Bergmann takes to be the more fundamental notion and, in fact, equivalent to justification. So, is reasonable disagreement possible when it is internal rationality that constitutes the epistemic status of beliefs? Again, the full answer to this question requires not only the investigation of the possibility of reasonable disagreement after full disclosure but also the possibility of the mutually recognized variety.

Recall that a belief p is said to be internally rational if and only if it is an appropriate epistemic response to the agent's evidence (mental states) regardless of whether or not this experiential evidence is the result of cognitive malfunction on the part of the agent's cognitive mechanisms. Let us concede that S_1 's and S_2 's beliefs are internally rational before disclosure. The question is whether they also remain internally rational after full disclosure. After disclosure, S_1 's and S_2 's bodies of evidence expand to take into account what each learns about the other's view about p . S_1 then has, in her body of evidence, not only her insight for the truth of p as well as the associated error theory but also the belief that S_2 has an apparent insight that not- p . According to Bergmann, given the strength of S_1 's apparent insight for the belief p and for her error theory, her recognition of S_2 's insight for not- p will fail to undermine her convictions and so she will remain internally rational in believing that p . In other words, S_1 's belief that p remains an epistemically appropriate response to her mental states after full disclosure. Likewise for S_2 's belief that not- p . So reasonable disagreement after full disclosure is possible.

There is, however, a slight complication here. To argue for this conclusion, Bergmann needs to assume that S_1 and S_2 maintain 'rational high confidence in their respective error theories' after full disclosure. For it is only because they believe their respective error theories with such high confidence that each is able to resist the negative epistemic influence of her newly found evidence that the other holds a strong insight for an incompatible belief and, thus, remain internally rational in her attitude towards p :

[T]hese beliefs – of S_1 in...the error theory of O_1 and of S_2 in...the error theory in O_2 are each partially based on apparent insight that the propositional content of the belief so based is true (Bergmann 2009, 339).

So S_1 's and S_2 's beliefs in their respective error theories are internally rational. They are, on the other hand, incompatible as they pass different judgments on what counts as a genuine insight about p . Taken together, what these remarks suggest is that S_1 and S_2 can reasonably disagree over what content a genuine insight about p must have. This means that, as with Goldman's account, it is only by assuming that reasonable disagreement is possible at one level (the level of error theories) that Bergmann is able to show that it possible at the level of the target proposition p .

Of course, Bergmann does not deny that sometimes recognized disagreement can provide one with a defeater of one's belief. Under such circumstances reasonable disagreement is not possible as peers are forced to temper their views. But what are these circumstances? According to Bergmann, these are circumstances in which one *should* disbelieve or seriously doubt that (C) "one is, on this occasion, more trustworthy than one's peer who holds a different attitude towards p ." However, he refrains from giving a general recipe as to when an attitude of disbelief or doubt is the epistemically appropriate response to one's mental states (after full disclosure). Each case, he says, should be examined individually.

But it is not clear why the recognition of peer disagreement should be thought to undermine, say, S_1 's or S_2 's confidence in (C), that is, in believing that each is more trustworthy than the other with respect to p . (C) is what the error theories in their outlooks entail. When, for example, S_1 believes on the basis of her error theory that the insight supporting S_2 's belief is not genuine, she would consequently take herself to be more trustworthy than S_2 on the topic of p . Given Bergmann's own arguments, S_1 's recognition of S_2 's disagreement about p should not undermine her belief in her comparative reliability. If S_1 is to be vulnerable to S_2 's adverse influence, one should make sure that she is no longer protected by her error theory. Once she loses her epistemic immunity, her beliefs will become vulnerable to the adverse influence of peer disagreement. But to deny that she is entitled to her error theory in such circumstances is to hold that peers cannot reasonably hold on to their beliefs in their incompatible error theories under such conditions. Once again, one has to assume, though in the opposite direction, that peer disagreement is *not* possible at one level (the level of error theories) to show that it is *not* possible at another level (the level of the target proposition).⁷

⁷ Interestingly, this seems to be what happens in Feldman's 2006 conciliationist argument against the possibility of reasonable disagreement. Feldman denies that one can appeal to the different insights that the agents might have to explain the possibility of reasonable disagreement after full disclosure. If, prior to disclosure, S_1 had good reason to think her belief that p is well supported, after disclosure, says Feldman, S_1 's expanded body of evidence would no longer support that belief. Rather, "it makes suspending judgment on this matter [' S_1 's belief that p is well supported'] the reasonable attitude" (232). It follows, Feldman claims, that the expanded body of evidence would no longer support p either because "if still does support p , then it supports [S_1] reasonably having a complex attitude that she would express as

