Deliberation and Group Disagreement

Fernando Broncano-Berrocal & J. Adam Carter

Abstract: Suppose an inquiring group wants to let a certain view stand as the group's view. But there's a problem: the individuals in that group do not initially all agree with one another about what the correct view is. What should the group do, given that it wants to settle on a single answer, in the face of this kind of *intragroup disagreement*? Should the group members deliberate and exchange evidence and then take a vote? Or, given the well-known ways that evidence exchange can go wrong, e.g., by exacerbating pre-existing biases, compromising the independence of individual judgments, etc., should the group *simply* take a vote without deliberating at all? While this question has multiple dimensions to it—including ethical and political dimensions—we approach the question through an *epistemological* lens. In particular, we investigate to what extent it is epistemically advantageous and disadvantageous that groups whose members disagree over some issue use deliberation in comparison to voting as a way to reach collective agreements. Extant approaches in the literature to this 'deliberation versus voting' comparison typically assume there is some univocal answer as to which group strategy is best, epistemically. We think this assumption is mistaken. We approach the deliberation versus voting question from a pluralist perspective, in that we hold that a group's collective endeavor to solve an internal dispute can be aimed at different, albeit not necessarily incompatible, epistemic goals, namely the goals of truth, evidence, understanding, and epistemic justice. Different answers to our guiding question, we show, correspond with different epistemic goals. We conclude by exploring several ways to mitigate the potential epistemic disadvantages of solving intragroup disagreement by means of deliberation in relation to each epistemic goal.

1 Setting the stage: deliberative versus non-deliberative agreement following intragroup disagreement

Many disagreements take place in group settings. Over the years, religious groups (e.g., Christians) have internally disputed topics they consider significant (e.g., the real presence of Christ in the Eucharist). More often than not, political parties (e.g., the Tories) go through internal divisions over issues of societal importance (e.g., a no-deal Brexit). A brief look at the history of science reveals how scientists (e.g., physicists) disagree over factual issues in their fields (e.g., the Copenhagen vs. the many-worlds interpretation of quantum mechanics). More mundanely, disputes over practical matters are the order of the day in many families. On occasions, such internal disagreements end up badly, with a split in the relevant group or a punishment for the less influential. Sometimes, however, they result in a consensus[1] or an agreement of sorts to take a particular course of action or to let some view stand as the group's view.[2] It is this latter kind of intragroup disagreement we are interested in: the one that gets resolved.

How members of a group internally disagree matters for many reasons, not only for the stability or survival of the group but also *epistemically*. In general, there are two epistemically significant ways

in which intragroup disagreement can be resolved, i.e., in which members of a divided group can come to *agree to let a certain view stand as the group's view*: (i) they can *deliberate* and/or (ii) take a *vote*.

In this chapter, we are interested in investigating the epistemic significance that the key difference between deliberative and voting procedures has for the resolution of intragroup disagreement: namely, the fact that only deliberation necessarily requires that group members communicate with each other and, more specifically, the fact that, by doing so, they exchange their evidence. Thus, the paper aims to assess, in general, the epistemic significance that such an exchange (or lack thereof) has for the resolution of intragroup disagreement.

This is of course not to say that deliberation and voting are mutually exclusive mechanisms for groups to resolve their internal disputes. In practice, groups settle their disagreements by *mixed methods of decision-making*, i.e., methods that both involve deliberating and voting—as is, for instance, the mixed method for judging articles of impeachment in the United States House of Representatives.

That said, to better pin down the epistemic significance of each, it is best to keep them apart, at least theoretically. Thus, the kind of cases we will mainly focus on (whether real or ideal) have the following structures:

> **Deliberative cases of intragroup disagreement**: Some operative members[3] of group G hold p and some not-p at t_1; at t_2, G's operative members deliberate among themselves (i.e., they exchange reasons, evidence, arguments, and so on) with an eye towards settling whether p or else not-p should stand as G's view; at t_3, as a result of this process, they settle on either p or not-p.

> **Non-deliberative cases of intragroup disagreement**: Some operative members of G hold p and some not-p at t_1; at t_2, G's operative members aggregate their views by taking a vote given some voting rule (e.g., majority rule), absent any communication among each other, with an eye towards settling whether p or else not-p should stand as G's view; at t_3, as a result of this process, they settle on either p or not-p.[4]

Some clarifications are in order. First, at t_2, both in the case of deliberation and voting, members who initially believed that p should stand as G's view *may change their opinion*, and vice versa. Second, we leave unspecified the *number of group members* that respectively hold p and not-p to make it compatible with several possibilities—as we will see, this factor marks a distinction in terms of reliability between deliberation and voting. Third, lack of communication among group members in the case of voting is compatible with there being common knowledge (perhaps implicit) of the existence of an internal disagreement or of the fact that it is to be solved by taking a vote.

Finally, certain cases will not be our main focus. Quite often, members of a group settle on a collective view pursuing *non-epistemic goals*—regardless of whether this collective agreement is reached by deliberation or vote. For example, the board of directors of a pharmaceutical company might settle on the view that a newly marketed drug is not the cause of the death of many, even if they know it, to prevent huge financial losses. A government might systematically deny that the country's secret services have been used for morally contentious surveillance activities, even if known to be true, to prevent protests and media pressure. A religious organization might conceal criminal activities by its members—and thus uphold the collective view that such activities never happened—to avoid criminal charges and loss of reputation.

The reason we won't focus on such cases is that deliberative and voting procedures have little or no epistemic value when aimed at non-epistemic goals.[5] Instead, the kind of cases we are interested in are cases of intragroup disagreement in which members of a group reach a collective agreement pursuing *epistemic goals*.[6] Although this certainly reduces the scope of our inquiry, by idealizing our

focus in this way we will be in a better position to rule out pragmatic noise when answering the two key *epistemological* questions of the paper.7 For ease of reference, call these the *resolution question* and the *deliberation question.*

> **Resolution question:** What is the most epistemically appropriate way to resolve intragroup disagreement: by means of deliberation or by taking a vote? More specifically, to what extent is it epistemically advantageous and disadvantageous that group members *exchange evidence* when it comes to reaching a collective agreement?

