

REASONABLE DISAGREEMENT AND RATIONAL GROUP INQUIRY

Marc A. Moffett

(penultimate draft—forthcoming in *Episteme: A Journal of Social Epistemology*)

ABSTRACT. According to one widely held view, a belief is fully justified only if it holds up against the strongest available counterarguments and we can be appropriately confident that it does hold up only if there is free and open critical discussion of those beliefs amongst our epistemic peers. In this paper I argue that this common picture of ideal rational group inquiry interacts with epistemic problems concerning reasonable disagreement in a way that makes those problems particularly difficult to resolve. By focusing on this idealized context, we get a clearer picture of the epistemic principles at issue. In the end, I argue that the best way to resolve the resulting epistemic conflicts is by appeal to the underdetermination of theory by evidence together with a principle of epistemic conservatism.

There is a widely held view in both philosophy and science about the relationship between individual knowledge and justification and what might be called rational group inquiry. In his essay *On Liberty*, Mill (1956 [1859]) does a respectable job of articulating the theory:

Complete liberty of contradicting and disproving our opinion is the very condition which justifies us in assuming its truth for the purposes of action; and on no other terms can a being with human faculties have any rational assurance of being right (24).

And again later he adds:

Nor is it enough that he should hear the arguments of adversaries from his own teachers, presented as they state them, and accompanied by what they offer as refutations. That is not the way to do justice to the arguments or bring them into real contact with his own mind. He must be able to hear them from persons who actually believe them, who defend them in earnest and do their very utmost for them. He must know them in their most plausible and persuasive form... (45)

The Millian theory of rational group inquiry goes something like this: (1) x 's theoretical beliefs are fully, perhaps even adequately, justified at a given time t only if they hold up against the strongest counterarguments practically "available" at that time; (2) x can be adequately justified in thinking that x 's current theoretical beliefs hold up against the strongest counterarguments practically available at t only if there is free and open critical discussion of those beliefs amongst x 's epistemic peers with whom x genuinely disagrees.¹ I will call the first of these principles the *Millian Platitude* and the second the *Collective Criticism Condition* (or, [CCC] for short). Taken together, the Millian platitude and the CCC suggest that we can be justified in believing that our theoretical beliefs are adequately justified only if we engage in mutual critical discussion with a community of experts at least some of whom hold views that conflict with our own.

¹ The scope of both theses is explicitly restricted to theoretical beliefs. Without trying to provide a fully adequate definition, the intention here is that theoretical beliefs are those that are non-deductively inferred from other beliefs. I take it that the property of *being theoretical* is a gradable property so that beliefs can count as more or less theoretical depending on the nature of their evidential support.

I have two preliminary goals. In §1, I argue that the Millian platitude entails the Collective Criticism Condition, at least for creatures with our cognitive limitations (i.e., creatures having finite cognitive resources, processing limitations, etc.). This claim is *prima facie* quite plausible and interesting insofar as it establishes that there exists a “deep” internal connection between individual rational investigation and (in Goldman’s phrase) social epistemics. My second goal is to show that the view appears to lead to troubling, even paradoxical, results. Specifically, I will argue that on plausible (though, I think ultimately unacceptable) assumptions the [CCC] leads directly to generalized skepticism about our theoretical beliefs because it forces us to confront the problem of reasonable disagreement (e.g., Feldman 2006; Elga forthcoming) in a particularly intransigent form. The problem here is that arguably one should give equal weight to the contrary views of one’s epistemic peers and this implies that we should become agnostic about the overwhelming majority of our substantive theoretical beliefs. But given the argument of §1, this means that the Millian platitude itself leads to such skepticism. Thus, to the extent that our epistemic situation approximates the conditions required by ideal rational inquiry, we will effectively undermine what is presumably the goal of such inquiry, namely, individual-level theoretical knowledge. If we are to avoid skepticism, we must find a way to reconcile the competency of our peer critics with what we take to be persistent and systematic error on their part.

In the final section, I explore a way out of the paradox by exploiting the widely accepted thesis that evidence rationally underdetermines theory in order to make room for epistemically blameless variation in belief. Specifically, if evidence rationally underdetermines theory, we are not driven to the unhappy dilemma of either abandoning our own views or downgrading the epistemic competency of our critics. In this way, underdetermination appears to allow for blameless (and so, reasonable) disagreement. The result is, I believe, a plausible picture of how critical, but non-skeptical, group inquiry is possible.

I. The Millian Theory of Rational Group Inquiry

According to the Millian theory, the justification of our theoretical beliefs depends not only on the positive evidence we can muster in its favor, but also on the extent and quality of our *critical* evaluation of those beliefs. Thus, the Millian theory is centered around the following apparent platitude of ideal rational inquiry:

[MP] x ’s theoretical beliefs are fully, perhaps even adequately, justified at t only if x is able to defend those beliefs against the strongest counterarguments practically “available” at that time.

