

The Procedural Epistemic Value of Deliberation

Forthcoming in *Synthese* special issue on “Epistemology of Inclusiveness”

Fabienne Peter¹

f.peter@warwick.ac.uk

February 2012

Abstract

Collective deliberation is fuelled by disagreements and its epistemic value depends, inter alia, on how the participants respond to each other in disagreements. I use this accountability thesis to argue that deliberation may be valued not just instrumentally but also for its procedural features. The instrumental epistemic value of deliberation depends on whether it leads to more or less accurate beliefs among the participants. The procedural epistemic value of deliberation hinges on the relationships of mutual accountability that characterize appropriately conducted deliberation. I will argue that it only comes into view from the second-person standpoint. I shall explain what the second-person standpoint in the epistemic context entails and how it compares to Stephen Darwall’s interpretation of the second-person standpoint in ethics.

I. Introduction

An important question one can ask about collective deliberation is whether it increases or decreases the accuracy of the beliefs of the participants. But this instrumental approach, which only looks at the outcome of deliberation, does not exhaustively account for the epistemic value that deliberation

¹ I’m very grateful to Peter Poellner and an anonymous referee for their comments on earlier versions of this paper. I also benefitted greatly from discussions I’ve had with Herlinde Pauer-Studer and Hans Bernhard Schmid about the main ideas of this paper.

might have. My aim in this paper is to show that collective deliberation may also have procedural epistemic value. The idea that collective deliberation has procedural value is familiar from the practical context. The procedural value of practical deliberation is typically captured in terms of the relationships of reciprocity (Rawls 1993), mutual accountability (Darwall 2006), or equal respect (Larmore 2008) that characterize – appropriately conducted – collective deliberation. I shall draw an analogy between the practical case and the epistemic case and argue that, under certain circumstances, deliberation has procedural epistemic value.

My argument links the procedural epistemic value of deliberation to the – second-personal – notion of relationships of mutual accountability among epistemic agents. The second-person standpoint has received powerful advocacy by Stephen Darwall in his recent book with that title. The main claim his book defends is that moral reasons, as a subcategory of practical reasons, are irreducibly second-personal. They stem from relationships of mutual accountability between persons who treat each other as self-originating sources of valid claims. Darwall also maintains that while practical reasons can be irreducibly second-personal, epistemic reasons cannot. What one has reason to believe depends on truth. Since truth is third-personal, Darwall argues, epistemic reasons are fundamentally third-personal.

I think this dichotomy overstates the case. I will show that the second-person standpoint is relevant not just in the practical context but in the epistemic context as well. In certain circumstances, deliberation gives rise to second-personal epistemic reasons. The relevant circumstances arise when deliberation is characterized by disagreements among epistemic peers.² In a disagreement among peers, each party needs to consider whether they can justifiably stick to their guns or whether it is appropriate that they adjust their beliefs in the direction of the other. It turns out not only that there are situations in which it is appropriate that they both adjust their beliefs in direction of the other,

² I shall comment later, in section II, on what I mean by epistemic peers. I shall work with a very weak notion of epistemic peers, one that can be applied to large-scale collective deliberation as it occurs in a democratic context.

but also that the reason that they have to do so is to some extent independent of the subject matter of their deliberation. What this shows is that accuracy of belief is not the only source of epistemic reasons. Some form of mutual accountability between the deliberative parties becomes an additional factor. I call this thesis about the epistemic value of deliberation the accountability thesis. Mutual accountability between epistemic peers is a procedural value and an appropriate account of mutual accountability invokes the second-person standpoint. The second-person standpoint thus helps us understand how the epistemic value of deliberation may not merely be instrumental but procedural as well.

II. Deliberation and Disagreements

Collective deliberation is fuelled by disagreements and is, at least in part, a tool to evaluate such disagreements. This is evident in a wide range of deliberative contexts – e.g. in science, in politics, or in any group endeavour such as committee work in public or private associations and organizations. What characterizes a disagreement is that the parties hold mutually incompatible beliefs. Such disagreements can occur for many reasons. Sometimes, not all parties have considered all the available facts. Sometimes, one party to the disagreement has misinterpreted the evidence or drawn the wrong conclusions from it. In this case, well-functioning deliberation may serve as an instrument to demonstrate the mistake of the wrong party and lead to an agreement on the correct belief. Once the disagreement has been eliminated in this way, deliberation ends. But the opposite can happen as well, of course. Dysfunctional deliberation may end in a consensus on the wrong conclusion. The literature on group polarization, for example, documents such dysfunctional deliberative dynamics (e.g. Sunstein 2002).