> **Deliberation question:** Which conditions should *deliberative* disagreement comply with to be epistemically appropriate? More specifically, what it would take to overcome, or at least mitigate, the epistemic disadvantages of resolving intragroup disagreement by means of deliberation?

Our methodological approach to answering these questions is based on a simple *working assumption*: a group's collective endeavor to solve an internal dispute can be aimed at different (albeit not necessarily incompatible) epistemic goals. More carefully:

> **Assumption**: Possibly, for two epistemic goals, E and E^*, and for two groups, G and G^*, members of G would let p or else not-p stand as the G's view only if the collectively accepted view has epistemic property E and members of G^* would let p or else not-p stand as G^*'s view only if the collectively accepted view has epistemic property E^*.

With this assumption in place, each epistemic goal can be interpreted as providing a particular *standard* for assessing the epistemic significance of deliberating and voting in the resolution of intragroup disagreement. More specifically, the way we propose to assess this epistemic significance is in terms of *goal-conduciveness*: for each goal, we can assess to what extent the fact that group members exchange (or refrain from doing so) reasons and evidence are conducive to it.

The following are four salient candidate epistemic goals we will consider, albeit they are not exhaustive (see fn. 9). For any group G in which some operative members hold p and some hold not-p, in trying to settle whether p or else not-p should stand as G's view by means of method M, members of G would let p or else not-p stand as the group's view only if:

1. *Truth*: the collectively accepted view is true.
2. *Evidence*: the collectively accepted view is better supported by the best evidence individually possessed by group members than the opposite view.
3. *Understanding*: the collectively accepted view leads to more understanding than the opposite view.
4. *Epistemic justice*: the fact that G's members let such a view stand as G's view does not wrong any member specifically in her capacity as an epistemic subject (e.g., as a giver of knowledge, in her capacity for social understanding, and so on) or any other person outside the group in that capacity.8 9

Before assessing each epistemic goal, a final methodological caveat is in order. Our approach to the epistemic significance of deliberation and voting in terms of goal-conduciveness does not entail that the different epistemic goals are incompatible with each other, nor does it imply any stance on a number of debates, including (i) whether the relevant goals are finally or instrumentally valuable (or fundamentally or derivatively) in the case of collectives (cf., Goldman 1999; Fallis & Mathiesen 2013); or (ii) whether deliberation has a constitutive epistemic aim in terms of one of these goals; (iii) or

whether deliberation has procedural in addition to instrumental epistemic value (cf., Peter 2013). Our results might of course be relevant to these debates, but we stay neutral on them.[10]

Here is the plan. In §2, we address the truth goal, explain what the different kinds of evidence involved in deliberation are and how they bear on the individual reliability of deliberators; compare the collective reliabilities of deliberation and voting drawing on social choice theory, and show how complex it is to give a straightforward answer to the question of whether deliberation is reliable due to, among other things, the existence of several reliability-undermining group phenomena—which are widely investigated in social psychology. In §3, we explain why it shouldn't be assumed that deliberation always achieves optimal results, nor that voting always produces suboptimal outcomes *vis-à-vis* the goal of evidence. In §4, we offer two interpretations of the understanding goal and argue that on both interpretations deliberation outperforms mere voting. In §5, we argue that voting is more efficacious than deliberation with respect to the goal of epistemic justice. In §6, we propose several ways to mitigate the potential epistemic disadvantages of solving intragroup disagreement by means of deliberation in relation to each epistemic goal.

2 Assessing for truth

As we have seen, in settling for a collective view, groups may pursue non-epistemic goals (e.g., preventing financial losses), but they sometimes pursue epistemic goals. One example of an epistemically respectable goal, if not the most fundamental epistemic goal,[11] is *truth*.

In scientific disagreement, for example, members of a research group that internally disagree over some factual issue would not let a view stand as the group's view unless they considered it true, or at least, more likely to be true than any competitor view. In a quiz show, members of a divided team would not let an answer stand as the team's answer unless they considered it correct (or likely to be correct). Thus, for any given method that members of a divided group may use to reach a collective agreement, we can assess its epistemic propriety in terms of how conducive towards truth this method is. Crucially, the *reliability* of deliberative and voting methods depends on the kind of individual and collective conditions under which they are employed. Fortunately, these conditions have been widely investigated in disciplines such as social psychology and social choice theory. That being so, we will review some of their results (with an eye on truth as the relevant epistemic standard) so as to provide an answer to the resolution question on a safe theoretical and empirical ground.

Before that, it is worth pointing out an *epistemic difference* concerning the reliability of deliberative and voting procedures in general. This will allow us to subsume some relevant results of the aforementioned disciplines under a broader epistemological framework. Consider, first, the following general idea: the reliability of a group in letting only a view that is true stand as the collective view is to some extent premised upon the reliability of individual group members in choosing the right view *both* in the case of deliberation and voting.

To see this, consider a group (e.g. a flat Earth society) such that all members are utterly unreliable (e.g., almost always, they get things wrong) regarding the question of whether p (e.g. whether Earth is flat or spherical). Suppose that this kind of group internally disagrees on whether p or else not-p should stand as the group's view. Even if all members aim at settling on the true view, most will end up defending the false view because of their utter unreliability, whereas those who unlikely end up upholding the right view, will do it by luck. In such a situation, it doesn't matter whether the group deliberates or takes a vote: whatever the procedure for settling their internal disagreement, it will be an unreliable one. Therefore, group members need to be individually reliable to a minimum degree for them to reliably reach a correct collective agreement as a group—in §2.2, we will see what the minimum required degree of individual reliability is according to social choice theory.

