When Conciliation Frustrates the Epistemic Priorities of Groups

Mattias Skipper · Asbjørn Steglich-Petersen

Penultimate draft. Final version to appear in *The Epistemology of Group Disagreement*, F. Broncano-Berrocal and J. Adam Carter (eds.). Routledge.

Abstract: Our aim in this chapter is to draw attention to what we see as a disturbing feature of conciliationist views of disagreement. Roughly put, the trouble is that conciliatory responses to in-group disagreement can lead to the frustration of a group's *epistemic priorities*: that is, the group's favoured trade-off between the "Jamesian goals" of *truth-seeking* and *error-avoidance*. We show how this problem can arise within a simple belief aggregation framework, and draw some general lessons about when the problem is most pronounced. We close with a tentative proposal for how to solve the problem raised without rejecting conciliationism.

Keywords: Conciliationism, group disagreement, epistemic priorities, belief aggregation, epistemology of disagreement

1. Introduction

Here is a question to which epistemologists have devoted much attention in recent decades: does the disagreement of others give you epistemic reason to reduce your confidence in your own views? Here is a relatively modest answer to this question: "Yes, at least sometimes." Let's stipulate that any epistemological theory of disagreement that entails this answer (or something stronger) counts as a version of conciliationism.

Conciliationism, even in this minimal form, is a controversial view, and we won't here try to make a final judgment on it.[1] Rather, our aim is to draw attention to what we see as a disturbing feature of conciliationism, which has (as far as we're aware) gone unnoticed in the literature. Roughly put, the trouble is that conciliatory responses to

[1] Different versions of conciliationism have been defended by Christensen (2007), Elga (2007), Kelly (2010), Lackey (2008), among others. For critics of conciliationism, see Titelbaum (2015), Tal (forthcoming), Smithies (2019), and Weatherson (2019).

in-group disagreement can lead to the frustration of a group's *epistemic priorities*: that is, the group's favoured trade-off between the "Jamesian goals" of *truth-seeking* and *error-avoidance*. We'll say more about what this "Epistemic Priority Problem" (as we'll henceforth call it) amounts to later on. But before we dive into the details, we'd like to put the problem into a slightly broader context.

One of the most exciting ideas to have emerged from the recent flurry of work in collective epistemology (and one of the main reasons, to our mind, why collective epistemology is an interesting and important field of study in its own right) is that epistemically well-performing individuals might make up epistemically ill-performing groups and that, conversely, epistemically well-performing groups might be made up of epistemically ill-performing individuals. This idea, broadly construed, is what Mayo-Wilson et al. (2011) refer to as the "Independence Thesis." The idea has been defended in various forms not only in the epistemological literature but also in the philosophy of science and social choice theory. For example, Goodin (2006) has argued that biased individuals may be able to pool their information in ways that give rise to unbiased groups; Zollman (2010) has argued that scientists who hold on to their theories despite strong evidence against them may help ensure that the broader scientific community doesn't abandon those theories prematurely; and numerous authors have contributed to the now extensive literature on the "wisdom of crowds" (see Lyon and Pacuit (2013) for an overview).

The Epistemic Priority Problem, as we will describe it, can be seen as lending further support to the Independence Thesis: it illustrates, in yet another way, how a seemingly rational epistemic practice at the individual level—namely, the practice of conciliating with those who disagree with us—can have adverse epistemic effects at the group level. Whether this result constitutes a problem for conciliationism *per se*, or whether it simply shows that the true epistemic norms for individual believers can have adverse epistemic consequences at the group level is not a question that will occupy us much. As we see it, the problem raised is an important one to address, even if it doesn't give us reason to doubt that conciliationism is true.

For the sake of clarity and definiteness, we'll embed our discussion of the Epistemic Priority Problem within a general "belief aggregation" framework (more on this framework below). This is not to suggest that the problem is specific to the aggregation-based way of understanding the relationship between the beliefs of a group and those of its members. Indeed, we suspect that very similar problems will arise for alternative ways of understanding group belief as well (though we won't try to defend this claim here). But the aggregation framework provides a simple and tractable way of making the Epistemic Priority Problem vivid.

Here, then, is the plan of attack. In §2, we begin by covering some basics of the belief aggregation framework, and explain in more detail what we mean by saying that a group can have "epistemic priorities." In §3, we go on to show, in a preliminary way, how, given certain idealizing assumptions, the Epistemic Priority Problem can arise as a consequence of conciliatory responses to in-group disagreement. In §4, we generalize the problem by showing how it can arise even if we relax the various idealizing assumptions. At this point we'll have established our main negative lesson. We close on a more positive note in §5 by offering a tentative proposal for how to solve the problem raised without rejecting conciliationism.

2. Preliminaries on Belief Aggregation

As many authors have pointed out, groups are often said to believe things.[2] For example, a jury might be said to believe that the defendant is guilty; UNESCO might be said to believe that education is a human right; and so on. This raises a natural question: how (if at all) do the beliefs of a *group* relate to those of its *members*?

According to a familiar answer, a group's belief state is (or may be represented as) a *function*, or *aggregate*, of the belief states of its individual members. This "aggregation model" of group belief has featured prominently in the literature on the "doctrinal paradox" and related impossibility results (Kornhauser and Sager 1986; List and Pettit

[2] For some good entry points into the literature on group belief, see Gilbert (1987; 1989), Tuomela (1992), List and Pettit (2011), and Lackey (2016).

2002), but has also been used to investigate questions about, e.g., the epistemic merits of co-authorship (Bright et al. 2017), the nature of group justification (Goldman 2011), and the normative significance of group disagreement (Skipper and Steglich-Petersen 2019).[3] Here we'd like to use the aggregation framework to illustrate how the Epistemic Priority Problem can arise as a consequence of conciliatory responses to disagreement among the members of a group. Below we introduce some nuts and bolts that will facilitate the discussion.

Let a *belief state* be a set of propositions (intuitively: the set of propositions believed by the agent in question), and let a *Belief Aggregation Function* (henceforth just a "BAF") be a function from sets of belief states to single belief states (intuitively: the function taking the belief states of the individual group members as input and returns the belief state of the group as a whole). Familiar BAFs include majority voting, unanimity voting, and dictatorship, but there are many other BAFs that a group might in principle use, and we won't make any limiting assumptions at the outset about which BAFs are admissible. For example, we won't assume that a group believes a proposition, p, only if a large enough proportion of its members believe that p. Moreover, we won't assume that the members of a group must be explicitly aware of which BAF they adhere to. For all we are concerned, a group's BAF might rather be a tacit feature of the group's practice.