We need not make the strong claim that x cannot be adequately justified unless x is able to defend her belief against the strongest counterarguments available, though it is not entirely implausible to do so. For present purposes, the important point is that it is *more epistemically optimal* to be able to defend one’s beliefs against the strongest available counterarguments than not to be able to so defend them. Obviously, if one is aware of a certain counterargument but cannot adequately respond to it, this is an epistemic problem. But it is also epistemically undesirable to simply fail to be aware of the counterarguments, quite independently of

whether or not one could respond if made aware of those arguments. Avoiding or otherwise failing to engage in critical appraisal of one's beliefs is *not* a means to becoming better justified in those beliefs!

Given this, it is crucial that we get clear on what constitutes "the strongest counterarguments practically available" at a given time. I take it to be obvious that it doesn't mean simply those counterarguments that one has considered or those counterarguments of which one is aware. Otherwise, one could (absurdly) improve one's epistemic situation by simply failing to consider counterarguments or by remaining ignorant of those on offer. Thus, the notion of availability at play here seems to be an externalist one of some stripe. But even if this is correct, it can't simply be identified with those theories that our community happens to have hit upon, for some communities might simply be epistemically unfavorable. This could happen if political laws or social mores were to generally stifle the intellectual pursuit of certain lines of thought. This, of course, is what prompted Mill's remarks concerning freedom of expression. It could also happen if one's community happened to have an unfavorable distribution of appropriately keen and engaged minds.

At least to a first approximation, the best gloss seems to be something along the following lines: a counterargument is available at time t iff either (i) x is aware of the counterargument at t , or (ii) it is one which it would be reasonable to expect creatures of our x 's abilities to come up with relative to the current overall state of our philosophical and scientific theories at t under circumstances of free intellectual pursuit. I will suppose that this, or something very like it, is correct.²

But given that the notion of availability is in this way externalist, it follows that the Millian platitude [MP] entails an externalist conception of justification. For even if one is able to defend one's theoretical beliefs against the strongest counterarguments practically "available" at a given time, one may not know or even have reason to believe that one is able to do this. After all, one may not know or even have reason to believe that one has considered the strongest counterarguments practically available.

However one feels about epistemic externalism generally, this situation is surely less than ideal. The problem isn't simply the common internalist complaint that the justification of one's theoretical beliefs is "out of one's control", but rather that the conditions under which we are prone to fail are disconcertingly easy to realize. In particular, it is empirically plausible claim that even scrupulously honest researchers are sufficiently biased in their assessment of their own views that they cannot be expected to reliably generate the most telling objections to those views. Moreover, it is plausible that someone who approaches a given issue from a different theoretical perspective is likely to spot deficiencies that an advocate of the view will miss and that individuals with whom we disagree are more likely to have such a difference in approach. Thus, if we simply rely on our own ability to dream up counterarguments to our views, we might reasonably suspect that we do not satisfy [MP]; as a result, we will have reason to suspect that we are not (adequately or at least fully) justified in our beliefs.

² If contextualism is correct, then it might be that this definition requires qualification; for instance, it might be that we need to add on various *sotto voce* clauses which tell us what counterarguments we are properly ignoring (Lewis 1996).

Perhaps the only way for creatures such as us to reliably avoid these sorts of pitfalls is to rely on the critical evaluation of our theories by others, specifically those with whom we genuinely disagree on the very issue in question. For these critics are the ones who will articulate the counterarguments in “their most plausible and persuasive form,” and will not themselves be open to the very biases which we are attempting to avoid. Given this, it appears that for creatures like us who possess limited cognitive resources and who are prone to theory-bias, [MP] implies a general connection between individual epistemology and social epistemics (Goldman 1986; 1987). This is precisely what Mill claims.

In sum: from the point of view of individual epistemology, it is an epistemic good and a practical necessity to rely on a network of genuine critics.

Moreover, the point of having a group of critics would be largely lost if we had reason to suspect that the critics themselves were relatively poor. Indeed, it appears that in the ideal case our critics would be our epistemic peers, at least with respect to the domain in question. After all, it is extremely hard to see what epistemic advantage there could be in having *worse* rather than better critics. Of course, there would be epistemic advantages to having one’s epistemic superiors as critics (though there might also be epistemic downfalls); but it seems that the most we should require is that one’s critics are one’s epistemic peers. If not, we would effectively be forced to say that being fully justified at a given time requires being able to respond to the best *possible* counterarguments, as opposed to merely the best counterarguments practically available at that time. Clearly this is far too demanding a constraint. After all, we are surely sometimes fully justified in accepting a theory or principle which is later shown to be open to insurmountable objections. For instance, it is quite plausible to claim that Frege was fully justified in accepting the naïve comprehension schema even though this was later shown to lead to paradox. But if full justification at a given time required that we be able to respond to the best possible counterarguments, such cases would be impossible.