Nonetheless, while the reliability of a group in letting only a view that is true stand as the collective view is to some extent premised upon the reliability of individual group members both in the

case of deliberation and voting, individual reliability is fixed differently in deliberative and non-deliberative cases. The reason, as we will argue next, has to do with the fact that the former involve *different kinds of evidence* besides private evidence and hence different individual competences are required to evaluate them (this will also have a bearing on our discussion in §3).

2.1 Individual reliability

In non-deliberative cases (at least as we have conceived them), the only evidence that group members use to establish which of the two options in a given dispute (*p* and not-*p*) is true (and therefore which one is the one that the group should uphold) is their own *private evidence*. In deliberative cases, by contrast, group members possess not only private evidence, but they are also exposed to *shared evidence*—i.e. evidence bearing on *p*/not-*p* shared by other group members during deliberation. In addition, as a consequence of this sharing process, they are also exposed to evidence about the distribution of opinions within the group, or *social evidence* for short—i.e. evidence that *n* number of group members are in agreement and *m* number in disagreement with one.

Thus, one plausible idea is that, for any given group member, her overall individual reliability concerning the disputed matter will be determined by how reliable she is in accurately judging to what extent *each* kind of evidence supports *p* or not-*p*. Interestingly, precisely because these are different kinds of evidence, the degree of reliability in assessing them need not coincide, hence the divergence with non-deliberative cases. Let's consider each in turn.

First, group members can be more or less reliable at seeing how relevant to the disputed matter their *private evidence* is, and on how much it supports or counts against the views in conflict. If, for instance, a group member's private evidence is misleading evidence for *p* because, say, it comes from a seemingly trustworthy but ultimately unreliable source, she will hardly assess in a reliable fashion that her evidence does not actually count in favor of *p*. Suppose that voting is the relevant procedure for resolving intragroup disagreement in a given case. In the difficult cases of misleading evidence just mentioned—as well as in cases where a group member has no evidence whatsoever—not voting for either option might be the best action to avoid collective error.

When it comes to *shared evidence*, matters are more complex. When group members put their private evidence on the table during deliberation, all involved members are exposed to two different things: (i) information that may back up, conflict, be redundant with or even irrelevant to their private evidence, and (ii) judgments from other group members to the effect that the shared information supports *p* or else not-*p* to such-and-such degree. Accordingly, and as in the case of private evidence, a group member can be more or less reliable at assessing to what extent the information provided by other members is relevant and corroborative of *p* or of not-*p*, and this can be done by, among other things, correctly assessing to what extent those other members are *competent* information-gatherers. Interestingly, as some have noted (e.g. Elga 2010; Weatherson 2013; Asunta-Eder [this volume]), being competent at acquiring evidence is independent from being competent at correctly judging the confirmational import of the evidence. Thus, on top of being more or less reliable at assessing the evidence shared during deliberation, group members can be more or less reliable at assessing to what extent their fellow members' assessments of such shared evidence are accurate.

The last kind of evidence involved in deliberative cases, *social evidence*, is somehow different, as it does not directly bear on the question of whether *p*. As we have defined it, social evidence is evidence about, specifically, the distribution of opinions within the group, i.e., evidence about how many group members believe that *p* (or else not-*p*) should stand as the group's view because *p* (or else not-*p*) is true. Interestingly, social evidence can have a defeating effect on its own even if it does not directly bear on the question of whether *p*, and namely even if it carries no more information than that of assertions of the type "I think that *p* should stand as the group's view".

To see this, suppose that you have conclusive private evidence for the truth of *p* and that on that basis you believe that *p* should stand as the group's view. Furthermore, suppose that you are the only person in your group in possession of evidence that is relevant to the disputed matter. You share your evidence with your fellow members, which you regard as your epistemic peers. Suppose, next, that no one is moved and all of them (e.g., 999 members), except for you, individually assert "I believe that not-*p* should stand as the group's view because not-*p* is true". Many in the epistemological literature on disagreement agree that you should reduce your confidence in your belief simply because of being exposed to social evidence to the effect that a majority is in disagreement with you.12 Furthermore, this defeating effect may occur even if private evidence directly bearing on *p*/not-*p* has not been put on the table yet.13

Turning to reliability, the kind of competence required to judge whether the social evidence available in the group is misleading or on the right track is a competence to judge whether the other group members are being *sincere* in asserting things such as "I believe that *p* should stand as the group's view". Suppose that, during deliberation, someone in your group asserts that. The questions you should ask yourself, *qua* group member, are like these: Is this person being sincere? Does she really care about the truth? Or is she making that assertion for strategic or pragmatic reasons? If one can answer these kinds of questions correctly for all (or at least many) group members, one is reliable at processing the group's available social evidence.14 By contrast, if one conciliates with the majority for a non-epistemic reason such as social comparison (e.g., to maintain a socially favorable position within the group), one is not reliable at processing the group's social evidence.15

Which of these different bodies of evidence (i.e., private, shared or social) and which of the corresponding degrees of reliability in processing them should have a greater weight in the overall individual reliability of a given group member is a question whose answer hangs to a great extent on the correctness of the different views of the epistemology of disagreement. In general, *steadfast theorists* will be more inclined to assign a greater weight to the group members' assessments of their own private evidence (or even of the shared evidence), whereas *conciliationists* will lean in the direction of giving a greater significance to the judgments of other group members and to the distribution of opinions within the group.16

To summarize the discussion so far, we've seen that there is a significant difference between how individual reliability is fixed in deliberative and non-deliberative cases of intragroup disagreement. This difference has to do with the fact that in non-deliberative cases group members only need to evaluate the confirmational import of their own private evidence to choose between the true and the false collective view. In deliberative cases, by contrast, group members need to process, besides their private evidence, the evidence shared by others as well as the available evidence about the distribution of opinions within the group, which can have a defeating effect by itself. In such deliberative cases, however, it is an open question which kind of evidence should carry more weight in fixing the individual reliability of group members, or what the interplay between these three types of evidence might be.