To simplify the discussion, we'll assume that each group member must either believe or disbelieve any given proposition (suspension of judgment isn't allowed). Accordingly, we'll count any credence above 50% as a belief, and we'll count any credence below 50% as a disbelief (a credence of exactly 50% isn't allowed). Needless to say, this is a rather strained use of the ordinary term "belief." But as we'll see, nothing of importance is going to turn on this simplification.

Though not uncontroversial, the aggregation model of group belief is a highly general one, and one that comes with very few substantive assumptions about the metaphysical nature of group belief. For example, as hinted at above, it doesn't commit

[3] For a general introduction to the theory of belief aggregation, see Pigozzi (2016).

us to a "summativist" view of group belief, according to which a group's believing that *p* is just a matter of all (or a sufficiently high proportion) of its members' believing that *p*.[4] Indeed, for all the aggregation model says, a group might believe those and only those propositions that its members unanimously agree are false. This would obviously require an odd BAF (one saying that the group believes *p* iff all of the group members believe ~*p*). But, in principle, the aggregation framework is general enough to accommodate any BAF you like.

The aggregation model is also silent on whether groups literally have minds of their own, or whether our talk of group belief should be understood in a less metaphysically committal way. To ease the exposition, we'll continue to talk as if groups, like their members, have genuine beliefs. But all we assume on the official story is that a group's belief state, whatever its metaphysical status, may be usefully represented as a function of the belief states of its members.

Just as individuals can be more or less epistemically reliable, groups can be more or less epistemically reliable as well: that is, their beliefs can "line up" with the truth in more or less accurate ways. They can do so in (at least) two different ways, corresponding to two different kinds reliability. On the one hand, there is an agent's *positive reliability*: that is, the probability that the agent believes that *p* given that *p* is true. On the other hand, there is the agent's *negative reliability*: that is, the probability that the agent doesn't believe that *p* given that *p* is false. Formally (cf. List 2005):

Positive reliability: $Pr(Bp|p)$
Negative reliability: $Pr(\sim Bp|\sim p)$

As William James (1896) famously pointed out, these two kinds of reliability do not always go hand in hand. In fact, they can come arbitrarily far apart. A highly credulous agent—someone who is willing to believe even the most improbable of propositions—will have a very high positive reliability, but a very low negative reliability. Such an agent will rarely miss out the truth, but at the cost of forming lots

[4] See Gilbert (1989) for an early discussion of summativist *vs* non-summativist views of group belief.

of false beliefs. Conversely, a highly incredulous agent—someone who is unwilling to believe even the most probable of propositions—will have a very low positive reliability, but a very high negative reliability. Such an agent will rarely form false beliefs, but at the cost of often missing out on the truth.

Here is an uncontroversial fact that will be important for what follows: different BAFs contribute in different ways to a group's positive and negative reliability. To take a simple example, consider a group with n members, where each member has the same positive and negative reliability, r. Given this, the group's positive and negative reliability will vary quite significantly, depending on what BAF the group uses. For example, unanimity voting will tend to yield a much higher negative reliability than majority voting, whereas majority voting will tend to yield a much higher positive reliability than unanimity voting (see Table 1 and Figure 1).[5]

We take this to suggest that what BAF it is advisable for a group to use depends, at least in part, on the group's *epistemic priorities*: that is, the group's preferred trade-off between *believing what is true* and *not believing what is false* (or: the group's preferred trade-off between positive reliability and negative reliability). For example, if the group described above (consisting of n members with identical positive and negative reliability, r) places more weight on negative reliability than positive reliability, unanimity voting will be preferable to majority voting. By contrast, if the group places equal weight on positive reliability and negative reliability, majority voting will be preferable to unanimity voting.

So far, so good. But what is the *right* trade-off between positive and negative reliability? In other words, what epistemic priorities *should* a group have? A natural first reaction to this question would be to say, on grounds of uniformity, that groups should simply have whatever epistemic priorities *individuals* should have. And many epistemologists have been inclined to think that individuals should be epistemically *risk-averse*, that is, that they should place more weight on error-avoidance than on

[5] Here and elsewhere we assume that the group members are independent of each other: that is, any given member's belief about p isn't affected by any other member's belief about p.

truth-seeking.[6] If so, it would be natural to think that groups should likewise be epistemically risk-averse, that is, that they should place more weight on negative reliability than on positive reliability.

BAF	Positive reliability	Negative reliability
Unanimity voting	r^n	$1-(1-r)^n$
Majority voting	$\sum_{i=.5(n+1)}^{n} \frac{n!}{i!(n-i)!} r^i (1-r)^{n-i}$	

Table 1: *A group's positive and negative reliability as a function of its size (n) and member reliability (r), depending on whether the group uses unanimity or majority voting.*

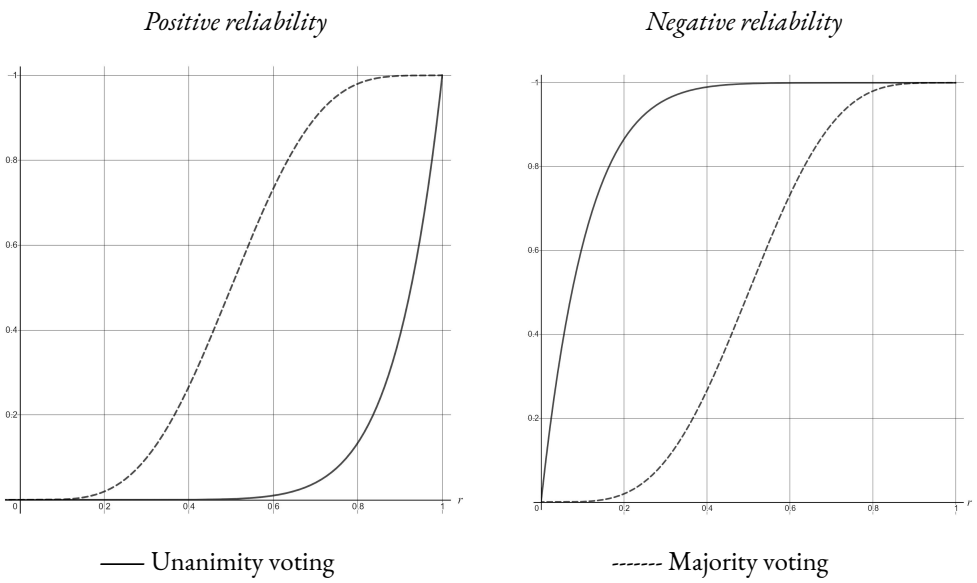

Figure 1: *A group's positive reliability (left graph) and negative reliability (right graph) as a function of the members' reliability, r (setting n = 9, for illustration).*