Thus, on very plausible assumptions about our actual cognitive situation, Mill’s first principle appears to imply the second, the Collective Criticism Condition:

[CCC] x can be adequately justified in thinking that x ’s current theoretical beliefs hold up against the strongest counterarguments practically available at t only if there is free and open critical discussion of those beliefs amongst x ’s epistemic peers with whom x genuinely disagrees.

One way of thinking about [CCC] is as follows: the best way for us to ensure that we have adequately satisfied principle [MP] is to present our work to a large community of peer-critics who differ from us as widely as would be reasonable in their fundamental assumptions. The counterarguments that such a community generate are a reasonable measure of (and might even be identified with) the set of counterarguments “which it would be reasonable to expect creatures of our cognitive abilities to come up with relative to the current overall state of our philosophical and scientific theories at t under circumstances of free intellectual pursuit.”

It follows that in the absence of epistemic peers with whom x reasonably disagrees, x 's epistemic situation will be significantly compromised. At the very least, the lack of such a community would generate significant higher-order uncertainties about one's epistemic situation.³ In this sense, the "Millian context" (as I will call it) appears to be epistemically ideal. We therefore have the following result:

- For most of x 's theoretical beliefs, if x is to be fully or ideally justified in those beliefs, then x must have a community of epistemic peers with whom x genuinely disagrees on those beliefs.

II. The Problem with Peers

To this point, I have argued in that, given certain quite plausible background assumptions about our cognitive situation, [MP] implies [CCC]. Unfortunately, however, [CCC] appears to have very untoward epistemic consequences. The problem, in a nutshell, is that the existence of epistemic peers with whom we genuinely disagree appears to undermine our claims to know those propositions on which we disagree (Feldman 2006; Elga forthcoming). But if this is correct, then the very conditions of ideal rational inquiry undermine what is presumably the purpose of such inquiry, namely, individual-level knowledge. Or, to formulate the point somewhat more paradoxically: wide-spread skepticism concerning our theoretical beliefs is the natural consequence of the conditions required for ideal rational inquiry!

This would be a surprising result to say the least. So why believe that the existence of individuals with whom we (rightly) take to be our epistemic peers undermines our knowledge claims? Here is one argument, due to Elga (forthcoming). Intuitively, we take an individual x to be an 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. Something like this is a very plausible account of the notion of an epistemic peer. For suppose we use a different criterion which does not imply this one. For instance, we might try to define the concept of an epistemic peer in terms of general intelligence or general level of expertise or both, but nevertheless maintain that this criterion does not imply an equal likelihood of being correct given the same evidence. Then there must be some further truth-conducive factor F which we believe is relevant to judgments in this domain and we must believe that we are superior to x with respect to F . But this is counterintuitive. If we genuinely believe that F is a truth-conducive factor in judgment in a given domain and we genuinely believe that we are superior to x with respect to F , then we do not regard x as an epistemic peer in the relevant domain.

³ There are, of course, a number of theoretical beliefs for which our evidence is so strong that lack of dissent does not generate higher-order uncertainties (e.g., the shape of the earth or the function of the heart). The existence of such beliefs shows that there are important qualifications to the Millian theory need to be made, though I will not attempt to formulate those qualifications here.

Despite such qualifications, however, I believe that it is quite appropriate to characterize the resulting skepticism as wide-spread or even massive. For one thing, it is easy to *overestimate* the extent of uncontested scientific agreement under ideal conditions because much of the agreement we actually witness is not grounded in fully rational considerations, but is often based (implicitly, of course) on prudential concerns (e.g., funding and employment issues), educational structure, and the culture of science. My point is not that science is an irrational, Feyerabendian political struggle, but the more modest claim that these factors play a significant role in scientific consensus. (Obviously the concern over excessive consensus is less important in the case of the social sciences and philosophy.) In addition, a huge number of the scientific and philosophical theories for which disagreement is reasonable are of central importance to our understanding of the world and of our place in it. Finally, while the disagreement in question may not be insurmountable over the very long run, it is nevertheless highly persistent; few, if any, of these disagreements will be resolved in the course of any one person's lifetime.

So I will take the equal likelihood condition as a necessary condition for being an epistemic peer. However, Elga argues that once we accept this condition we cannot avoid the conclusion that in case of disagreement we should give our peer's judgment and our own judgment equal weight. For suppose that it were reasonable for us to weight our own judgment somewhat higher than x 's judgment (assuming the same evidence). In that case, we would by definition no longer regard x as our epistemic peer. Thus, the mere fact that we disagree would have to provide some evidence that x was not after all an epistemic peer and so we could, as it were, bootstrap into own epistemic superiority simply on the basis of our disagreement. But according to Elga this is wrong. If we antecedently believe that x is as likely to be correct as we are, then the mere fact that we disagree does not provide any reason to think that we are more likely to be correct! If we do believe this, we believe it without evidence.