2.2 Group reliability
One question we can ask concerning the reliability of deliberative and voting procedures is how reliable individual members of a group undergoing an internal dispute need be in order for one such procedure to reliably lead the group to settle on the true view. The results of social choice theory become useful on this score.

Let's consider voting first. In general, political scientists assess voting rules in terms of *fairness criteria*, i.e., how sensitive they are to all of the voter's opinions and preferences in the right way (Pacuit 2019). However, interestingly for our epistemological purposes, they can also be assessed in terms of how well they track the truth, i.e., in terms of how much the resulting collective view approximates it

(List 2013; Pacuit 2019). A voting rule that is often referred to as a *collective truth-tracking device* in the two-option case (the one we are concerned with) is *majority rule* (see, e.g., List & Goodin 2001).

As it is well-known, one prominent argument for adopting majority rule comes from the Condorcet Jury Theorem (CJT), which maintains that, given two possible positions *p* and not-*p* with respect to a given topic (e.g., a verdict, a diagnosis, a factual issue), where only one of the options is correct given some standard (in our case, truth), the probability that a majority votes for the correct option increases and converges to one as the size of the group grows. Crucially, CJT is premised upon two conditions: (i) that the probability (viz. reliability) that each group member identifies the correct position is greater than 0.5 and the same for all voters (*voter competence condition*); and (ii) that all correct votes are mutually independent conditional on the truth,[17] which is either *p* or not-*p* (*voter independence condition*).[18]

Whether or not majority rule reliably yields epistemically appropriate results (true or accurate group agreements) crucially depends on the voter competence and independence conditions being met. But this seldom happens. For instance, factors that have been cited as leading to correlated votes include opinion leaders, schools of thought, communication among voters or common information (cf., e.g., Ladha 1992). Moreover, as Dietrich and Spiekermann (2020) point out, *any* common cause of votes is a potential source of dependence, including non-evidential (e.g., situational) factors such as distracting heat.[19]

Lack of independence has implications not only for whether or not CJT applies to a group that aims to resolve an internal dispute by taking a vote according to majority rule, but also for how the nature of such a disagreement should be conceived. After all, if the votes of every member in the two disagreeing subgroups are correlated, intragroup disagreement comes down to a one-to-one disagreement situation, as there would be two sets of mutually dependent votes: those for *p* and those for not-*p*. This is epistemically significant. For one epistemic benefit of CJT is that the *larger* the group, the *better* at tracking the truth it is. Therefore, if all votes are correlated in the two disagreeing subgroups, the size of the group no longer has a bearing on its reliability.

Turning to voter competence, multiple specific factors can bear on the individual reliability of voters. From a general epistemological point of view, the quality of their private evidence is probably the most significant factor. But note that voter reliability can be low even when the evidence privately possessed is good evidence. For the probability that a voter *correctly* judges that her good evidence is supportive of *p* rather than not-*p* is independent of the epistemic goodness of the evidence (e.g., someone with conclusive private evidence might fail to notice that the evidence is conclusive because of not being sufficiently attentive). Conversely, a voter with misleading evidence might uncritically follow her evidence, thus making it unlikely that she votes for the correct view. Finally, voters can also be unlikely to vote for the right view when they possess no evidence whatsoever (e.g., by casting their votes for *p* or not-*p* merely on the basis of the results of tossing coins that unbeknownst to them are independently biased in favor of the false view). Interestingly, all this can happen while all group members vote with the aim of choosing the correct view.

Lack of voter competence bears on the epistemic appropriateness of majority vote as a way of solving a group's internal dispute. One thing that the literature on CJT shows is that when the votes are independent but the competence of all voters is lower than 0.5 and the same for all, or when the average judgmental competence of voters in the group is lower than 0.5 (cf., e.g., Grofmam et al. 1983), the probability that a majority votes for the correct option increases and converges to 0 as the size of the group grows. So majority rule can be an epistemically *inappropriate* procedure for solving intragroup disagreement after all.

Of course, the literature is filled with jury theorems that relax the independence (Ladha 1992, Dietrich and List 2004, and Dietrich and Spiekermann 2013; Goodin and Spiekermann 2018) and the competence (e.g. Boland, P.J. 1989; Grofmam et al. 1983) conditions while still serving as truth-

tracking devices, and hence as epistemically appropriate ways to resolve intragroup disagreement, at least in the case of large groups.[20] But, in general, voters need to be individually reliable to a sufficient degree, where in most cases this means being better than random.[21]

For some internal disputes, when individual reliability is an issue, groups can opt for some sort of proxy voting system that allows delegation of the votes to the most competent or well-informed in the group, or to weighted majority rules (e.g., expert rules) that assign different weights to different competence distributions (e.g., more weight to the votes of the most competent members). In general, for any competence distribution, there will be an optimal voting truth-tracking procedure for the group to solve its internal disagreement (for optimal voting rules see, e.g., Nitzan & Paroush 1982; Gradstein & Nitzan 1986; Dietrich 2006).

Beyond specific voting rules groups might use to resolve their internal disputes, one question we can ask is this: does communication and evidence sharing among group members represent a significant epistemic advantage over members simply taking a vote on the basis of her private evidence? What are the epistemic benefits and drawbacks of deliberation in general *vis-à-vis* the goal of reaching a true collective agreement?

One way to answer these questions is to offer a formal analysis of deliberation and compare it to voting procedures. Hartmann and Rafiee Rad (2018) do precisely this and show that deliberation is truth-conducive in a similar way as majority voting as per CJT. It is worth considering their proposed Bayesian model of deliberation, not only because its results are relevant to the subject matter, but also because it will serve to illustrate the many complexities that communication among group members may give rise to and, therefore, that any formal model of deliberation might need to incorporate if deliberation is to be compared to voting in a realistic way.