Before I turn to, what I take to be, a fundamental problem with Bergmann's defense of the steadfast view, it would be worthwhile to note another dimension with respect to which his account resembles Goldman's. We noted that, according to Goldman, S_1 and S_2 can reasonably disagree with each other because, though not first-order justified, they are nevertheless iteratively justified in holding different attitudes towards the target proposition p . This does not seem to be very different from what Bergmann claims about the internal rationality of those agents after disclosure. Although a great deal needs to be said about the notion of internal rationality and how it differs from external rationality, Bergmann's sketchy remarks and his examples can still give us a rough idea of how the distinction is to be understood. An internally rational belief, he says, has to do with what goes on in belief formation 'downstream from experience' regardless of whether or not the experience itself is due to cognitive malfunction. He cites, as examples of internally rational beliefs, the beliefs of the victims of an evil demon or a superscientist. As we have seen, however, such beliefs are often construed as involving iteratively justified beliefs (involving a truth-conducive conception of epistemic justification). Goldman himself calls such beliefs 'weakly,' as opposed to 'strongly,' justified.⁸ Weak justification does not imply strong justification just as, according to Bergmann, internal rationality is distinct from external rationality. So, at most, what follows from Bergmann's account is that S_1 's and S_2 's beliefs are, after disclosure, responsibly formed and non-culpable. But, as with Goldman's account, this is a fact that one could already acknowledge about S_1 and S_2 , having recognized them as rational and responsible peers.

Finally, there is something about Bergmann's way of implementing the standard strategy that threatens to violate, what I take to be an adequacy condition on an acceptable theory of reasonable disagreement. The idea is that an epistemology of disagreement must be sensitive to the distinction between reasonable disagreement in isolation and reasonable disagreement after full disclosure. Suppose S_1 comes to believe that p as it seems to her that her evidence supports p . Let us further assume that since S_1 has carefully examined her body of evidence and there is no countervailing evidence against p , S_1 is justified in believing that p . Being aware of her fallibility and her own history belief formation, S_1 can obviously conceive of others contradicting her belief on the basis of the same body of evidence. The mere conceivability of disagreement, however, is no bar to S_1 's maintaining her belief if the belief has been responsibly formed. We may thus grant the rationality of S_1 's belief in isolation. Suppose now S_1 learns that S_2 has come to believe not- p on the basis of the same body of

follows: I believe p , but I suspend judgment on whether my evidence supports p " (232). Whatever the merits of Feldman's argument, it shows the impossibility of reasonable disagreement at the level of the target proposition only by assuming it another level, namely, for the belief that the belief that p is well supported.

⁸ See footnote 6.

evidence. S_1 realizes that at least one of them is in error in their evaluation of the evidence. Who is to blame? Given that the interesting cases of disagreement are those where it is not obvious which belief the evidence really supports, S_1 and S_2 are in a symmetrical position in so far as the question of blame is concerned.⁹ It would not thus be acceptable if S_1 were to think that the error lies with S_2 (simply because it is her own belief that is at stake) without *explaining* why this is so.

To condone S_1 's attitude is to fail to attach any epistemic significance to the distinction between reasonable disagreement in isolation and reasonable disagreement after full disclosure. The central question in the epistemology of disagreement is what one should do with one's own belief when one learns of an epistemic peer's contrary opinion. So I take it that any steadfast view of reasonable disagreement that claims that one can reasonably maintain one's belief in the face of disagreement merely on the a priori ground that views incompatible with one's own are flawed and ought to be ignored is violating an adequacy condition on an acceptable theory of reasonable disagreement. In other words, any account of why it is reasonable to stick to one's view after full disclosure which rules out the information about peer disagreement as relevant merely on the ground that views contrary to one's own are in error is inadequate. This conclusion can be reinforced by noting that this way of securing reasonable disagreement would also undermine the epistemic significance between possible and actual disagreement. For what is of importance for the question of the possibility of reasonable disagreement is actual, rather than, possible peer disagreement. As noted above, the fact that one's belief might possibly be contradicted by similarly intelligent and intellectually virtuous people merely reflects our fallible and imperfect epistemic predicament. Now, it seems to me that that, by incorporating error theories in his theory of epistemic justification (rationality), Bergmann's defense of the steadfast view violates the aforementioned adequacy condition on an acceptable theory of reasonable disagreement.

To clarify further, suppose, having argued for the view that free will is incompatible with determinism, van Inwagen learns of David Lewis's contrary opinion. *After* carefully examining Lewis's reasons, he decides that Lewis is in error. To explain, he claims that he has insights that are denied to Lewis. Now, while this might be, methodologically speaking, a legitimate thing to do if one wishes to defend a steadfast position regarding a particular dispute, it raises problems of the sorts mentioned above when it is turned into a general strategy for defending the steadfast views. For to show that agents can reasonably maintain their views in the face of disagreement, we must make sure that their cognitive systems are armed, as it were, with something like defense mechanisms that are, *ceteris paribus*, automatically activated whenever the agents' beliefs become vulnerable to adverse epistemic influence as in the peer

⁹ See Feldman 2006.

disagreement cases. Such built-in defense mechanisms tend to epistemically immunize the agents against possible threats to their cognitive systems. This way an agent will be allowed to maintain her belief whenever she is in an epistemically hostile situation. Indeed, given such built-in defense mechanisms, the actual cases of disagreement would be as innocuous for one's views as are the possible ones.