In my opinion, Elga's bootstrapping argument doesn't succeed generally. The problem is that Elga does not adequately take into account the holistic relation between evidence and theory (Quine 1951). It is a presupposition of the bootstrapping argument that it is illegitimate to reduce our credence in the proposition that an individual with whom we disagree is an epistemic peer on the basis of the fact that we disagree. But this, in effect, is to make a mistake which is directly opposite to the one made by the verificationists; it is to presuppose that a certain (nonlogical) proposition, that x and I are epistemic peers with respect to p , is immune to revision in light of further information, specifically, the information that x and I disagree about whether or not p . But it might, in fact, be the case that the best overall adjustment in theory is to reassess x 's status as an epistemic peer.

The reasonableness of this sort of theory-to-theory transition is brought out clearly in cases of very strong evidence. Suppose that you and I are looking out the window onto the quad where I clearly and distinctly see a young man playing Frisbee with his dog. I remark on the fact only to discover that you reject my claim that there is any such scene occurring in the quad. Prior to this exchange, I had good reason to regard you as an epistemic peer with respect to such observational matters. But in such a case, my visual evidence is so strong, that it simply makes more sense to jettison my antecedent belief that you are an epistemic peer than to accept that I am mistaken.⁴

Moreover, the bootstrapping argument rests on a false dilemma. The argument relies on the assumption that, given some disagreement between x and I, either I must not really have regarded x as an epistemic peer to begin with or I ought to give x 's assessment and my own assessment equal weight. But this overlooks a number of alternatives. For instance, we might take the disagreement to be evidence that our counterpart's evidence is in some way degraded (e.g., because of temporary inattention, very local conceptual misunderstanding, etc.) or we might simply take our counterpart to be lying about her conclusions. Either of these responses would allow us to place additional weight on our own conclusions without impugning the status of our counterpart as an epistemic peer. Of course, as Christensen (2007)

⁴ Note, that my belief that there is a young man playing Frisbee with his dog in the quad need not have credence 1 in order to take disagreement with me on the issue as a good sign that you are not an epistemic peer on such matters. All that is required is that the overall changes to my theory of the world would have to be more radical than the changes required by my reassessment of your peer status. In cases such as this one, that seems likely since error on my part will be possible only under extremely unusual circumstances.

notes, the mere fact of disagreement might rationally require that I lower my credence in the proposition that x is my epistemic peer. But unless the disagreement in question is persistent (or, more carefully, unless the lowering of my credence value must be iterated an appropriate number of times), there is no reason to think that the lower credence value must drop below the level required for full belief.⁵

So I take it that as a general argument for the equal weight view, the bootstrapping argument fails (see Kelly forthcoming for a similar line of argument). Unfortunately, however, the sorts of moves suggested above do not appear to be available under the conditions which are characteristic of ideal rational inquiry. There are two characteristics of the Millian condition which are relevant. First, as argued in §1, it is a necessary condition on the [CCC] that our critics be peer critics and that we regard them as such. Consequently, it would do us no good to avoid the equal weight view by abandoning our true belief that our critics are our epistemic peers.

The second characteristic of the Millian condition which generates problems is the fact that the peer critics invoked by the [CCC] are in *persistent* disagreement with us. Suppose then that upon first encountering disagreement with our views, we attempt to account for this disagreement by hypothesizing that our critics' evidence is in some way degraded. But given that the disagreement will persist as we proceed further in our collective inquiry we will, after a very short period of time, have to regard these critics as simply being incapable of attaining the same quality of evidence as us. But this amounts to positing a truth-conducive factor on which we are superior to our critics and so, to not regard them as genuine peers after all! Alternatively, suppose we take our critics to be kidding (or lying) about their disagreement with us. In such a case, we would take our critics to be in agreement with us about our conclusions; the disagreement would be only apparent. But then we would not satisfy Mill's seemingly well-motivated demand that the disagreement in question be genuine. Consequently, this response would not allow us to overcome the concerns we had about our propensity to factors like theory-bias.

What these considerations suggest is that, whether or not disagreement among peers is as general a problem as some philosophers believe, it is a much more difficult problem when we move to the idealized conditions of the Millian context. The reason for this is that the most plausible responses to the equal weight view depend, either directly or indirectly, on the presence of an epistemic shortcoming of some sort. As we idealize away from those limitations, it becomes much more difficult to avoid epistemic stalemate with our peers. Moreover, the recalcitrance of the problem of reasonable disagreement in idealized settings is, in itself, a philosophically important result since it illuminates certain general characteristics of our epistemic situation which, though perhaps swapped in the actual world, are nevertheless operative.