Hartmann and Rafiee Rad's Bayesian model of deliberation is based on several assumptions. First, all evidence is put on the table *before* deliberation (i.e., no extra evidence shows up during deliberation, so all evidence is shared evidence). Second, group members are assigned a first-order reliability value that measures how correctly they judge the disputed matter. Third, they are assigned a second-order reliability value that reflects how well they estimate the first-order reliability of the other group members. The latter is kept fixed during the course of deliberation, while the former may increase as members learn to better judge the reliability of other group members. In this way, deliberation, as they model it, consists in the following process:

> The group has to decide on the truth or falsity of a hypothesis H. Each group member assigns a certain probability to H. Then each group member casts a vote on the basis of this probability. Then each group member updates her probability on the basis of the votes of the other group members, weighted according to the estimated reliabilities (...). The procedure is iterated, and in each round the second order reliabilities are increased which leads to a more accurate estimation of the reliability of the votes of the other group members. After a number of rounds, this process converges (Hartmann & Rafiee Rad 2018: 1278)

Their results show that the truth-tracking properties of deliberation are very similar to those of majority vote. As they summarize them:

> The deliberation process results in a consensus and correctly tracks the truth for groups of large size in the following cases: (i) homogeneous groups with a first order reliability greater than 0.5 and with a high second order reliability (ii) inhomogeneous groups with average first order reliabilities above 0.5 and with a high (initial) second order reliability. In this sense the deliberation procedure manifests the same epistemic properties as the majority voting while adding the benefit of a group consensus (...) We furthermore provided simulation results that indicate that the deliberation procedure tracks the truth even in cases that do not fall under the

conditions stated in the Condorcet Jury Theorem for majority voting as well as for groups with low second order reliabilities (Hartmann & Rafiee Rad 2018: 1289)

In sum, if Hartmann and Rafiee Rad are right, although majority vote may be more easily implemented as a procedure for solving intragroup disagreement in the case of large groups, deliberating *vis-à-vis* the goal of reaching true collective agreements is roughly as epistemically appropriate as voting by majority rule. This gives an answer to the resolution question. As we will see next, however, this answer is incomplete, since real-life deliberation cases may involve many complexities that make giving a general, straightforward answer to that question a complex matter.

2.3 Deliberation in non-idealized conditions

Formal models are surely a great approach to the question of whether deliberation or voting is the most reliable way to solve intragroup disagreement. But deliberation involves many complexities—not present in voting cases—that have a bearing on its reliability as a collective method for solving intragroup disagreement, and which can make it difficult to give a straightforward answer. To illustrate this, consider Hartmann and Rafiee Rad's model again. As they acknowledge, in order to capture the reliability of more realistic deliberative situations, several assumptions of the model need to be relaxed, such as the assumption that the deliberators are independent, i.e., that the only cause for a group members' verdict is the truth or falsity of the hypothesis in question—whereas other members' verdicts are evidence for the truth or falsity of that hypothesis that don't necessarily break such an independence. Indeed:

(i) In real-life deliberation cases, the individual judgments of members of deliberating groups may not be independent from each other.

Or the assumption that the first-order reliability of group members remains unchanged during deliberation. After all:

(ii) In real-life deliberation cases, the probability that a given group member is right or wrong about the disputed matter may change along the deliberative process.

Or the assumption that the first-order reliability of group members is independent from their second-order reliability. However:

(iii) In real-life deliberation cases, how well a group member estimates how reliable, concerning the disputed matter, other group members are may be influenced by the judgment of those other members.

Other complications that a formal model of deliberation might need to incorporate to better reflect how deliberation compares to voting in real cases include the following. For instance, one crucial assumption of Hartmann and Rafiee Rad's model is that there is full disclosure of the evidence among all group members before deliberation. But:

(iv) In real-life deliberation cases, group members may gradually disclose their private evidence to other group members, and may not even disclose any evidence at all.

This is relevant, because in so proceeding deliberating groups can and often times fall prey to *shared information bias*, a tendency to discuss shared evidence, i.e., evidence that most group members possess, in detriment to discussing potentially relevant evidence privately possessed by individual members or only shared by a few of them (cf., Stasser & Titus 1985). If group members have good private evidence but they do not draw on it during discussion—a situation that is often referred to as a

hidden profile—, the reliability of the deliberative process can be compromised. Several factors can help groups overcome shared information bias (see §3). Interestingly, the kind of groups we are concerned with—those whose goal is to find a correct answer—see this bias diminished by devoting more of their discussions to considering critical clues thus becoming more likely to adopt a correct view when relevant private information remains unshared (Stasser & Stewart 1992).

Another crucial assumption of Hartmann and Rafiee Rad's model is that deliberation proceeds in a series of iterations in which each group member first assigns a probability to the hypothesis in question, then casts a vote, and then each member updates her probability on the basis of the votes of other group members, weighted according to the estimated reliabilities. However:

> (v) In real-life deliberation cases, group members may discuss the relevant issue one or several times and then take one final single vote to decide which view should stand as the group's view, or simply reach consensus without voting at all.

Other complexities have to do with the different types of evidence distinguished in §2.1. For example, during deliberation group members put their private evidence on the table, which becomes shared evidence. The complication, as we have already pointed out, is that:

> (vi) In real-life deliberation cases, group members may need to assess two things: how good the evidence shared by other members is (e.g., by judging, among other things, how reliable those members are in gathering good evidence) and how good those other members' assessments of their own shared evidence are (e.g., by judging how reliable they are in assessing the confirmation import of their evidence).

This means that a more realistic model may need to include two measures of second-order reliability, instead of one (see Asunta-Eder, this volume, for this kind of approach).