I want to focus on a special case of deliberation here, namely the case of deliberation among parties who consider each other epistemic peers. An epistemic peer is someone who you take to be equally likely to make a mistake (Elga 2007). This is a very weak definition of what it means to be a peer, since it only takes the form of an all-things-considered criterion and doesn't invoke any input

conditions such as equal formal qualifications or equal computational abilities, etc. On this definition of peers, deliberation among parties who consider each other peers can occur not just in contexts of academic or expert inquiry, but in any small or large social collective, for example on issues which are too wide-ranging and complex for anyone to count as an expert, or when relevant information is dispersed across all deliberative parties.

While you judge your peers, by definition, to be equally likely to make a mistake in a particular situation, this doesn't mean that two peers always perform equally well. Sometimes you have information about the circumstances of the disagreement that makes it appropriate for you to discount their judgment to a certain extent. For example, of two scientific colleagues who compare their different conclusions about the validity of a hypothesis, one may have double-checked the data and the calculations and asked an assistant to do the same while the other was pressed for time and admits that he only ran what he was given through an off-the-peg computer programme. In a context of social deliberation, imagine a case of committee work, even though all participants consider each other to be equally able to take up the available evidence, some may have carefully thought about the implications of the evidence presented and this is evident from how they argue in support of their views, others respond with a gut reaction. But what is important is that, absent independent information about their respective ability or inclination to perform, peers need to consider the implications of their disagreement. Independence here refers to information other than the ex post information of a disagreement with a peer. The thought is that the fact that someone who you ex ante regarded as a peer disagrees with you is not sufficient ground to dismiss his or her belief.³

³ On independence, see Christensen (2011: 1): "In evaluating the epistemic credentials of another's expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn't rely on the reasoning behind my initial belief about P."

What is the rational response for someone who realizes that a person she considers an epistemic peer disagrees with her? Is it possible for epistemic peers to have a reasonable disagreement or does rationality force them to converge towards a particular doxastic attitude? Is any revision of one's initial beliefs required?

One way to answer the last question negatively is by arguing that what justifies someone's belief is their response to the available evidence and by denying that information about the beliefs of peers carries independent epistemic weight.⁴ And if the available evidence is such that two persons, including two peers, can justifiably form incompatible beliefs, no adjustment is required. On this view – sometimes called the Steadfast View of disagreements – the issue of whether the beliefs of the two parties are justified can be dissociated from the fact that they disagree.⁵ The consequence is that reasonable disagreements are possible and the parties to the disagreement are not rationally required to give any weight to each other's beliefs.⁶ If two peers hold conflicting but permissible views – say p with credence 0.2 and $\neg p$ with credence 0.8 – and if each has good epistemic reasons to regard their original belief as justified, they may simply end up agreeing to disagree.

If this were the correct view of disagreements among peers, the case for the procedural epistemic value of collective deliberation is closed from the start. If epistemic value only depended on how epistemic agents, unilaterally, respond to the available evidence, it would imply that deliberation in itself has no epistemic import. If deliberation itself is to have a hold on the deliberating parties, it must be because an alternative account of disagreements among peers is possible and indeed more plausible.

⁴ I can't get into a discussion here of what conditions are necessary and sufficient for the justification of belief. I'm simply bracketing this question and argue from the premise that there is some set of such conditions that can be met.

⁵ The label is from Christensen (2009 and 2011); but see Kelly (2005) for the view itself.

⁶ Kelly (2010: 116) has a very clear articulation of this view. On what he calls the "symmetrical no independent weight view", there are at least some cases of peer disagreement in which "both parties to the dispute might be perfectly reasonable even if neither gives any weight at all to the opinion of the other."

As a first step towards an account that explains the epistemic significance of deliberation itself, consider the case where the available evidence warrants a unique belief. Suppose, after carefully considering this evidence, two epistemic peers hold mutually incompatible beliefs. If the Uniqueness Thesis is true, i.e. if the body of evidence warrants a unique belief, and if the parties to a disagreement have full access to this body of evidence, then they cannot both be justified to hold the beliefs they do. The evidence will warrant p with a certain credence and so at least one of them holds an incorrect belief.