Consequently, the equal weight view enjoys a fairly high degree of plausibility within the context of ideal rational group inquiry. But, of course, if we have equally good epistemic grounds for accepting

⁵ The importance of the distinction between persistent and transient cases is elaborated below. It seems to me that the literature on reasonable disagreement suffers from a general failure to recognize that there are a wide variety of distinct kinds of cases of disagreement and that what may count as a reasonable move to make with respect to one of those kinds of cases, may not be reasonable with respect to the others.

both p and not- p , it seems that we should (at least, epistemically speaking) refrain from believing either.⁶ From this it follows that if there exists an individual who we take to be our epistemic peer and with whom we share our evidence but with whom we disagree concerning p (hereafter, epistemic peer critics), then we cannot know that p . For in that case, either we will not believe that p or, if we do, we will do so without adequate epistemic warrant.

But this would mean that the [CCC] entails widespread skepticism under what on independent grounds we would consider epistemically ideal conditions. For [CCC] says that in order for our theoretical beliefs to be fully or ideally justified, we must have a group of epistemic peer critics for those beliefs. And as we have just seen, the existence of such a group of epistemic peer critics undermines our claims to knowledge. So, paradoxically, the conditions which are apparently ideal for rational inquiry undermine the apparent purpose of such inquiry (specifically, individual-level knowledge).

Of course, this line of reasoning depends on our *taking* the critics to be epistemic peers. So, if [CCC] were to merely require the existence of a group of epistemic peers and not require that we take them to be epistemic peers, the paradoxical conclusion could be avoided. However, if we don't take our critics to be epistemic peers, this will generate second-order doubts regarding our ability to rebut the strongest available counterarguments. For our peer critics are intended to produce counterarguments which approximate what it would be reasonable for someone with our own level of cognitive ability to come up with given our evidential situation and general state of knowledge. This second-order doubt is mooted if we recognize our critics to be epistemic peers, since the judgment of our peers is as good as our own. Thus, in the ideal case we would (correctly) regard our critics as our epistemic peers and the puzzle is once again up and running.

III. Disagreement, Underdetermination, and Epistemic Conservatism

The underlying tension which the preceding discussion has exposed is this: our need to rely on others for the critical evaluation of our views conflicts with our need to regard those critics as intellectually

⁶ It might be thought that the move from the equal weight view to general agnosticism is too quick. Specifically, it seems that we should distinguish between two, independent claims:

- (i) The Summation View: The view that x should give the same weight to the opinion of any epistemic peer as x gives to x 's own opinion in the sense that one should average the credences of the believers, and
- (ii) The Stalemate View: The view that if x believes that p and x knows that x has an epistemic peer who believes not- p , then x should become agnostic about p .

It might be thought that these views are quite different, and indeed, incompatible in their recommendations except for the special case in which the overall distribution of opinion among the relevant group or peers is averages to 0.5. Moreover, it might be thought that it is really the Summation View that we ought to accept.

Unfortunately, the Summation View is not a viable position. The reason for this is that the question of who counts as an epistemic peer is not independent of how we assign our credences. Suppose, for instance, that on the basis of evidence E , x comes to believe that p with a credence of 0.7. Suppose then that y concludes that not- p on the same evidence. It seems that x must either say that the probability that x is wrong is 0.7 or x must take her original credence assignment to be flawed (i.e., she revises her claim that E supports p to degree 0.7). But, by the definition of an epistemic peer, if x and y are epistemic peers, then they are equally likely to be correct about p ; and this means that when x 's and y 's judgments are incompatible the probabilities must be set at 0.5. But by Elga's bootstrapping argument, it is not permissible for x to revise her assessment of y as an epistemic peer merely on the basis of disagreement. Her only option, therefore, is to regard the original credence assignment as flawed. And in this case, it will make no epistemic sense whatsoever to progress to the Summation View.

In any event, under the idealized circumstances under consideration here, it is reasonable to suppose that the relative proportion of advocates of various opposing views is roughly equal.

trustworthy and competent. This means that if we are to avoid skepticism, we must find a way to reconcile the competency of our peer critics with what we take to be persistent and systematic error on their part.

One possible way of doing this is to take advantage of the transient underdetermination of theory by evidence in order to “loosen up” the concept of an epistemic peer in a way that makes room for epistemically blameless variation in our conclusions.⁷ Specifically, if one were to accept that one’s current evidentiary position (perhaps including whatever theoretical virtues one wishes to posit) does not rationally determine a commitment to a unique theory, then one would be free to maintain the competence of one’s critics despite their genuine disagreement. Against this backdrop, an epistemic peer would be construed as someone who would be as likely as us to judge correctly given the same evidence *and the same background theoretical beliefs*.⁸ Because evidence rationally underdetermines theory, we would not be driven to the unhappy dilemma of either abandoning our own views or downgrading the epistemic competency of our critics. In this way, underdetermination appears to allow for blameless (and so, reasonable) disagreement.⁹

Feldman (forthcoming) argues that this apparent way out of the paradox is only apparent. According to Feldman, once we acknowledge that there is no rational way of choosing between two or more distinct theories, the only rational course of action is to suspend judgment since we will then recognize that our starting point is epistemically arbitrary. If this is correct, then the appeal to theoretical underdetermination simply brings us to the same conclusion by way of a different path: rather than suspending judgment because we should ultimately give equal weight to the contrary assessment of our peers we should instead suspend judgment because the total available evidence does not determine a unique theory.