In addition, as we have also argued, social evidence—i.e., evidence about the distribution of opinions within the group—can have a defeating effect on its own (i.e., independent of the group's shared evidence) even if it carries no information directly bearing on the question of whether *p* (like shared evidence does). Relatedly, the very distribution of the disagreement matters and, in particular, when the relevant intragroup disagreement is between a majority and a minority. This is illustrated by extensive research in social psychology on group conformity pressures, and in particular, on *majority influence*. For example, in a famous study by Sherif (1936), subjects were asked to perform a visual task. Subjects whose estimations diverged from those of the majority gradually converged to the latter after being exposed several times to the opinions of the majority. In later studies by Asch (e.g., 1952), the relevant visual task had an obvious correct answer and conformity to the majority was also observed (although to a lesser extent). Accordingly, it is plausible that:

> (vii) In real-life deliberation cases, group members who hold a different view to the one held by most members of the group may conform to the majority opinion by repeatedly being exposed to it.

Judging whether a disagreeing majority is right or wrong might be a complex issue. In particular, to judge whether majority influence is epistemically appropriate, one needs to determine whether it is informational (i.e., due to the fact that there is more evidence supporting the relevant opinion) or normative (e.g., due to a desire to fit and avoid social exclusion). Incorporating a corresponding realistic measure of reliability in a formal model of deliberation might accordingly be a complex issue as well.

The issue is even more complex considering the fact that *minorities* also exert influence on majorities. For instance, Moscovici, Lage, & Naffrechoux (1969) observed this kind of effect in a visual

task with an obvious correct answer when a minority of subjects gave consistent and unanimous answers that diverged from those of the majority. Thus:

> (viii) In real-life deliberation cases, group members who unanimously hold a different view to the one held by most members may make the latter conform to their opinions by consistently exposing it to them.

The existence of minority dissent is not necessarily negative at the collective level. Quite the contrary: minority dissent previous to group discussion has been observed to improve the quality of the resulting collective judgments and decisions (e.g., Hightower & Sayeed 1996; Brodbeck et al. 2002; Schulz-Hardt et al. 2006)

Another specific condition widely investigated in social psychology that may affect the reliability of deliberation is the *group polarization* phenomenon (e.g., Stoner 1961; Burnstein & Vinokur 1977; Isenberg 1986):22

> (ix) In real-life deliberation cases, the individual members of like-minded groups may adopt, on average, more extreme views after group discussion than those held before deliberation.

This means, for instance, that in an intragroup disagreement where most group members lean, on average, towards p and only a few towards not-p chances are that if group members discuss whether p or not-p should stand as the group's view, the group's average will more strongly lean towards p. This is a source of collective unreliability, at least in the cases where p is false. Thus, the initial distribution of opinions in a group featuring an internal disagreement matters for how reliable deliberation is in solving it.

Finally, when it comes to the different kinds of evidence involved in deliberation, the most difficult issue to solve is this:

> (x) In real-life deliberation cases, it may be unclear what exactly the interplay between the different kinds of evidence (private, shared or social) is, and which one should play a more significant role in whether a group ends up adopting a true or else a false view following deliberation.

As we noted in §2.1, it is an open question which of these three kinds of evidence should have a greater weight in fixing the reliability of individual deliberators. This question might be difficult to address with formalization or empirical research only, and further philosophical investigation is required.

So where does this leave us? Is deliberation as appropriate as a method to solve intragroup disagreement *vis-à-vis* the truth goal as voting is when CJT-style theorems apply? According to Hartmann and Rafiee Rad's model, the answer is 'yes'. However, this answer to the resolution question, albeit on the right track, is not (as they also acknowledge) fully satisfactory: deliberating groups can be affected by a variety of factors that bear negatively (but also positively) on the reliability of deliberation. Some such factors that we have discussed are: (i) the interdependence between the judgments of group members; (ii) changes in their first-order reliability along the deliberative process; (iii) the first-order and second-order reliabilities of group members not being independent; (iv) shared information bias and hidden profile situations; (v) different modes of deliberating, such as several iterations of deliberation and voting, deliberation followed by a single final vote, or deliberation followed by consensus absent voting; (vi) the group members' need to assess the epistemic quality of the evidence shared by others and of the judgments they make about such evidence; (vii) majority influence; (viii) minority influence; (ix) group polarization; and (x) the complex interplay between private, shared and

social evidence. If anything involves complexity, that is the question of whether deliberation is a truth-conducive method for solving intragroup disagreement.

3 Assessing for evidence

Truth is not the end of the story, however. It is not unusual that members of a group featuring an internal disagreement are not merely interested in settling on a true collective view, but on a view that is supported by the best evidence individually possessed by them. Of course, the truth and evidence goals are not incompatible and are in fact oftentimes pursued simultaneously, so that members of a group would let a certain view stand as the group's view only if it were true and supported by the best private evidence available in the group.

However, the two goals are also independent from each other and cases of groups whose primary goal is not truth but evidential support are conceivable. For example, consider a group of high-profile members of the Bush administration back in 2003 having a disagreement about the exact location of Saddam Hussein's weapons of mass destruction. Even if all are aware or suspect that there are no WMDs, they might still be interested in collectively agreeing on the view that is supported by the best evidence privately possessed by them (such as the most credible military reports about the possible locations of WMDs). Or consider a tobacco company's board of directors back in the '50s having a disagreement on which kind of evidence provides the best epistemic justification against the (now proven) fact that there is a causal link between smoking and lung cancer. Each board member might be in possession of different bodies of evidence, such as different statements from physicians against that fact or different scientific reports to the effect that there is no conclusive scientific proof of a link between smoking and cancer. Even if all board members might individually suspect that smoking causes cancer, they might still be more interested in agreeing on a collective view that is false yet supported by their best private evidence than on a true collective view with worse or no intragroup evidential support.[23] The motives of these groups might be non-epistemic—e.g., convincing the public opinion that Saddam Hussein has WMDs or that smoking is not causally linked to lung cancer—, but their goals, insofar as they prime evidential support, can be considered epistemic.[24]

Thus, with the goal of evidence (not truth) in mind, we can ask: what method for solving intragroup disagreement is best *vis-à-vis* the evidence goal: deliberation or voting? At first sight, deliberation seems a better method for solving intragroup disagreement when a group is mainly seeking evidential support. After all, it is at the core of any deliberative process that group members communicate their opinions and share their evidence with other members. Thus, in ideal deliberative conditions, no collective decision is made or no collective view is adopted unless all members share their private evidence and everyone processes it.