The Uniqueness Thesis implies that there is epistemic content in the comparison of the beliefs of epistemic peers. If I realize that someone I consider a peer holds a different belief from the one that I assumed was warranted by the available evidence and absent independent information about the performance of my peer, Uniqueness forces me to ask the question of who is right. Since we cannot both be right, the comparison warrants a response. This brings me into a deliberative context with my peer and vice versa.

If we grant Uniqueness, the rational response to a disagreement between two persons who consider each other epistemic peers and have fully disclosed the evidence available to them is to reduce confidence in their original beliefs. How much their confidence should be reduced depends on additional considerations, especially considerations about the relative weights that it is rational to attribute (i) to one's own ability to come to the correct conclusion and (ii) to the fact that someone one has reason to regard as a peer has come to a different conclusion.

Take the case in which two peers find that they hold diametrically opposed views – p and not p . Assume also that they have independent information that lets them consider each other as peers and that suggests that each is just as likely to perform well. In this case, suspending belief is the rational response for each. (ii) becomes decisive in this case, as there is no reason to give more weight to one's own original belief than to the belief of the peer. What has come to be called the Equal Weight View (Elga 2007) focuses on this case to claim that reasonable disagreements are not

possible. On this view, absent any further information, the rational response to a disagreement between peers is for each to meet the other halfway and thus to dissolve the original disagreement. As long as Uniqueness holds and each has access to the same body of evidence, two peers who are were likely to perform well cannot both be justified to hold the original beliefs they do and so a reasonable disagreement between them is not possible. On this view, rational deliberation between peers forces them to reach an agreement.

The Equal Weight View is a special case of a broader position called the Conciliatory View (Elga 2010; Christensen 2011). On the Conciliatory View of peer disagreements, you are required to give some weight – but not necessarily equal weight – to the belief of your peer. It makes room for the possibility that, in spite of the information that the other party in the disagreement is a peer, you also have reason to give some extra weight to your own belief—not because it’s your own, but because of evidence you have about your own performance. There may even be clear cases in which you are entitled to discount the belief of your peer completely, say because only you, but not your peer, cross-checked several sources or because your peer came to a result that is obviously impossible (Christensen 2007; Elga 2007). But when independent information about the ability and willingness of your peer to perform well gives you good reasons to believe that there is a positive probability that your peer is correct, you will be required to move some distance towards the belief of your peer.

But even without upholding Uniqueness, it is possible to argue for the need to revise one’s belief in response to a peer disagreement. On Thomas Kelly’s revised version of the Steadfast View – he calls it the Total Evidence View – a disagreement with a peer is treated as a piece of – higher-order – evidence. On this view, “what it is reasonable to believe depends on both the original, first-order evidence as well as on the higher-order evidence that is afforded by the fact that one’s peers believe as they do” (Kelly 2010: 142). Just like you have reason to adjust your original belief after you double-check a mental math exercise on your calculator and come up with a different result, so you

have reason to adjust your belief if you find yourself in a disagreement with a peer. This holds independently of whether the evidence you get from the calculator or the disagreement with the peer is conducive to accuracy or misleading.

The Total Evidence View and the Conciliatory View come apart with regard to how they explain the epistemic role that disagreement with a peer plays. According to the Total Evidence View, if the disagreement receives any weight, it is as a piece of evidence that a rational epistemic agent needs to consider together with the other available evidence. On the Conciliatory View, by contrast, the disagreement gives you a reason to revise your belief that is of a different kind than the reasons that you had to form your original belief.⁷

I can't fully assess the debate here, and I don't need to as both sides accept that in some cases of peer disagreement, you have reason to adjust your belief in direction of your peer (and vice versa). But I have some sympathies for the line that the Conciliatorists take. Here is why. What makes it plausible to claim that a disagreement with a peer gives you epistemic reasons of a different kind than those that led you to your original belief is our fallibility as epistemic agents. However hard we try to appropriately respond to the evidence we have, we often fail. We share this predicament with our peers and so we all have reason not just to try and respond to the evidence in the best way we can but also, under some circumstances at least, to give weight to the belief of a peer. David Christensen (2011: 33) expresses the thought well: "Rationality requires that I take seriously evidence of my own possible cognitive malfunction in arriving at my beliefs. But insofar as I'm willing to do this, I must evaluate evidence for that possibility in a way that is (at least somewhat) independent of some of my reasoning."⁸

⁷ In other words, what is at stake is the independence condition that Conciliatorists uphold and that I've cited in footnote 2 above; see Christensen (2011) for this diagnosis of the debate.