However, we don't think this parity between individuals and groups can be easily maintained. Note that the idea that individuals should place more weight on error-avoidance than on truth-seeking is usually motivated on distinctly *epistemic* grounds, e.g., by appeal to considerations about the irrationality of contradictory

[6] See, e.g. Dorst (2019) and Easwaran (2016). See also Skipper (forthcoming), Steinberger (2019), and Hewson (forthcoming) for critical discussion of this idea.

beliefs (Dorst 2019, p. 185), the rationality of suspending judgment (Easwaran 2016, p. 824), or the rationality of imprecise credences (Konek forthcoming). By contrast, the kinds of considerations that are naturally taken to bear on questions about what BAF it would be advisable for a group to use are often *practical* in nature. Here are three hypothetical (but, we hope, not too far-fetched) examples:

> **Criminal Trial:** The jury in a criminal trial must reach a collective verdict about whether the defendant is guilty. The jury is required to deem the defendant guilty iff all of the jurors believe that the defendant is guilty beyond reasonable doubt. This unanimity procedure is justified on the grounds that it is much more important to avoid punishing the innocent than it is to punish the guilty.
>
> **Quiz Show:** A group of friends appear on a quiz show. The host asks the group to collectively answer *yes*, *no*, or *pass* in response to a series of questions. Each right answer gives +1 point, each wrong answer gives -1 point, and no points are awarded or subtracted if the group says "pass." The friends decide to base their answers on simple majority voting on the grounds that they suspect this to maximize their expected score.
>
> **Anti-Terror Unit:** An anti-terror police unit must decide whether to treat an apparent threat as real. The unit conforms to a policy of treating apparent threats as real as long as at least *one* member of the unit believes that the threat is real. The rationale behind this policy is that it's much worse to treat a real threat as merely apparent than to treat a merely apparent threat as real.

In each case, the group's choice of BAF seems perfectly reasonable given the circumstances in which the group is to form a collective verdict. Yet, only in the first case does the group place more weight on negative reliability than on positive reliability. In the two other cases, the group places at least as much weight on positive reliability as on negative reliability.

We take this to suggest that, whatever might be said about the epistemic priorities of individuals, there isn't is a unique trade-off between positive and negative reliability that groups should always try to make. It seems much more plausible to suppose, as we'll henceforth do, that different contexts call out for different epistemic priorities, and hence different BAFs.[7]

3. The Epistemic Priority Problem

With these preliminaries in place, we are now ready to show how the Epistemic Priority Problem can arise within a simple aggregation framework. We begin by making some additional idealizing assumptions, which will help to simplify the exposition (the assumptions will be relaxed later on).

First, we assume that the *agenda*—that is, the set of propositions that the group members are to form beliefs about—consists of just a single proposition, *p*. This allows us to sidestep certain problems that can arise when the agenda contains two or more logically interconnected propositions (as exemplified in the famous *doctrinal paradox*).[8] These problems are of obvious interest and importance in their own right, but they are orthogonal to our present concerns.

Second, we assume that all group members are (and consider themselves to be) *epistemic peers* with respect to *p*. We will take this to mean that all group members have the same positive and negative reliability: that is, no two members differ in their positive reliability, and no two members differ in their negative reliability.[9] Furthermore, we'll assume that each group member has the same positive and negative

[7] There are also less practically founded reasons to think that no one BAF will fare well in all contexts; reasons coming from various impossibility results in social choice theory (see Pacuit (2019) and List (2013) for overviews). See also Kelly (2013), Horowitz (2017), and Pettigrew (2016) for recent discussions of how rational individuals might trade off the Jamesian goals of truth-seeking and error-avoidance in different ways.

[8] See List and Pettit (2002).

[9] One might instead adopt an "evidentialist" notion of epistemic peerhood, whereby two agents are said to be epistemic peers with respect to *p* iff they have the same evidence about *p* and are equally competent at judging how that evidence bears on *p*. This notion of epistemic peerhood has often been operative in the "peer disagreement" literature (e.g., Christensen 2007). As previously noted, we suspect that the Epistemic Priority Problem (or a very similar problem) will arise equally for such an evidentialist notion of peerhood, but we won't here try to defend this claim in any detail.

reliability, r, where $r > 50\%$ (which is just to say that the group members are at least slightly more reliable than the flip of a fair coin).

Third, we assume that all group members practice a particularly strong form of conciliationism akin to the familiar "Equal Weight View" defended by Christensen (2007), Elga (2007), and others. More specifically, we'll assume that the group members "split the difference" in response to peer disagreement. This is clearly not the only available interpretation of the Equal Weight View, nor perhaps the most plausible one.[10] But the exact interpretation of the Equal Weight View won't matter for present purposes. As we'll see, the Epistemic Priority Problem can in any case arise in much less conciliatory environments.

Fourth, we assume that the reliability of any given member doesn't depend on how confident that member is that his or her opinion is correct. In other words, members with *more* extreme credences are assumed to be neither more nor less reliable than members with *less* extreme credences. This assumption may seem egregiously unrealistic. (Wouldn't it be more realistic to assume that people who are highly confident in their opinion on some matter are also more likely to be correct about that matter?) But the assumption isn't meant to be realistic, since (as we'll see) the Epistemic Priority Problem can in any case arise without it. For now, we are just looking to make things as simple as possible.

Finally, we'll work with a very sparse set of possible degrees of confidence (or "credences"). More specifically, we will assume that each group member is either *slightly more* confident of p than $\sim p$ (which we'll write as "$Cr(p) > Cr(\sim p)$") or *much more* confident of p than $\sim p$ (which we'll write as "$Cr(p) \gg Cr(\sim p)$"). The converse is obviously also allowed: members may be *slightly less* confident of p than $\sim p$ (written "$Cr(p) < Cr(\sim p)$") or *much less* confident of p than $\sim p$ (written "$Cr(p) \ll Cr(\sim p)$"). One small complication of this way of modelling credences is that it doesn't involve real numbers, which means that the idea of "splitting the difference" can't be taken to mean "taking the arithmetic mean." All we need to assume in what follows, however, is

[10] See, e.g., Fitelson and Jehle (2009) and Rasmussen et al. (2018).

that if you're *much less* confident of p than $\sim p$ while I'm *slightly more* confident of p than $\sim p$, then the way for us to split the difference is by both becoming *slightly less* confident of p than $\sim p$. And, conversely, if you're *much more* confident of p than $\sim p$ while I'm *slightly less* confident of p than $\sim p$, then the way for us to split the difference is by both becoming *slightly more* confident of p than $\sim p$.