Now, this seems to be what Bergmann is doing with his postulation of error theories as essential ingredients of the agents' outlooks. These error theories are formulated in terms of the contrary views of the agents' peers. For example, S_1 's error theory rejects as not genuine the insight of any possible peer who believes not- p and who subscribes to another incompatible error theory. Thus, when S_1 learns that there *is* an epistemic peer, S_2 , who actually believes the ingredients of O_2 , her cognitive defense mechanism automatically starts to kick in to protect her against the adverse epistemic influence of S_2 . Although, S_1 's error theory does not particularly target S_2 , it is intended to apply, as we have seen, to "those roughly equal in intellectual virtue who...reject p " (Bergmann 2009, 339). This is why, the news of the S_2 's disagreement would hardly catch S_1 by surprise since given " S_1 's rational high confidence in her theory of error," her recognition of S_2 's belief that not- p would not count as a "fairly powerful potential defeater" for S_1 's belief that p (Bergmann 2009, 340).¹⁰ S_1 's error theory would entitle S_1 to her belief at any time regardless of whether or not she has learned of any peer disagreement. Accordingly, on Bergmann's steadfast view, reasonable disagreement in isolation and reasonable disagreement after disclosure are rendered epistemically on a par. In so far as the agents are able to maintain their confidence in their error theories they need not worry about whether or not others disagree with them.

Another factor that ensures the irrelevance of the distinction between reasonable disagreement in isolation and reasonable disagreement after disclosure, on Bergmann's account, concerns the type of evidence that he takes to be germane to the rationality of one's beliefs. He takes insights to be what support the key ingredients of the agents' outlooks. Insights are, however, not the sort of things that can be fully shared:

[E]ven after full disclosure, S_1 and S_2 do not have the same evidence. Reporting their apparent insights to each other is not the same as giving those apparent insights to each other (Bergmann 2009, 339).

So, given that all the new evidence that S_1 and S_2 might obtain after full disclosure are the *reports* of their insights, the strength of their evidential situations could hardly change after they learn of each other's views. This would

¹⁰ I think the same problem afflicts Goldman's OBR-account of reasonable disagreement. Once epistemic justification is taken to be norm-relative, agents who subscribe to a particular set of E-norms can safely ignore the views of those peers who adopt different E-norms and regard them as unjustified.

provide another reason why, on Bergmann's account, it hardly matters for determining the reasonability an agent's' view whether she is considered in isolation or after exposure to peer disagreement. I conclude therefore that Bergman's account fails to respect an adequacy condition on acceptable theories of reasonable disagreement.¹¹

Conclusion

We have examined a number of steadfast views about how one should epistemically behave in the face of peer disagreement. Despite their differences, all such views employ, what I have called, the standard strategy. Reflecting on these accounts, a certain pattern emerges which explains why they are susceptible to the same problems. It will also help illuminate the conciliationists' reactions to such views. Let me explain.

To show how epistemic peers can reasonably maintain their beliefs after full disclosure, the steadfast views, employing the standard strategy, tend to postulate some sort of epistemic asymmetry between the parties involved in a dispute. We have seen that different steadfast views interpret the asymmetry differently. Some construe it in terms of pragmatic considerations (Moffett) and some in terms of different epistemic norms (Goldman) while others appeal to different error theories (Bergmann). Thus, it seems that, to explain the rationality of the agents' first-order beliefs about a particular proposition *p* in the case of a disagreement, we have to go one level up and assume that the agents have certain fundamental beliefs about, say, which epistemic norms are correct and justification-conferring or which evidence or insights are genuine. For Goldman it was the postulation of the objectively justified beliefs about the correct E-norms that grounded incompatible but rational attitudes towards the target proposition while, for Bergmann, it was the agents' rational beliefs about the pertinent error theories that allowed them to rationally maintain their beliefs after full disclosure. Despite its clear relevance to the question of the possibility of reasonable disagreement, the standard strategy is apt to make the steadfast views that employ it vulnerable to the following general problems.