⁸ Our fallibility is compounded by the fact that disagreements might also arise from the fact that the evidence we have for our beliefs is not fully transparent to us. So far, I've only focused on cases in which both parties to the disagreement have been able to fully disclose the evidence that they have for the beliefs they hold. But

The upshot of this discussion is that in some circumstances there are good epistemic reasons for epistemic peers to be responsible towards each other's claims and to consider some revision of their original beliefs on the basis of these claims. Call this the accountability thesis about the epistemic value of deliberation. Both the Total Evidence View and the Conciliatory View are compatible with the accountability thesis. But the latter is also compatible with a stronger version of it. On the Total Evidence View, you are accountable to the (higher-order) evidence that you get from the fact that someone you take to be a peer disagrees with you together with the pieces of (first-order) evidence that you have. On the Conciliatory View, you have reason to respond to the disagreement with a peer independently of the first-order evidence you had for your original view. Your accountability thus involves the standing that you attribute to your peer in a more immediate sense.

The epistemic reasons that are generated in these circumstances are second-personal: it is not the first-order evidence about the object considered but the claim of your peer that gives you a reason to adjust your original belief. To flesh out the accountability thesis, I now want to turn to a discussion of the second-person standpoint and its interpretation in the epistemic context.

III. The Second-Person Standpoint

As I understand Darwall's idea, the second-person standpoint in the context of moral reasons is a web of four concepts – reason, claim, practical authority, and accountability – in which each entails the others. Someone's valid claim gives someone else a reason to act accordingly. My claim, for example, that you move your foot from on top of mine, gives you a reason to do so. And it does so because we mutually acknowledge our authority in cases like that, i.e. I have a right to make claims of this sort on you (and vice versa).

problems with fully disclosing the relevant evidence suggest yet another take on disagreements among peers, one that is particularly relevant for deliberation about complex social problems. I've argued for this view, which I call the Opacity View, in Peter (unpublished), drawing on contributions by Alvin Goldman (2010) and Ernest Sosa (2010). But since the Opacity View, as I understand it, only adds to but doesn't change the claim I'm focusing on here – i.e. that we sometimes have reason to adjust our beliefs in direction of our peers – I will not discuss it here.

This web characterizes the relationship between moral agents and exhaustively explains, according to Darwall, the normative grip of each of the component concepts. What is distinctive about the second-person standpoint is that it locates the source of normativity in the relationship between moral agents – not in individuals as such (first-person standpoint) and not outside of their relationship (third-person standpoint).

The divisions that Darwall draws between the second- and third-person standpoints, on the one hand, and, on the other, between the second- and first-person standpoints are not equally sharp. The significance of the second-personal standpoint in morality is linked to a denial of a third-personal source of moral normativity. So the distinction between the second- and the third-personal standpoint in ethics is, on Darwall's account, a sharp one. First-personal considerations, by contrast, are part of the second-personal standpoint (Pauer-Studer 2010). It is, after all, me who makes certain claims. But this doesn't imply that the second-person standpoint reduces to the first-person standpoint. Instead, the second-person standpoint serves to qualify first-personal considerations. It identifies which such considerations have moral weight or, in other words, what we can validly claim from each other.

In parallel with the distinction between the second- and the third-personal standpoint, Darwall also sets up a dichotomy between practical reasons on the one hand and epistemic reasons on the other. He argues that while some of our practical reasons, specifically, our moral reasons, are irreducibly second-personal, in the sense that they do not make reference to a third-personal source of normativity, epistemic reasons are not. In the epistemic case, it is truth that is reason-giving. But truth is third-personal, Darwall argues, and so epistemic reasons are fundamentally third-personal.

Underlying this dichotomy between moral and epistemic reasons are two possible sources of authority. Darwall explains the thought with a reference to Hobbes' distinction between "command"

and “counsel”.⁹ With a command, the source of authority is in someone else’s will. With counsel, the source of authority lies in an object outside of the relationship between the person in authority and the person who is accountable. Epistemic reasons, in Darwall’s view, are always a form of counsel; they never simply arise from a claim someone else makes and so never take the form of command. In the moral case, by contrast, when people recognize and respect each other as self-originating sources of valid claims (Darwall 2006: 21), practical authority will take the form of command. Authority, in this case, doesn’t derive from something outside of the relationship between the two agents. It resides entirely in the relationship between people. As Darwall puts it, “when you demand that someone move his foot from on top of yours, you presuppose an irreducibly second-personal standing to address this second-personal reason. ... [This] standing itself neither is, nor simply follows from, any form of third-personal or epistemic authority” (Darwall2006: 13).