I believe that Feldman is wrong in his assessment in part because his view rests on a too narrow conception of the epistemic.¹⁰ I will proceed as follows. First, I will sketch an argument that epistemic conservatism is the correct general theory of when we are justified in relying on our standing beliefs. I will then argue that awareness of the underdetermination thesis does not constitute a defeater for one’s prima

⁷ Sklar (1975) distinguishes between radical and transient underdetermination of theory by evidence. Roughly, the distinction is between underdetermination “in the evidential limit” and underdetermination given present evidence. While radical underdetermination is a highly contentious theoretical claim, transient underdetermination is hardly controversial in many, and perhaps most, cases.

⁸ A certain amount of care is required in stating this condition. In saying that the individual would judge “correctly” we cannot mean that the individual would judge “truly” unless we are also willing to embrace some form of anti-realism. What is intended, rather, is that an epistemic peer is as likely as we are to arrive at the most rationally supported conclusion given those features which fix such a conclusion uniquely. This way of framing the issue has the evident virtue of making correctness an issue of epistemic rationality rather than truth, which is what we would expect of a discussion concerning reasonable disagreement.

⁹ Of course, it might well be the case that there is no *radical* underdetermination of theory and that the best total theory of all the evidence “in the limit” is fixed in stages. If so, then the relevant background information at a given time will fix a unique best partial theory of the world and peer disagreement will no longer be possible with respect to that domain of inquiry. In such a case, it might well be epistemically desirable for disagreement to persist, but it could not constitute disagreement among peers. Nevertheless, even under such circumstances the present view would caution against concluding that our critics are not our epistemic peers since it will rarely be obvious that the evidence does, in fact, uniquely determine a partial theory.

¹⁰ Strictly speaking, Feldman does not consider the exact conjunction of philosophical positions which I will endorse, namely the underdetermination of theory by evidence and epistemic conservatism. In particular, Feldman never considers the possibility of invoking conservatism as a way of over coming his arguments against variations in (in his rubric) “starting points”. The ensuing discussion, therefore, must be understood as my own reconstruction of what Feldman would say given his published work on the issue.

facie reliance on one's standing beliefs. Finally, I will argue that these considerations are, in fact, properly epistemic (rather than pragmatic or prudential).

Methodological conservatism (Sklar 1975) is the thesis that we need not abandon our theories simply because we become aware of an incompatible theory which is equally warranted by the available evidence.¹¹ Arguably, some such conservative principle is an essential component of our epistemic situation. 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 (Harman 1986), we will be unable to make justifiable progress in our theoretical undertakings. In fact, insofar as our theoretical understanding aspires to comprehensiveness, re-verifying all relevant standing beliefs would make anything like a justified comprehensive worldview practically impossible: we could make progress only by failing to discharge our epistemic duties (viz., by failing to re-verify all relevant standing beliefs) and we could discharge our epistemic duties only by stifling further development. Thus, unless we have some plausible account of the present justification of our standing beliefs which is less demanding than the outright possession of a justification for those beliefs, an especially pernicious form of skepticism looms: to the extent that we develop general scientific and philosophical worldviews, they will turn out to be epistemically ill-founded.

Considerations such as these give us reason to posit the following static conservative epistemic principle:

For any individual x and proposition p , x 's believing that p is both necessary and sufficient for x 's being prima facie justified in believing that p .¹²

But it seems reasonably clear that the static version of epistemic conservatism entails a corresponding dynamic version:

For any individual x and proposition p , if x believes that p at t , then x should not abandon this belief at t unless one has an adequate reason to do so.

After all, the following principle of belief revision is hardly controversial:

For any individual x and proposition p , if x believes that p and x has a prima facie justification for believing that p at t , then x should not abandon this belief at t unless one has an adequate reason for doing so.

¹¹ The principle of methodological conservatism articulated in the text might be dubbed “weak” methodological conservatism since it merely says that belief retention is permissible. Strong methodological conservatism would be the following thesis: we *should* not abandon our theories simply because we become aware of an incompatible theory which is equally warranted by the available evidence. Feldman's position requires the denial of weak methodological conservatism.

¹² It seems to me that the only viable alternative to conservatism is a Burge-style preservationist principle of the following sort:

A standing belief b for which one does not have an adequate set of justifying reasons at t is prima facie justified for x at t if and only if x was justified in forming b and this justification has not, in the meantime, been defeated.