1. Although the postulation of certain antecedent rational beliefs, on the part of the epistemic peers, is, as we have just seen, necessary to explain how

¹¹ At the end of his discussion, Bergmann makes some remarks which seem to suggest that he is not willing to attach much epistemic significance to the distinction between possible and actual disagreement. First, he points out, correctly, that the fact that it is possible for things to seem to us perceptually just as they do even if we are the victims of an evil demon does not imply that we should question the reliability of our perceptual beliefs. Likewise, he claim, the fact that we think our peer is mistaken, despite being equally confident as we are in our beliefs, does not justify revising those beliefs. He realizes, however, that while his first case concerns the mere possibility of skepticism, the second case concerns an actual scenario in which we are being challenged. He thinks, however, that "the lesson is the same nevertheless" (Bergmann 2009, 545).

they can rationally disagree about a particular proposition, it makes the steadfast views susceptible to the charge of begging the question against conciliationism. For to assume that agents rationally hold such fundamental and incompatible beliefs is to assume that rational disagreement is possible. We saw that, on Goldman's account, agents are O-justified in believing incompatible E-norms to be correct while, for Bergmann, they enjoy rational confidence in different error theories. So both accounts try to show that rational disagreement is possible at the level of the target proposition only by assuming that it is possible at another (more fundamental) level.

2. Another problem for the standard strategy is whether it can establish the possibility of *both* reasonable disagreement and mutually recognized reasonable disagreement after full disclosure. For what seems necessary for the possibility of reasonable disagreement seems to render mutually recognized reasonable disagreement impossible. For example, while, for Bergmann, the assumption that agents hold different error theories is necessary in order to explain how they can remain rational after full disclosure, that very same assumption seems to hinder the possibility of *mutually recognized* peer disagreement. This is, of course, to be expected if an agent is to remain justified in believing her error theory after full disclosure. What is, at most, mutually recognizable by epistemic peers is that the other enjoys iterated justification in her belief about the target proposition. The same thing is true about Goldman's OBR-account which explicitly endorses only iterated justification for the target proposition. Regardless of whether such an attenuated notion is sufficient to establish the possibility of reasonable disagreement, the problem with this suggestion is that one could already acknowledge its tenability without advertent to the various epistemological frameworks promulgated by the steadfast views such as Goldman's or Bergmann's. Iteratively justified belief, as non-culpable belief that is responsibly formed, is what is to be expected from agents who recognize each other as rational peers.

3. Finally, the implementation of the standard strategy is apt to lead to the violation of an adequacy condition on an acceptable theory of reasonable disagreement, namely, the recognition of the epistemological significance of the distinction between reasonable disagreement in isolation and reasonable disagreement after disclosure. We saw that once the standard strategy is used at the service of explaining the reasonability of disagreement it is bound to turn into a defense mechanism that would eventually obviate the epistemological significance of the aforementioned distinction.

For these reasons, I conclude that the steadfast views had better look into options other than adopting the standard strategy.

References

- Alston, William P. 1988. "The Deontological Conception of Epistemic Justification." Reprinted in his *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca: Cornell University Press, 1989.
- Bergmann, Michael. 2009. "Rational Disagreement after Full Disclosure." *Episteme: A Journal of Social Epistemology* 6, 3: 336-353.
- Christensen, David. 2007. "Epistemology of Disagreement: the Good News." *The Philosophical Review* 116: 187-217.
- Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41, 3: 478-502.
- Feldman, Richard. 2006. "Epistemological Puzzles about Disagreement." In *Epistemology Futures*, ed. Stephen Hetherington, 216-236. New York: Oxford University Press.
- Feldman, Richard. 2007. "Reasonable Religious Disagreements." In *Philosophers without Gods: Meditations on Atheism and the Secular Life*, ed. Louise M. Antony, 194-214. Oxford: Oxford University Press.
- Goldman, Alvin I. 1988. "Strong and Weak Justification." In *Philosophical Perspectives*, Volume 13, ed. James E. Tomberlin. Atascadero, CA: Ridgeview.
- _____. 2010. "Epistemic Relativism and Reasonable Disagreement." In *Disagreement*, eds. Richard Feldman and Ted Warfield. New York: Oxford University Press.
- Harman, Gilbert. 1986. *Change in View: Principles of Reasoning*. Cambridge: The MIT Press.
- Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*, vol. 1, eds. Tamar Szabo Gendler and John Hawthorne, 167-96. New York: Oxford University Press.
- Moffett, Marc A. 2007. "Reasonable Disagreement and Rational Group Inquiry." *Episteme: A Journal of Social Epistemology* 4, 3: 352-367.
- Sklar, Lawrence. 1975. "Methodological Conservatism." *Philosophical Review* LXXIV: 186-191.
- Sosa, Ernest. 2010. "The Epistemology of Disagreement." In *Social Epistemology*, eds. Adrian Haddock, Alan Millar and Duncan Pritchard. Oxford: Oxford University Press.
- Vahid, Hamid. 2004. "Varieties of Epistemic Conservatism." *Synthese* 141: 97-122.