Darwall grants that second-personal considerations may have a place in epistemic contexts as well, for example in cases of testimony. But because the authority of the testifier is third-personal – it depends on whether or not the testifier has appropriately responded to the evidence he or she has – my second-personal reason to believe what a testifier tells me is not entirely contained in our relationship. On his view, epistemic reasons can only be derivatively second-personal. The authority to address moral claims to others, by contrast, is, in Darwall’s view, “fundamentally second-personal” (*ibid.*).

Is this dichotomy between second-personal practical reasons and third-personal epistemic reasons plausible? Benjamin McMyler (2011: 146ff) has recently argued that epistemic reasons, too, can be irreducibly second-personal. Here is how he adapts the distinction between second and third-personal reasons to the epistemic case: “a second-personal reason for belief is a consideration that justifies a belief in virtue of interpersonal relations of authority and responsibility existing between

⁹ As Hobbes puts it, “counsel is a precept, in which the reason of my obeying is taken from the thing itself which is advised; but command is a precept, in which the cause of my obedience depends on the will of the commander” (Hobbes *De Cive*, quoted in Darwall 2006: 12n25).

an addresser and an addressee, and a third-personal reason for belief is a consideration that justifies a belief but not in virtue of such interpersonal relations.” He follows Darwall in focusing on the case of testimony, but then argues that the authority of the testifier takes the form of command, not counsel. In his terminology, an irreducible second-personal reason to believe that *p* is one that derives from the testifier testifying that *p*, just like a second-personal moral reason to *phi* derives from the claim someone has on us to *phi*. On his terminology, irreducibly second personal reason are those that necessarily involve a second epistemic agent.

I share McMyler’s view that the dichotomy between second-personal practical reasons and third-personal epistemic reasons that Darwall sets up is too stark. But McMyler’s proposal for how to dissolve the dichotomy is not satisfactory. This is so for two reasons. The first is that his disagreement with Darwall about whether or not there are irreducibly second-personal epistemic reasons is primarily semantic. Darwall defines irreducibly second-personal reasons as those which have their source entirely within relationships of mutual accountability. They arise from second-personal authority. His claim that epistemic reasons cannot be irreducibly second-personal is based on the further claims that truth is third-personal and that the second-person standpoint is thus not sufficient to account for epistemic reasons. McMyler defines irreducibly second-personal reasons more weakly as reasons which necessarily (but not sufficiently) involve a second-personal element. So he grants to Darwall that it is the third-personal epistemic authority of the testifier that gives the addressee the reason to believe. But in his terminology, this doesn’t imply that epistemic reasons derived from testimony are not irreducibly second-personal.

To dissolve the dichotomy between the epistemic case and the practical case McMyler then introduces the further claim that both epistemic and practical reasons ultimately bottom out in third-personal reasons (McMyler 2011:151). This substantive part of the argument against the dichotomy between moral and epistemic reasons rests on a denial of second-personal practical

authority. On Myler's view, practical authority is like epistemic authority, ultimately, third-personal. If this were correct, then Darwall's dichotomy would indeed collapse.

But I don't find this part of the argument satisfactory either. There may be cases in which practical authority is, indeed, correctly thought of as third-personal. But that doesn't imply that practical authority always is. A weaker claim than the one that Darwall makes still goes through. If there are at least some cases in which it is appropriate that persons treat each other as self-originating sources of valid moral claims, then there is second-personal authority and the way in which Darwall draws the distinction between reducibly and irreducibly second-personal reasons still makes sense.

Be this as it may. In what follows, I will simply accept Darwall's claim that there are irreducibly second-personal reasons, in the sense that he defines them. I will also agree with Darwall that epistemic reasons that derive – even indirectly – from third-personal authority, are not irreducibly second-personal. But I will show in the next section that there is such a thing as second-personal epistemic authority and hence that there are epistemic reasons that are irreducibly second-personal in the sense that they stem from second-personal authority relations. Testimony is not, however, the paradigm case for such authority relations. Deliberation among peers is.

IV. Mutal Accountability and Second-Personal Epistemic Authority

I'm granting that epistemic agents aim at truth and that truth is third-personal. But it doesn't follow from the third-personal character of truth that epistemic authority is necessarily third-personal.