While I think that there are good reasons for preferring conservatism to preservationism (see Moffett 2007 for a discussion), for present purposes either principle will suffice, assuming that the appropriate adjustments are made below.

But if, as the static version claims, merely having a belief constitutes a *prima facie* justification for believing it, then dynamic conservatism follows.¹³ Consequently, there is an epistemic presumption in favor of the theories to which we are antecedently committed.

Dynamic conservatism, however, is a weaker thesis than methodological conservatism. For methodological conservatism entails, not only that one should not abandon a standing belief without adequate reason, but also that the realization that there is an equally well-supported alternative does not constitute an adequate defeating reason. This is the claim to which Feldman objects, rather than to any more general principle of epistemic conservatism.¹⁴ That is, even if Feldman grants that belief confers *prima facie* justification, he would nevertheless deny that this *prima facie* justification holds up in light of the knowledge that one's (theoretical) belief is rationally underdetermined by the evidence. In this case, Feldman argues that one epistemically speaking ought to go agnostic. After all, once we accept the underdetermination thesis, we have no more reason for thinking that our belief is true than some other alternative belief.

But consider the cognitive situation to which this epistemic advice would lead us. Assuming (as we should) that transient underdetermination is a widespread phenomenon, the result would be something like an impartial, non-committal survey of a vast portion of logical space: the world could be like this, or like that, or like that, and so on, *ad infinitum*. While such theoretical detachment is perhaps feasible with respect to limited domains—for instance, my own epistemic stance concerning death or religion might be of this sort—it would be epistemically devastating if carried through to the extent required by Feldman's position. Believing some proposition (or set of propositions) typically involves more than simply taking that proposition to be true (Cohen 1992), it also involves conceptualizing the world in accordance with the truth of that proposition, seeing the world as being thus and so. But this phenomenon of conceptualizing the world in a certain way is not epistemically inert; rather, it is fundamental to our ability to produce and develop new ideas and direct the course of further investigation. The point here is not the familiar one that observation is theory-laden, though that point too is relevant. But whether or not there is a theory neutral description of our basic evidence, the capacity to conceptualize the world from within a given perspective is pivotal to our ability to develop a systematic worldview. In short, there are non-trivial theoretical benefits

¹³ The converse entailment, however, is more tenuous. If a capricious, arbitrary or otherwise unmotivated change of belief was epistemically prohibited only if the belief was justified, then the dynamic version would plausibly entail the static version. However, other epistemic considerations might equally well tell against such changes regardless of the justificatory status of the belief in question. For instance, if diachronic belief stability is an independent epistemic desideratum, then this fact might be sufficient in and of itself to ground the dynamic version of epistemic conservatism.

¹⁴ There are, of course, various arguments which purport to show that the more general versions of epistemic conservatism are wrong. Consider, for instance, Foley's (1982) bootstrapping argument. According to Foley, epistemic conservatism has the following counterintuitive consequence: it entails that a "borderline justified" belief which would not otherwise be justified could become justified simply because it is believed. In my opinion, Foley's argument is specious. Its plausibility turns on an equivocation between one's being justified in forming a belief and one's being justified in relying on a standing belief. Epistemic conservatism is plausible only with respect to the second concept. Consequently, when told with the proper care, the correct description of the case is one in which the individual is not justified in forming the belief but is justified in relying on it. So understood, the correct moral of Foley's example is that the mere fact that one believes something cannot provide one with a reason for coming to believe it. But that is hardly surprising and certainly not something to which the conservative is committed.

associated with genuine belief which would be lost if one were to adopt a theoretically neutral psychological standpoint.¹⁵

Now one might doubt that this alleged benefit is really an epistemic benefit (Foley 1983). After all, what is being claimed is that genuinely believing a theory helps us to elaborate and develop that theory even though we acknowledge that the evidence in support of it makes it no more likely to be true than some set of alternatives. In what sense, then, is this an epistemic reason for belief? As Kuhn (1962) saw in the case of scientific theorizing, the elaboration of our theories plays a crucial role in overall theoretical progress. Not only can such development expose anomalies which ultimately lead to new, more adequate theories (revolutions), but partially developed theories are hardly the best for purposes of comparative evaluation. It would be only by construing the epistemic in an inappropriately narrow manner, as reasons which make the belief itself more likely to be true, that one can deny that the preceding are properly epistemic reasons for belief.