Taken together, these assumptions make for a highly idealized setting in which to study the Epistemic Priority Problem. This should not be taken to suggest that the Epistemic Priority Problem is a mere theoretical curiosity with little practical relevance. As already mentioned, we'll eventually argue that the problem can arise under much less idealized circumstances as well. But we'd like to begin by showing how the problem can arise in a very clean and simple setting.

To show this, we'll proceed in a case-based manner. Each of the cases will feature a group whose members start out with a given set of individual credences in p—we'll call them their *pre-conciliation* credences. The group will then undergo a *conciliation process*, whereby the members learn about each others' credences and respond to any potential disagreements by splitting the difference in the way described. As a result, the members end up with identical credences in p once the conciliation process is completed—we'll call them their *post-conciliation* credences.

That's the basic setup. Now for the cases (of which there are three).

Case 1: Unanimity Voting

Consider a group with the following characteristics (in addition to those listed above):

(i) The group consists of two same-sized subgroups, g_1 and g_2.

(ii) The group members' pre-conciliation and post-conciliation credences in p are as stated in Table 2.

(iii) The group uses unanimity voting: that is, the group believes p iff all of its members believe p.

Let's begin by asking: what does the group believe about p before and after the conciliation process? Before the conciliation process, all members of g_1 are more

confident of p than $\sim p$, and so they all believe p. By contrast, all members of g_2 are more confident of $\sim p$ than p, and so they all believe $\sim p$. Hence, due to the lack of unanimity, the group as a whole neither believes p nor believes $\sim p$.

However, note that the members of g_1 all start out being *much more* confident of p than $\sim p$, whereas the members of g_2 start out being only *slightly more* confident of $\sim p$ than p. As a result, the members of g_1 and g_2, after having conciliated with each other, all end up being slightly more confident of p than $\sim p$. Consequently, the group as a whole ends up believing p after the conciliation process is completed.

	g_1	g_2
Pre-conciliation	$Cr(p) \gg Cr(\sim p)$	$Cr(p) < Cr(\sim p)$
Post-conciliation	$Cr(p) > Cr(\sim p)$	

Table 2: *Before the conciliation process, all members of g_1 are much more confident of p than $\sim p$, whereas all members of g_2 are slightly more confident of $\sim p$ than p. After the conciliation process, all group members are slightly more confident in p than $\sim p$.*

This change in the group's belief state may not seem like much of a problem. But consider what has happened to the group's positive and negative reliability, respectively. Recall that the distinctive feature of unanimity voting is that it secures a high negative reliability in comparison to other BAFs (e.g., in comparison to majority voting, as illustrated by Figure 1). In other words, unanimity voting is supposed to be an effective guard against false belief. But the conciliation process puts a crack in the guard: it leads the group to form a new belief which, in turn, introduces a new error-possibility. Now, the mere introduction of a new error-possibility obviously isn't enough to show that the group's negative reliability has decreased as a result of the conciliation process. We also need to consider whether any existing error-possibilities have been eliminated. This would be the case if the conciliation process made the group drop an existing belief. But since the group starts out neither believing p nor

believing ~p, there is no such belief to be dropped.[11] This is the qualitative reason why the conciliation process harms the group's negative reliability: it introduces a new error-possibility without eliminating any existing ones.

How significant is this problem from a quantitative point of view? To get a feel for this, let's put some numbers on the table. Suppose that g_1 and g_2 each have 5 members (n = 10), and suppose that all members have a positive reliability and negative reliability of 70% (r = .7). We can then ask: once the conciliation process is completed, how likely is it that the group's belief that p is false? In other words, how likely is it that p is false given the members' post-conciliation beliefs about p? On the face of it, this might seem like an intractable question, since we haven't said anything about how the conciliation process might affect the reliability of the group members. But we can approach the question in a more indirect way, by considering a slightly different question: how likely is p to be false given the members' pre-conciliation beliefs about p? This is a question that we *can* answer. But before we do, let's explain why the two questions must have the same answer.

Suppose that you're a third party—someone not a member of the group in question—who seeks to use the group members' beliefs about p as evidence bearing on p. And let's say that, upon learning about the group members' pre-conciliation beliefs about p, you should have such-and-such a credence in p. Now suppose you learn that the group members have been through a conciliation process: that is, you learn that the group members have adopted their average credence in p (and that's *all* you learn). Should you revise your credence in p in light of this new piece of information? It seems not. After all, the mere fact that the group members have conciliated doesn't seem to have any bearing on whether p is true or false. This suggests that there is no difference between, on the one hand, the probability that p is false given the group members' *post-conciliation* beliefs about p, and, on the other hand, the probability that p is false

[11] More generally: as long as the group uses unanimity voting, it's impossible for the conciliation process to eliminate any existing error-possibilities, since, if the members unanimously agree on a proposition prior to the conciliation process, they will also unanimously agree on that proposition after the conciliation process.

given the group members' *pre-conciliation* beliefs about *p*. Hence, to determine the former probability, we need only determine the latter.

What, then, *is* the probability that *p* is false given the group members' pre-conciliation beliefs about *p*? In the case at hand, the answer is simply "50%" since there are equally many, equally reliable members who believe *p* and ~*p*, respectively.[12] Thus, the group ends up with a belief that, by the lights of its own members, is no more likely to be true than false. This already looks like a severe blow to the group's negative reliability.

We can harden the blow even more by considering just how unlikely ~*p* would have to be in order for the group to believe *p before* the conciliation process. The relevant scenario is one where all ten group members falsely believe *p*, which, in spite of their relatively modest reliability, is extremely unlikely: $(1 - r)^n = (1 - .7)^{10} = .000006$. Thus, the conciliatory effects of in-group disagreement can in fact have a very significant, adverse impact on a group's negative reliability.