Authority, but not truth, is fundamentally agential. Epistemic authority is a concept that captures the right to make claims about what ought to be believed, just like practical authority captures the right to make claims about what ought to be done. To assume that epistemic authority is third-personal is thus to assume that an epistemic agent's right to make claims about what ought to be believed is determined by third-personal truth. But that's the assumption I want to question. Of course, it is often the case that the right to make claims about what ought to be believed derives directly from third-personal truth. When it is the case, then authority is appropriately characterized from the

third-personal standpoint. As the case of testimony discussed above shows, third-personal epistemic authority may even give rise to second-personal epistemic reasons. But what I want to argue here is that there are cases in which epistemic authority itself is – necessarily, even if not sufficiently – second-personal.

Above we saw that in some cases of disagreements among epistemic peers that can arise in deliberation, the appropriate response is, for each epistemic agent, to give (some) weight to the claims of the other. Just to recapitulate, deliberation among peers may lead to one of the following outcomes. First, well-functioning deliberation may demonstrate that at least one of them failed to appropriately take on board relevant evidence. In this case, deliberation dissolves the disagreement. The party who has made the mistake is rationally required to adjust his or her belief. But if the disagreement persists, they each will have to adjust their belief in direction of the other. Under certain circumstances, they will even have to suspend belief altogether and come to an agreement that neither of their original beliefs was justified.¹⁰

Epistemic authority under the circumstances of a disagreement among deliberative parties who consider each other to be peers is not merely the fallout of them appropriately responding to the available first-order evidence. To respect others as epistemic peers, even when they disagree with you, is to acknowledge them as a source of valid epistemic claims beyond what you came to believe to be true.

The case of deliberation among epistemic peers is located in between the practical case that Darwall focuses on and the testimony case. In testimony, the testifier is accountable to truth and his or her epistemic authority is third-personal. Qua this authority, the testifier has a claim on the person to whom the testimony is addressed and the addressee then has a second-personal reason to adjust his belief. In a disagreement between peers, there is a link to third-personal truth, but the “chain of

¹⁰ See Kelly (2010: 150) on how the Total Evidence View concurs with the Conciliatory View on this.

command” doesn’t flow unidirectionally from truth to the first agent involved in deliberation and then to the other participant. Instead, we have a case in which the epistemic agents are mutually accountable to each other. Each has a claim to epistemic authority and each has reason to give some weight to the claims of the other.

One might object that what is going on in theoretical deliberation between epistemic peers will not be a case of command, but a case of counsel, where what the parties to the deliberation have reason to believe depends on what the deliberation is about and not on who makes which claims. But my discussion of the case of deliberation among peers has shown that we sometimes have epistemic reason to adjust our beliefs in response to the claims of others. And what gives me this reason is not the object on which deliberation focuses, but the belief of my peer. In this sense, deliberation involves authority relationships which take the form of command.¹¹

What the case of a deliberation about a potentially reasonable disagreement between epistemic peers shares with the practical case is that authority –practical in one case and epistemic in the other – is vested in the agents and each does, and should, acknowledge the potential authority of the other. Where the two cases come a part is that in the practical case, as Darwall construes it, authority is irreducibly second-personal but in the theoretical case it is not. In the practical case, each is accountable to the other as a self-originating source of valid claims. In the theoretical case, although each has to acknowledge the other as a source of valid claims, the validity of their claims doesn’t stem entirely from the relationship between the agents. Instead, there is a triangulation between third-personal truth and the claims the agents make on each other. They are each accountable both to each other and to the truth they both aim for. What distinguishes the epistemic

¹¹ This holds for both the Total Evidence View and the Conciliatory View. The Total Evidence View distinguishes between the first-order evidence you have for the subject matter of deliberation and the higher-order evidence you get from the fact of a disagreement with a peer. But the Conciliatory View, because of its insistence on the independence condition, has an even better handle on the distinction between command and counsel.

case from Darwall's characterization of the practical case is that although the second-personal standpoint may be necessary to account for some forms epistemic authority, it is not sufficient.

V. The Procedural Epistemic Value of Deliberation

The case of a disagreement between epistemic peers highlights our fallibility as epistemic agents. It is not unusual that we make mistakes. And although we are often able to discover our own mistakes – say by becoming aware of inconsistencies in our beliefs or by experiencing unexpected consequences when we act based on false beliefs, deliberation with epistemic peers is, as illustrated by scientific inquiry, an effective means to identify our misconceptions and to highlight our uncertainties.