By contrast, not all voting rules seem well-suited to reach collective views that are supported by the best evidence privately possessed within the group. Consider majoritarian rules. Suppose that a group of physicians disagree about whether they should give treatment A or B to a patient. All except one believe that they should apply A. Their opinions are based on their own physical examination of the patient and their clinical judgment. By contrast, the only dissenter in the group is in possession of conclusive evidence for applying B (e.g., evidence from randomized clinical trials). Without prior deliberation (something admittedly rare for medical decisions), the group takes a vote and they collectively accept that they should apply treatment A. This view, however, is not supported by the best evidence privately possessed within the group: the dissenter's evidence is discounted as a result of the voting procedure.

Of course, this neither implies that deliberation always achieves optimal results, nor that voting always produces suboptimal outcomes *vis-à-vis* the goal of evidence. As the social choice theory literature shows, some voting rules are conducive to this goal. In addition, as the social psychology

literature demonstrates, deliberating groups often operate in non-ideal conditions that prevent them from exploiting the full potential of deliberation.

Let's consider voting first. In a recent paper, Bozbay et al. (2014) propose a quota rule that aims to make correct collective decisions (or judgments) while being efficient in light of all the information privately possessed by group members. As Dietrich and List (2007: 392) explain, quota rules are judgment aggregation rules such that "a proposition is collectively accepted if and only if the number of individuals accepting it is greater than or equal to some threshold". Bozbay et al. focus on cases in which groups need to settle on the correctness of two propositions, e.g., a jury on whether a contract was broken and whether it is legally valid, or a hiring committee on whether a candidate is good at research and good at teaching. Their proposed quota rule for these simple preferences (i.e., choosing between a correct and an incorrect decision, making a right or a wrong judgment) is based on the idea that for any of these propositions to be more probably true than false given all information, at least a number of group members above a certain threshold—which they define formally—need to possess evidence for that proposition. This rule, they argue, effectively uses all the private evidence available in the group assuming that group members aim for correct decisions and judgments. In sum, even if majoritarian rules are not appropriate for the goal of evidence, voting cannot be discarded out of hand as a procedure for solving intragroup disagreement with the aim of settling on an evidentially well-supported collective view.

Let's consider deliberation now. While it is true that there are voting rules that are appropriate for the goal of evidence, it is also true that deliberation may not be conducive to it. As we have seen, deliberating groups often fall prey to shared information bias—recall: the tendency to discuss evidence that most members possess in detriment to discussing potentially relevant evidence privately possessed by individual group members or only shared by a few of them. When a group undergoes this kind of bias and the unshared evidence is the best or at least relevant evidence (a hidden profile situation), solving an internal dispute by deliberating may not be more conducive to the evidence goal than voting by majoritarian rules.

We have also seen that groups that aim to find a correct answer (at the truth goal), see this bias diminished because they devote more discussion time to consider critical clues (Stasser & Stewart 1992). Other factors that help reduce shared information bias include the involvement of team leaders (Larson et al. 1996) as well as of members with experience in the subject matter (Wittenbaum 1998)—they pay more attention to unshared information—, low time pressure or having access to sheets that either indicate which pieces of information are shared and unshared—rather than discussing the relevant issue from memory—(Bowman & Wittenbaum 2012). However, the most relevant factor in the context of intragroup disagreement is, precisely, the existence of dissent within the group.

 For instance, in a study by Brodbeck et al. (2002) they observed that groups featuring dissent before deliberation shared more information during discussion. In another study, Schulz-Hardt et al. (2005) found that groups featuring dissent are more likely to arrive at a correct collective decision or judgment than homogenous groups—especially when someone in the group advocates the right solution—by, among other things, increasing discussion intensity and better pooling of the unshared evidence. Interestingly, they also established a correlation between pre-deliberation dissent and better collective outcomes and less shared information bias in groups where none of the members favored the right solution—e.g., groups like the high-profile members of the Bush administration disagreeing about the possible locations of Saddam Hussein's WMDs or the tobacco company's board of directors disagreeing about which evidence best proves that there is no causal link between smoking and lung cancer. Another study by Greitemeyer et al. (2006) further confirmed the positive effects of intragroup disagreement, this time with artificially fostered controversy within target groups. In their study, they implemented an advocacy procedure in which each group member acted as an advocate for each

alternative for some time independently of their individual preferences. This procedure resulted in an increased exchange of both unshared and shared information.

Thus, if it is at the core of any deliberative process that private evidence is shared among group members and the very existence of intragroup disagreement is already a factor that cancels out (perhaps in combination with other factors we've seen) shared information bias, we have good reason to think that deliberating is an appropriate way to solve intragroup disagreement with the primary aim of settling on an evidentially supported collective view.

4 Assessing for understanding

Suppose that members of a group, *G,* aim to let *p* or else not-*p* stand as the group's view only if it facilitates *understanding.* Understanding—at least, as it is typically discussed in epistemology—is a genus with (at least) two distinct species: (i) *understanding-why* (e.g., I understand why the house burnt down, I understand why Caesar crossed the Rubicon); and (ii) *objectual understanding* (I understand chemistry, I understand Australian Rules football). It is an open question how these two species are related.