There is, however, a positive flip side: the *decrease* in negative reliability is accompanied by an *increase* in positive reliability. The qualitative reason is the same as above: given that the group uses unanimity voting, it isn't possible for the conciliation process to eliminate any existing, potentially true beliefs. By contrast, it *is* possible for the conciliation process to introduce a new, potentially true belief. Hence, the group's positive reliability goes up.

This also brings out a more general lesson about the problem we're facing. The problem isn't so much that the conciliatory effects of in-group disagreement can harm a group's *overall* reliability (although this may sometimes be the case, depending on how we determine an agent's "overall" reliability on the basis of the agent's positive reliability and negative reliability). Rather, the problem is that the conciliatory effects of in-group disagreement can lead to the frustration of a group's preferred trade-off between positive and negative reliability. That's why we began by naming it the "Epistemic Priority Problem."

[12] Assuming that the prior probability of *p* is 50%.

Case 2: Inverse Unanimity Voting

The same general problem can arise for groups that place more weight on positive reliability than negative reliability. Consider a group with the same characteristics as the one above except that it uses a different BAF:

(i) The group consists of two same-sized subgroups, g_1 and g_2.

(ii) The group members' pre-conciliation and post-conciliation credences in p are as stated in Table 2.

(iii) The group uses (what we'll call) "inverse" unanimity voting: that is, the group believes p iff at least one of its members believes p.

The operative BAF here—inverse unanimity voting—is less familiar than, say, unanimity voting or majority voting. It has some rather odd properties. For example, it entails that a group believes both p and $\sim p$ whenever its members do not unanimously agree about whether p is true or false. Yet, this is precisely what makes inverse unanimity voting conducive to a high positive reliability: just as unanimity voting is an effective way of avoiding false beliefs, inverse unanimity voting is an effective way of gaining true ones.

Let's ask again: what does the group believe about p before and after the conciliation process, respectively? Before the conciliation process, the group believes both p and $\sim p$, since the members do not unanimously agree about whether p is true or false (more specifically: all members of g_1 believe p, whereas all members of g_2 believe $\sim p$). But after the conciliation process is completed, the members unanimously agree that p is true. As a result, the group drops its belief in $\sim p$, but retains its belief in p.

What has happened to the group's positive and negative reliability here? On the one hand, the group's positive reliability has decreased, since the group has dropped a potentially true belief without forming any new ones. On the other hand, the group's negative reliability has increased, since the group has eliminated an existing

error-possibility without introducing any new ones. Thus, the group's epistemic priorities are once again frustrated.[13]

Case 3: Majority Voting

What about groups that place *equal* weight on positive and negative reliability? Can the problem arise for such groups as well? The short answer is "yes." But the details are a bit different from the previous two cases. Consider a group with the following characteristics:

(i) The group consists of two subgroups, g_1 and g_2, where g_1 has 4 members, and g_2 has 5 members.

(ii) The group members' pre-conciliation and post-conciliation credences in p are as stated in Table 2.

(iii) The group uses majority voting: that is, the group believes p iff more than half of its members believe p (which secures an equal weighing of positive and negative reliability, as illustrated by Figure 1).

What does the group believe about p before and after the conciliation process, respectively? Before the conciliation process, the group believes p, since all members of the majority group, g_2, believe p. But since the members of the minority group, g_1, are *much more* confident of $\sim p$ than p, whereas the members of g_2 are only *slightly more* confident of p than $\sim p$, the result of the conciliation process is that all members of the combined group end up being slightly more confident of $\sim p$ than p. So, after the conciliation process is completed, the group believes $\sim p$.

What has happened to the group's reliability here? Consider first the group's *negative* reliability. One effect of the conciliation process is that the group drops its belief in p, which eliminates an existing error-possibility. But the group also forms a new belief in $\sim p$, which introduces a new error-possibility. These opposing effects

[13] For a quantitative example pertaining to Case 2, the calculations provided in connection with Case 1 carry over, *mutatis mutandis*, to the present case. We omit the details.

might be thought to "cancel each other out," so as to leave the group's negative reliability unaffected. But things are a little more complicated than that.

Here is why: before the conciliation process, the majority group is more likely than the minority group to be right (assuming, as we do, that $r > 50\%$).[14] Thus, the conciliation process effectively leads the group to trade a belief that is *less likely* to be false for a belief that is *more likely* to be false, which means that the group's negative reliability decreases. The same goes for the groups' positive reliability: it also decreases, since the group effectively trades a belief that is *more likely* to be true for a belief that is *less likely* to be true. This stands in contrast to the previous two cases, where the group's positive/negative reliability decreased, whereas the group's negative/positive reliability increased, thereby leaving the group's *overall* reliability (at least potentially) untouched.[15]

However, this result should be taken with a pinch of salt. As we'll see in the next section, the result is sensitive to our background assumptions in a way that the previous two results are not. (More specifically: it doesn't fully generalize to settings where a high level of confidence is indicative of a high reliability.) Thus, we still take the main upshot of the foregoing considerations to be that the conciliatory effects of in-group disagreement can lead to the frustration of a group's epistemic priorities (rather than necessarily damage the group's overall reliability).

4. Generalizing the Epistemic Priority Problem

We have now seen how the Epistemic Priority Problem can arise in a highly idealized setting. The next thing we'd like to do is to generalize the problem by showing how it can arise even without the various idealizing assumptions introduced in the previous section. We will skip over some of the assumptions that clearly aren't responsible for

[14] This is a consequence of Condorcet's famous jury theorem (Condorcet 1785). For an accessible modern discussion of the result and its implications, see Goodin and Spiekermann (2018).

[15] A different but related point has been made by Hazlett (2016), who argues that the probability with which majority voting yields the correct result may be harmed if the voters defer to each other's beliefs prior to voting, because this compromises the "independence" assumption underlying the Condorcet jury theorem.

the problem (e.g., the assumption that suspension of judgment isn't allowed, and the assumption that there are only two levels of comparative confidence). This leaves us with three assumptions to consider.

The first assumption is the one saying that the reliability of any given member is independent of how confident that member is that his or her opinion is correct. There are two general ways in which this assumption might be modified: either (i) by assuming that members with more extreme credences are more reliable than members with less extreme credences, or (ii) by assuming that members with more extreme credences are less reliable than members with less extreme credences. While the latter of these dependencies might well obtain in certain kinds of situations,[16] we'll focus our attention on the former dependency here, since this carries no obvious presumption of irrationality on part of the individual group members.[17]

So, let's assume that members with more extreme credences are also more reliable, and let's ask: how, if at all, does this change affect the results from the previous section (other things being equal)? There are three cases to consider.