Deliberation has instrumental epistemic value if and only if participating in deliberation – comparing evidence and opinions and responding to the evidence and opinions of others – leads to more accurate beliefs. It has instrumental epistemic disvalue if it hinders the formation of accurate beliefs. What characterizes the purely instrumental approach is that it reduces the value of deliberation to its contribution to the – intrinsic – epistemic value of accuracy.

I don't mean to deny the importance of the instrumentalist approach. And I'm also granting that accuracy as defined in relation to third-personal truth is the only intrinsic epistemic value and that a complete account of the epistemic value of deliberation will necessarily make reference to third-personal truth. But what the circumstances I've described above show, I contend, is that the purely instrumental approach doesn't always fully account for the epistemic value that deliberation might have. In disagreement among peers, the epistemic value of deliberation no longer reduces to its value as a means to produce accurate beliefs. As we saw, epistemic agents are sometimes not just accountable to the truth they seek, but to each other as well. This implies that deliberation as such, not merely its outcome, may have epistemic value.

I understand procedural value in opposition to instrumental value. Deliberation has procedural value – epistemic or practical – if its value does not reduce to the value of its result. What the above discussion of peer disagreement showed is that in some cases of disagreements among peers, each has reason to move, even if it is only a tiny bit, in the direction of the other by reducing the confidence in their original beliefs. This is so even if it one side is – from a God’s eye point of view – correct and the other is not. And it is not so simply because either side presents evidence that the other should incorporate. The reason to move in each other’s direction stems from accountability to epistemic peers.

What the accountability thesis highlights is that sometimes the procedural value of accountability to each other is all that is left to epistemic peers who find that they cannot agree on which belief the evidence warrants. In such cases, to be a responsible epistemic agent is to be accountable to one’s peers in a deliberative process. Epistemic authority, the right to make claims about what ought to be believed, is shared by the peers and dependent on a relationship of mutual accountability. My account of the significance of the second-person standpoint in the epistemic context thus lends itself to the explication of the procedural epistemic value of deliberation.

Here’s how the procedural epistemic value of deliberation should be understood. Building on Rawls (1971: 85), we can distinguish among three main forms that proceduralism about deliberation might take. The first is pure proceduralism. According to pure proceduralism, the deliberative procedure is necessary and sufficient for the value of its outcome. The outcome of deliberation has value simply because it is the result of this procedure. It is a distinctive feature of this form of proceduralism that there is no procedure-independent standard that could confer value on the outcomes.

In the practical case it is possible to value deliberation in purely proceduralist fashion. For example, I have argued that democratic legitimacy should be understood in this way (Peter 2008). The thought

is that a democratic decision is legitimate if and only if it has been made through an appropriate deliberative democratic decision-making process. So it's only procedural features that make the outcome legitimate, not features of the outcome of this decision-making process.

In the epistemic case, the value of deliberation is never purely procedural. Accuracy of belief remains the only intrinsic epistemic value and plays a role in the determination of what counts as valuable outcomes. So the epistemic value of deliberation cannot be purely procedural. The other two forms that proceduralism might take are perfect and imperfect proceduralism. What they have in common is that they assume that there is a procedure-independent standard that confers value on the outcomes. According to perfect proceduralism, a certain procedure is necessary to realize valuable outcomes as identified by the procedure-independent standard. According to imperfect proceduralism, the procedure is necessary to approximate valuable outcomes. Although the procedure may fail to reach good outcomes, it gets its value from the irreducible role it has in approximating good outcomes.

The procedural epistemic value of deliberation takes the form of imperfect proceduralism. Accuracy remains the only intrinsic epistemic value and sets a procedure-independent standard for evaluating deliberation. But when epistemic peers' find themselves in a persistent disagreement about what counts as accurate belief, the value of their accountability to each other, and hence of procedural features of deliberation, come into view.

How could the procedural value of deliberation be cashed out? In other words, what are the conditions that specify appropriate relationships of mutual epistemic accountability? I don't have a full answer to this question, but the account I have given above suggests that at least the following conditions are important: (i) respect of epistemic equality, (ii) willingness to enter deliberation and

to explicate one's beliefs; and (iii) uptake.¹² The first condition states that deliberative parties who count each other as peers ought to recognize each other as such. It is then not permissible to give extra weight to one's own beliefs simply because they are one's own. This condition ensures that the participants are each aware of their own fallibility and acknowledge the possibility that their own beliefs may be wrong while their peers might be correct. The other two conditions characterize a deliberative process among peers. The second condition gets deliberation going by demanding that the participants spell out the evidence they have – or think they have – for holding the beliefs they do. Since we assumed from the start that the participants aim at truth and not at some other, non-epistemic goal, this condition orients deliberation to the evaluation of the accuracy of the beliefs they each hold.¹³ Uptake, finally, keeps deliberation going as it ensures that the participants adequately respond to each other. It follows from the view that I have sketched that uptake relates both to the need to respond to the evidence they each present and to the claims they each make as such.