There is a related, but perhaps philosophically less tractable, reason to refrain from agnosticism in light of the underdetermination thesis. From the point of view of one's intellectual development, our theoretical commitments shape, not just how we see things, but who we are. And while we can aspire to being more than mere apologists for our views, the mere intellectual awareness that there are feasible alternatives hardly seems like an adequate ground for abandoning those commitments (particularly, if one denies the radical underdetermination of theory). The point here is reminiscent of Bernard Williams' (1981) observations about the integrity of the self. I have traveled along a certain intellectual trajectory. The commitments I have taken on board in that process have, in part, made me who I am. To abandon them *merely* on the basis of some ideal of epistemic neutrality seems wrong in much the same way that abandoning my projects and goals on the basis of a principle of moral neutrality seems wrong. Specifically, theoretical inquiry is a temporally extended process requiring a more or less long-term commitment on the part of an investigator. Abandoning such commitments simply because an attractive competitor, even an otherwise equally attractive one, arises does not do justice to our status as epistemic agents, especially if (as Mill says) it is genuine believers who are best positioned to do their "do their very utmost" for a theory.

Of course, it might also reasonably be suggested here that this does not really constitute an epistemic reason, but rather a prudential or moral one. It is not clear to me that this is correct. After all, it is our status as epistemic agents which generate the commitments. But even if this is not so, it seems to me that such considerations have epistemic consequences. For it seems to me that epistemology ought to be

¹⁵ The reader should bear in mind here that the extent of agnosticism in question is quite extensive. While there are no doubt many benefits to agnosticism concerning limited areas of inquiry, such benefits accrue only against a backdrop of massive overall theoretical commitment. Indeed, it is arguable that the degree of agnosticism which would be required by Feldman's position is psychologically impossible. There is no theory-less cognitive "state-of-nature" to which we can return when our investigations fail to yield a determinate verdict. Rather, having a general (though perhaps implicit) conception of the world to which we are committed is a necessary precondition for having anything like the sort of mental lives which we do in fact have.

Of course, we are assuming throughout that the beliefs in question were formed on the basis of otherwise adequate evidence. The question is whether or not underdetermination epistemically gives us an adequate reason to overturn such beliefs. Obviously, agnosticism is permitted (in fact, required) when there are no adequately supported theories on offer.

responsive to moral and prudential considerations in the following way: what is morally/prudentially required ought to be epistemically permissible and, conversely, what is epistemically prohibited ought not to be morally/prudentially required. If this is right, then moral considerations can inform us as to what is epistemically permissible, even if they do not themselves constitute epistemic reasons.

Let me recap. I have argued that there is an epistemic presumption in favor of our standing beliefs which is not overridden by the existence of an alternative, but equally well-supported set of beliefs. If this is correct, then (weak) methodological conservatism is vindicated. And this in turn means that the transient underdetermination of theory by evidence underwrites a plausible resolution to the puzzle of ideal rational inquiry developed in §§1–2.

IV. Conclusion

In this paper, I have attempted to show that the problem of reasonable disagreement is tied directly to central issues in epistemology and is not simply a troubling but inessential feature of our epistemic situation; the problem persists even under conditions of ideal rational inquiry. Moreover, by locating the discussion in this idealized setting, I hope to have isolated what is central to the problem from what appear to be mere stopgap measures. If I am correct, the central problem is best resolved by appeal to the transient underdetermination of theory by evidence together with a principle of epistemic conservatism (or some related principle of epistemic inertia).

Acknowledgements

This paper began as a commentary on Richard Feldman's "Reasonable Disagreement" at the 2004 *Inland Northwest Philosophy Conference*. The author would like to thank audiences at Rutgers University and the University of Wyoming. Special thanks to John Bengson, Adam Elga, Richard Feldman, Thomas Kelly and Jonathan Kvanvig for comments.

References

- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review*, 116: 187–217.
- Cohen, L. Jonathan. 1992. *Belief and Acceptance*. Oxford: Oxford University Press.
- Elga, Adam. Forthcoming. "Reflection and Disagreement." *Noûs*.
- Feldman, Richard. 2006. "Epistemological Puzzles about Disagreement." In S. Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press.
- Foley, Richard. 1983. "Epistemic conservatism." *Philosophical Studies*, 43: 165–82.
- Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, Alvin. 1987. "Foundations of Social Epistemics." *Synthese*, 73: 109–144.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, MA: MIT Press.
- Kelly, Thomas. Forthcoming. "Peer Disagreement and Higher-order Evidence." In R. Feldman & T. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press.
- Kuhn, Steven. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: Chicago University Press.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy*, 74: 549–67.

Mill, John Stuart. 1956 [1859]. *On Liberty*. Indianapolis, IN: Bobbs-Merrill.

Moffett, Marc A. 2007. "Conservatism, Stability and the Metaphysics of Belief." Unpublished MS.

Quine, W. V. O. 1951. "Two Dogmas of Empiricism." *Philosophical Review*, 60: 20–43.

Sklar, Lawrence. 1975. "Methodological conservatism." *Philosophical Review*, 84: 374–400.

Williams, Bernard. 1981. *Moral Luck*. Cambridge: Cambridge University Press.