In Case 1, the Epistemic Priority Problem still shows up, albeit with mitigated strength. The basic mechanism is the same as before: the group's negative reliability decreases, since the group forms a new belief, which introduces a new error-possibility. But since the members of g_1 are more confident (and hence, by present assumptions, more reliable) than the members of g_2, the group's negative reliability doesn't suffer as much as before. In particular, the group now ends up with a belief that, by the lights of its own members, is at least slightly more likely to be true than false (unlike the original

[16] For example, one might wonder whether the dependency sometimes obtains as a consequence of the well-documented "Dunning-Kruger" effect, whereby (roughly) people who are less competent on a given matter are more prone to overestimating their own competence (Dunning and Kruger 1999).

[17] A small aside on this point: psychological studies have documented a robust and widespread "overconfidence bias," whereby people's confidence in their answers to a wide range of tests tends to exceed the actual frequency with which their answers are correct (Lichtenstein et al. 1982; Hoffrage 2004). In other words, of the answers people are *n*% confident in, less than *n*% are true. One might be tempted to see this overconfidence effect as evidence *against* the claim that a high confidence is typically indicative of a high reliability. However, this would be too quick. Something much stronger would be needed to show this, namely, that if we compare the answers people are *n*% confident to the answers they are >*n*% confident in, a higher proportion of the former answers are true. As far as we know, there is no evidence to support this stronger claim.

case where the group ended up with a belief that, by the lights of its own members, was no more likely to be true than false).

The same goes, *mutatis mutandis*, for Case 2: the group's positive reliability still decreases, since the group drops a belief, which eliminates an existing possibility of being right. But since the members of g_1 are more confident (and hence, by present assumptions, more reliable) than the members of g_2, the group's positive reliability doesn't suffer as much as before.

By contrast, the Epistemic Priority Problem need no longer arise in Case 3. The reason for this is that the members of the minority group, g_1, are (by present assumptions) more reliable than the members of the majority group, g_2, which means that it's no longer clear that the majority group is initially (that is, prior to the conciliation process) more likely to be right than the minority group. Rather, what subgroup is more likely to be right is going to depend on how much more reliable the members of the g_1 are assumed to be than the members of g_2. Thus, although the Epistemic Priority Problem will still arise on *some* ways of filling in the details of the case, the problem is no longer inevitable.

The next assumption we'd like to consider is the one saying that all group members are epistemic peers. As many authors have pointed out, this condition is rarely (if ever) met in real life.[18] It is therefore natural to wonder whether the Epistemic Priority Problem is affected (one way or the other) by relaxing the peerhood assumption. So, let's assume that the group members may differ in reliability, and let's also assume (which seems reasonable from an epistemic viewpoint) that members who are more reliable are also accorded more weight by the group's BAF. How, if at all, does this change affect the Epistemic Priority Problem (other things being equal)?

In Case 1, the Epistemic Priority Problem still arises with unmitigated strength. The reason for this is that the introduction of differential weights has no effect on the output of unanimity voting: all members still have to agree on p in order for the group to believe p. In consequence, the group's negative reliability still decreases as a result of

[18] See, e.g., King (2012).

the group forming a new belief, which introduces a new error-possibility. And, at least insofar as there is no reason to think that the members of g_1 are systematically more or less reliable than the members of g_2, there is no reason to think that the group's negative reliability suffers any more or less than in the original case.

The same goes, *mutatis mutandis*, for Case 2: the introduction of differential weights has no effect on the output of inverse unanimity voting, which means that the group's positive reliability still decreases as a result of dropping an existing, potentially true belief. And given that there is no reason to think that the members of g_1 are systematically more or less reliable than the members of g_2, there is no reason to think that the group's positive reliability suffers any more or less than in the original case.

Things get a bit more complicated in Case 3, since the introduction of differential weights *can* affect the output of majority voting. Whether it *does* affect the output in the case at hand depends on the specific weight allocation. It would take us too far astray to enter a detailed discussion of how different types of weight allocation would affect the Epistemic Priority Problem. But we'd like to consider one particularly natural weight allocation (or class of weight allocations) which turns out to have the potential to mitigate the Epistemic Priority Problem, at least to some extent. On this way of allocating weight, members with more extreme credences are given more weight than members with less extreme credences. The rationale behind this weight allocation is supposed to be that people who are more confident in their beliefs are also more likely to be right in their beliefs. As mentioned, this dependency might not always hold true. But we find it realistic enough in many cases for it to be worthwhile considering how the Epistemic Priority Problem might be affected by it.

The first thing to observe is that it's no longer clear what the group believes before the conciliation process, since the members of the minority group have more extreme credences (and hence, given present assumptions, are given more weight) than the members of the majority group. Rather, whether the group as a whole initially (that is, prior to the conciliation process) agrees with the majority group or the minority group

is going to depend on how much more weight is placed on the beliefs of the minority group than on the members of the majority group. This gives us two cases to consider.

The first (and simplest) case is the one where the group as a whole initially agrees with the minority group. Here it's clear that the Epistemic Priority Problem no longer arises, since the conciliation process leads to no change in the group's belief state, and hence leaves the group's positive and negative reliability unaffected.

The second (and slightly more complicated) case is the one where the group as a whole initially agrees with the majority group. Given this, the conciliation process *does* lead to a change in the group's belief state, since all members still end up agreeing with the minority belief once the conciliation process is completed. However, given that the members of the minority group are more reliable than those of the majority group, it's not immediately clear whether the group's reliability increases or decreases. Rather, whether the group's reliability increases or decreases depends on whether the majority group is initially more or less likely to be right than the minority group. And this, in turn, depends on just how much more reliable the members of the minority group are assumed to be than those of the majority group. Thus, the Epistemic Priority Problem may or may not arise, depending on how we fill in the details of the case.

The third (and final) assumption we want to discuss concerns the particular version of conciliationism practiced by the group members. Until now, we have assumed that the group members practice a form of "splitting the difference." However, there are various weaker versions of conciliationism which have been defended in the literature (perhaps the best-known example being Kelly's (2010) "Total Evidence View"). This makes it natural to wonder whether the Epistemic Priority Problem can also arise in more moderate conciliatory environments. So, let's suppose that the group members practice a moderate form of conciliationism: that is, they don't split the difference, but they do revise their credence at least to some extent in the face of peer disagreement. How, if at all, does this change affect the Epistemic Priority Problem (other things being equal)?