These conditions, however rough, are clearly all procedural conditions. They specify a relationship of mutual accountability between the deliberative parties and they do not reduce the value of this relationship to the outcome it produces.

I accept that in many contexts, we can assess the epistemic value of deliberation in instrumentalist fashion. These contexts are virulent when truth is the aim of deliberation and is available as a standard for assessing the participants' beliefs. The procedural value of deliberation moves into the foreground the less access we have to such a standard and the more epistemic value depends on

¹² Although the conditions are slightly different, my thinking on these matters was much influenced by Helen Longino's proceduralist social epistemology – see her book *The Fate of Knowledge*.

¹³ By assuming that participants aim at truth and by ignoring the non-epistemic goals the participants might pursue in deliberation I am, of course making a strong and very unrealistic assumption. But it is appropriate for the argument I'm building here, as this argument aims to show how deliberation may have procedural value not just in the familiar practical sense of acknowledging equal respect but even in the narrowest possible epistemic sense. The fact that actual deliberation usually involves both practical and epistemic aims is not a problem for my argument.

relationships of mutual accountability between epistemic peers. A full theory of the epistemic value of deliberation would have to spell out the extent to which instrumental and procedural concerns determine the epistemic value of deliberation in different contexts. This is beyond the scope of this paper. All I wanted to show here is that deliberation might have procedural epistemic value.¹⁴

References

- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116: 187-217.
- Christensen, David. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5):756-767.
- Christensen, David. 2011. "Disagreement, Question-Begging, and Epistemic Self-Criticism." *Philosopher's Imprint* 11(6); <http://hdl.handle.net/2027/spo.3521354.0011.006>.
- Darwall, Stephen. 2006. *The Second-Person Standpoint*. Cambridge: Harvard University Press.
- Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41(3): 478-502.
- Elga, Adam. 2010. "How to Disagree about how to Disagree." In Feldman, Richard and Ted A. Warfield (eds.) *Disagreement*. Oxford: Oxford University Press, pp. 175 – 186.
- Goldman, Alvin. 2010. "Epistemic Relativism and Reasonable Disagreement." In Feldman, Richard and Ted A. Warfield (eds.) *Disagreement*. Oxford: Oxford University Press, pp. 187 – 215.
- Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In Gendler, Tamar Szabo and John Hawthorne (eds.) *Oxford Studies in Epistemology* volume 1, pp. 1167 – 196.

¹⁴ I'm very grateful to Peter Poellner and Hans Bernhard Schmid for their feedback on the paper. I also benefitted from a very helpful referee report.

Kelly, Thomas. 2010. "Peer Disagreement and Higher Order Evidence." In Feldman, Richard and Ted A. Warfield (eds.) *Disagreement*. Oxford: Oxford University Press, pp. 111 – 174.

Larmore, Charles. 2008. *The Autonomy of Morality*. Cambridge: Cambridge University Press.

Longino, Helen. 2002. *The Fate of Knowledge*. Princeton: Princeton University Press.

McMyler, Benjamin. 2011. *Testimony, Trust, and Authority*. New York: Oxford University Press.

Pauer-Studer, Herlinde. 2010. "The Moral Standpoint: First or Second Personal?" *European Journal of Philosophy* 18(2): 296–310.

Peter, Fabienne. 2008. *Democratic Legitimacy*. New York: Routledge.

Peter, Fabienne. 2012. "Epistemic Foundations of Political Liberalism." Unpublished.

Rawls, John. 1971. *A Theory of Justice*. Cambridge: Harvard University Press.

Rawls, John. 1993. *Political Liberalism*. New York: Columbia University Press.

Sosa, Ernest. 2010. "The Epistemology of Disagreement." In A. Haddock, A. Millar, and D. Pritchard (eds.) *Social Epistemology*. Oxford: Oxford University Press, pp. 278 – 297.

Sunstein, Cass. 2002. "The Law of Group Polarization." *Journal of Political Philosophy* 10(2): 175–195.