The answer, to a first approximation, is the same in all three cases: the Epistemic Priority Problem can still arise, but it does so in a more limited range of cases. A little more precisely: the Epistemic Priority Problem still shows up as long as the members practice a form of conciliationism that is strong enough to ensure that all members end up favoring the same proposition once the conciliation process is completed. What counts as "strong enough" is going to depend on the pre-conciliation credences of the group members. For example, in Case 1, a fairly weak form of conciliationism will suffice to generate the Epistemic Priority Problem, since all members of g_1 are *much more* confident of p than $\sim p$, while all members of g_2 are only *slightly more* confident of $\sim p$ than p. By contrast, if some of the members of g_2 had instead been *much more* confident of $\sim p$ than p, we would have needed a stronger form of conciliationism to generate the problem. Thus, although the Epistemic Priority Problem is most prevalent in highly conciliatory environments, it can arise in more moderate conciliatory environments as well.

5. Solving the Epistemic Priority Problem

Although not the main focus of the chapter, we'd like to close on a more positive note by offering a tentative proposal for how to solve the Epistemic Priority Problem without rejecting conciliationism. The proposal relies on a distinction that has come up in various forms in the recent "higher-order evidence" literature: a distinction between, on the one hand, an agent's *credence* in p, and, on the other hand, (what we'll call) the agent's *first-order judgment* as to whether p.

An agent's first-order judgment as to whether p, as we'll understand it, is the agent's judgment of how likely it is that p is true given the first-order evidence available to the agent.[19] We won't here try to say anything very precise about what counts as "first-order evidence," but, as a minimum, the mere fact that someone disagrees with you is supposed to *not* count as first-order evidence, but is rather supposed to count as

[19] Variations on the notion of a first-order judgment have been employed by, e.g., Barnett (2019, §4) who talks about your "disagreement-insulated inclination" toward p, and Worsnip (ms, §4) who talks about your "personal take" on whether p.

a (higher-order) reason for you to doubt the reliability of your judgment of what your first-order evidence supports.[20]

Now, in many cases your credence in *p* will line up (at least roughly) with your first-order judgment as to whether *p*. For example, if it seems to you on the basis of your visual experience that it's raining outside, you will normally be quite confident that it's raining outside. Sometimes, however, your credence may come apart from your first-order judgment, precisely because you have reason to doubt the accuracy of your own first-order judgment. This is what can happen in cases of disagreement.

Suppose, for example, that you disagree with a trusted colleague about how strongly a given body of meteorological data supports the proposition that (*p*) it's going to rain this afternoon. In your judgment, the data strongly supports *p*. In your colleagues judgment, the data strongly supports ~*p*. Setting aside the fact that your colleague disagrees with you on this particular occasion, you don't consider your judgment to be any more or less likely to be accurate than your colleague's. Thus, you adopt a relatively low credence in *p* (say, around 50%), not because you have been persuaded by your colleagues first-order considerations, but because the disagreement itself has led you to doubt the accuracy of your own first-order judgment.

Distinction in hand, here is the proposal in rough outline: rather than aggregating the group members' *credences* in *p*, let's instead aggregate their *first-order judgments* as to whether *p*. Doing so would block the Epistemic Priority Problem by preventing the conciliatory effects of in-group disagreement from having any impact on the group's belief state in the first place (since the group members' first-order judgments are not supposed to be sensitive to higher-order considerations). And it would at the same time allow the group to take advantage of various other deliberative activities like knowledge sharing and critical argumentation (since the group members' first-order judgments are supposed to be sensitive to first-order considerations).

[20] For more detailed characterizations of the distinction between "first-order evidence" and "higher-order evidence," see Christensen (2010), Lasonen-Aarnio (2014), and Skipper (2019; ms).

The hope, then, is that by aggregating first-order judgments rather than credences, we can at once (i) avoid the Epistemic Priority Problem, (ii) retain conciliationism, and (iii) reap the epistemic benefits of group deliberation. Needless to say, there are various concerns one might have about the concrete implementation of this proposal. Most obviously, it is not immediately clear how easy it will be to elicit people's first-order judgments in real-world settings (say, a typical voting scenario). We are not ourselves in a position to give an informed assessment of the practical feasibility of the proposed solution. For now, we are content to leave the proposal on the table for our joint consideration.

6. Conclusion

Here, then, is the main takeaway: conciliatory views of disagreement have a disturbing feature. The trouble is that conciliatory responses to in-group disagreement can lead to the frustration of a group's *epistemic priorities*: that is, the group's favoured trade-off between the "Jamesian goals" of truth-seeking and error-avoidance. This is what we called the "Epistemic Priority Problem." The problem is most prevalent in highly conciliatory environments, but it can in principle arise whenever the members of a group practice at least a minimal form of conciliationism. Thus, we take the problem raised to flow from all versions of conciliationism, albeit with different severity.

As mentioned at the outset, this is not to say that conciliationism, understood as a view about how individuals should revise their beliefs in response to disagreement, is undermined (partly or wholly) by the Epistemic Priority Problem. The considerations put forth in this chapter might just show that the true epistemic norms for individual believers sometimes have adverse epistemic consequences at the group level. If so, a solution along the lines of the one outlined in §5 may be particularly apt, since it doesn't force us to give up conciliationism. But in any case, it seems to us that we need to face up to the problem raised in one way or another.

Acknowledgments: We'd like to thank Fernando Broncano-Berrocal and Adam Carter for providing very helpful comments on an earlier version of this chapter.

References

Barnett, Z. (2019): "Philosophy Without Belief." In: *Mind* 128, pp. 109–38.

Bright, L., H. Dang, and R. Heesen (2018): "A Role for Judgment Aggregation in Coauthoring Scientific Papers." In: *Erkenntnis* 83, pp. 231–52.

Christensen, D. (2007): "Epistemology of Disagreement: The Good News." In: *The Philosophical Review* 116, pp. 187–217.

——— (2010): "Higher-Order Evidence." In: *Philosophy and Phenomenological Research* 81, pp. 185–215.

Condorcet, N. (1785): *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*, Paris.

Dorst, K. (2019): "Lockeans Maximize Expected Accuracy." In: Mind 138, pp. 175–211.

Dunning, D. and J. Kruger (1999): "Unskilled and Unaware of it: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessment." In: *Journal of Personal and Social Psychology* 77, pp. 1121–34.

Easwaran, K. (2016): "Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities." In: *Noûs*, 50, pp. 816–853.

Elga, A. (2007): "Reflection and Disagreement." In: *Noûs* 41, pp. 478–502.

Fitelson, B. and D. Jehle (2009): "What is the 'Equal Weight View'?" In: *Episteme* 6, pp. 280–93.

Gilbert, M. (1989): *On Social Facts*. Routledge.

——— (1987): "Modelling Collective Belief." In: *Synthese* 73, pp. 185–204.

Goldman, A. (2011): "Social Process Reliabilism: Solving Justification Problems in Collective Epistemology." In: *Essays in Collective Epistemology*, J. Lackey (Ed.), Oxford University Press.

Goodin, R. (2006): "The Epistemic Benefit of Multiple Biased Observers." *Episteme* 3, pp. 166–74.

Goodin, R. and K. Spiekermann (2018): *An Epistemic Theory of Democracy*. Oxford: Oxford University Press.

Hazlett, A. (2016): "The Social Value of Non-Deferential Belief." In: *Australasian Journal of Philosophy* 94, pp. 131–51.

Hewson, M. (forthcoming): "Accuracy Monism and Doxastic Dominance: Reply to Steinberger." In: *Analysis*.

Hoffrage, U. (2004): "Overconfidence." In: *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*, pp. 235–254.

Horowitz, S. (2017): "Epistemic Utility and the Jamesian Goals." In J. Dunn and K. Ahlstrom-Vij (eds.), *Epistemic Consequentialism*. Oxford University Press.

James, W. (1896): "The Will to Believe." In: *The Will to Believe: And Other Essays in Popular Philosophy*, Cahn (ed.), NY: Longmans, Green, and Co., pp. 1–15.

Kelly, T. (2010): "Peer Disagreement and Higher-Order Evidence." In: *Social Epistemology: Essential Readings*, A. Goldman and D. Whitcomb (eds.). Oxford University Press, pp. 183–217.

——— (2013): "Evidence Can Be Permissive." In M. Steup, J. Turri and E. Sosa (eds.), *Contemporary Debates in Epistemology*. Blackwell.

King, N. (2012): "Disagreement: What's the Problem? *or* A Good Peer is Hard to Find." In: *Philosophy and Phenomenological Research* 85, pp. 249–72.

Konek, J. (forthcoming): "Epistemic Conservativity and Imprecise Credence." In: *Philosophy and Phenomenological Research*.

Kornhauser, L. and L. Sager (1986): "Unpacking the Court." In: *Yale Law Journal* 96, pp. 82–117.

Lackey, J. (2016): "What is Justified Group Belief?" In: *The Philosophical Review* 125, pp. 341–96.

——— (2008): "A Justificationist View of Disagreement's Epistemic Significance." In: *Social Epistemology*, A. Millar, A. Haddock, and D. Pritchard (eds.), Oxford University Press, pp. 145–54.

Lasonen-Aarnio, M. (2014): "Higher-Order Evidence and the Limits of Defeat." In: *Philosophy and Phenomenological Research* 88, pp. 314–345.

Lichtenstein, S., B. Fischhoff, and L.D. Phillips (1982): "Calibration of Probabilities: The State of the Art to 1980." In: *Judgment under Uncertainty*, Kahneman, Slovic, and Tversky (eds.), pp. 306–334. Cambridge University Press.

List, C. (2013): "Social Choice Theory." *Stanford Encyclopedia of Philosophy*, Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2013/entries/social-choice/>.

List, C. and P. Pettit (2011): *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.

——— (2002): "Aggregating Sets of Judgments: An Impossibility Result." In: *Economics and Philosophy* 18, pp. 89–110.

Lyon, A. and E. Pacuit (2013): "The Wisdom of Crowds: Methods of Human Judgement Aggregation." In: *Springer Handbook for Human Computation*, Springer New York, pp. 599–614.

Mayo-Wilson, C., K. Zollman, and D. Danks: "The Independence Thesis: When Individual and Social Epistemology Diverge." In: *Philosophy of Science* 78, pp. 653–77.

Pacuit, E. (2019): "Voting Methods." *Stanford Encyclopedia of Philosophy,* Zalta (ed.), URL = https://plato.stanford.edu/archives/fall2019/entries/voting-methods/

Pettigrew, R. (2016): "Jamesian Epistemology Formalized: An Explication of 'The Will to Believe'." In: *Episteme* 13, pp. 253–68.

Pigozzi, G. (2016): "Belief Merging and Judgment Aggregation." In: *The Stanford Encyclopedia of Philosophy*, E. Zalta (ed.).

Rasmussen, M. Skipper, A. Steglich-Petersen and J.C. Bjerring (2018): "A Higher-Order Approach to Disagreement." In: *Episteme* 15, pp. 80–100.

Russell, J., J. Hawthorne, and L. Buchak (2015): "Groupthink." In: *Philosophical Studies* 172, pp. 1287–1309.

Skipper, M. (2019): "Higher-Order Defeat and the Impossibility of Self-Misleading Evidence." In: *Higher-Order Evidence: New Essays*, M. Skipper and A. Steglich-Petersen (eds.). Oxford: Oxford University Press.

——— (forthcoming): "Belief Gambles in Epistemic Decision Theory." In: *Philosophical Studies*.

——— (ms): "Does Rationality Demand Higher-Order Uncertainty?" Unpublished manuscript.

Skipper, M. and A. Steglich-Petersen (2019): "Group Disagreement: A Belief Aggregation Perspective." In: *Synthese* 196, pp. 4033–58.

Smithies, D. (2019): *The Epistemic Role of Consciousness*. Oxford University Press.

Steinberger, F. (2019): "Accuracy and Epistemic Conservatism." In: *Analysis* 79, pp. 658–69.

Tal, E. (forthcoming): "Is Higher-Order Evidence Evidence?" In: *Philosophical Studies*.

Titelbaum, M. (2015): "Rationality's Fixed Point (Or: In Defense of Right Reason)." In: *Oxford Studies in Epistemology* 5, pp. 253–94.

Tuomela, R. (1992): "Group Beliefs." In: *Synthese* 91, pp. 285–318.

Weatherson, B. (2019): *Normative Externalism*. Oxford University Press.

Worsnip, A. (ms): "Compromising with the Uncompromising: Political Disagreement Under Noncompliance." Unpublished manuscript.

Zollman, K. (2010): "The Epistemic Benefit of Transient Diversity." In: *Erkenntnis* 2, pp. 17-35.