Moreover, it is a point of contention whether either of these species of understanding *reduces* to propositional knowledge, or to each other.[25] For the present purposes, we will remain neutral on these points. One assumption we will make, however, is that understanding involves—in some suitably specified sense—*grasping.* In the case of *understanding-why*, what one grasps when one understands why something is so is the relation between the *explanans* and the *explanandum.* In the case of *objectual understanding*, what one grasps when one understands something, X, which can be treated as a subject matter, is the explanatory and coherence-making relations between propositions making up the relevant body of information (e.g., Kvanvig 2003, Ch. 8; Gordon 2017).

In *both* cases, truth plays a constraining role, even if true beliefs don't suffice for understanding of either variety. For example, you don't understand why the house burnt down if the explanation you grasp is itself false—e.g., *The house burnt down because of arson* (false) versus *The house burnt down because of faulty wiring* (true). Likewise, in the case of objectual understanding, one doesn't understand combustion even if one grasps the coherence and explanatory relations between the (mostly false) propositions making up phlogiston theory—and this is so even if one, by such grasping this body of information, understands, merely, the phlogiston *theory* of combustion.

If true beliefs are necessary for understanding, why are they not sufficient? Here is the importance of grasping to understanding. A child might, for example, believe truly (e.g., via testimony from a parent) that the house burnt down due to faulty wiring without *understanding why* this is so, on account of failing to suitably grasp how the *explanans* relates to the *explanandum.* Likewise, one might fail to understand algebraic geometry even if one has memorized true axioms and formulae, if one fails to grasp how the relevant axioms and formulae hang together, e.g., by failing to grasp how the axioms and formulae are mutually supporting.

Against this (albeit brief) background: let's consider how voting and deliberation, respectively, fare as a means to achieving a group's aim to let *p* or else not-*p* stand as the group's view only if it facilitates *understanding.* For simplicity, we will focus on understanding *why*—using as a reference case the following: suppose the mayor of a city has appointed a committee to determine why city hall mysteriously burned to the ground in a fire. The two salient alternatives the committee are evaluating as the cause are faulty wiring and arson. Let 'Case 1' be a case where the group simply takes a vote (e.g., 'yea' for fire, 'nay' for arson), and let 'Case 2' be a case featuring deliberation. Which best facilitates understanding?

Interestingly, there are two very different senses in which a vote or a deliberation *vis-à-vis* arson or faulty wiring might (broadly speaking) 'facilitate' understanding, which need to be separated. Suppose, for example, the mayor goes further to instruct the committee to not settle on a group view

(*vis-à-vis* arson or faulty wiring) *until* the group understands why the hall burnt down. In such a situation, the settled view should be made *on the basis* of a certain kind of epistemic credential, and regardless of whether taking that vote *promotes* the group's understanding—e.g., regardless of whether reaching a settled view will *itself* help the group come to understand the cause of the fire, or to increase that understanding.

On this interpretation of the understanding goal, *merely* voting will be inefficacious, and this is so even if voters antecedently meet voter competence and voter independence conditions, and thus, even if voting would further the truth aim. Deliberation by comparison does much better. Put another way, the probability that the group will be positioned to reach a verdict on the basis of understanding conditioned on deliberation is higher than on voting, even if it is *not* particularly high in cases where subjects prior to deliberation fail to meet competence and independence conditions.

The rationale for why voting will do worse in comparison to deliberation (*vis-à-vis* the above interpretation of the understanding goal) is that the mere registering (by group members) *that* each member holds certain views [either the 'yea' or 'nay' view] is simply not the sort of thing that could ensure that a group *grasps* why something is so; deliberation, by contrast is. On this point, Kenneth Boyd's (2019) analogy between a group's physically grasping something and cognitively grasping something is helpful:

> [...] consider again the way in which individuals rely on one another in the case of physical group grasping: they rely on each other insofar as they are aware both that they must direct effort towards the same goal, and that if any other let go then they would not be able to pull their friend to shore safely on their own. In the epistemic case the situation is analogous: if the goal of the group is to understand (why/how/that) p, then members of the group are mutually p-reliant in the case that they recognize both that they are contributing towards the relevant goal (perhaps in the form of representing reasons and relationships between reasons), and that they would not be able to achieve that goal on their own (given the circumstances) (2019: 15-6).

Boyd's idea here is that kind of grasping that is germane to a group's understanding something involves, necessarily, some kind of *group reliance*, viz., reliance between group members on each other's contributions toward the common goal, as well appreciation that each other's contributions is necessary. Mere voting in the absence of sharing evidence is a paradigmatic example of a *non-reliant* contribution to a common goal. If group members fail to understand why city hall burnt down prior to voting, so they will fail to understand why after voting, and will thereby simply *register* viewpoints that are reached unreliantly on other group members' influences. Deliberation, by contrast, offers at least the kinds of conditions that could make such reliance possible, especially when deliberation involves the sharing of evidence and reasons.

The above articulates the situation, at least, if the idea is that the settled view of the group should be made *on the basis* of a certain kind of epistemic credential—viz., understanding. Interestingly, with respect to the understanding goal, we end up with the same result (viz., deliberation beats voting) even if we reject that the settled view should be made *on the basis* of a certain kind of understanding and instead ask whether voting or deliberation better *promotes* the group's understanding.

Voting does look, initially at least, as though it could promote understanding. Returning to our illustrative case of the committee appointed by the mayor to determine why city hall burned (viz., arson or fire): even if the committee simply takes a vote, with no deliberation whatsoever, group members, in virtue of their appreciation of what each other has voted, might gain some kind of intellectual traction on the situation. For example, if I am on the committee and antecedently think that other committee members satisfy a voter competence condition, then my coming to find out that there is a near-unanimous vote favoring the arson explanation might lead me to think it's more likely than not that it

was arson rather than faulty wiring. In fact, on the basis of simply gaining knowledge of this revealed voting distribution, I might even become highly epistemically justified in believing this.

However, no matter how high we raise my antecedent knowledge of the extent to which other voters on the committee satisfy a competence condition, and no matter how large the group size agreement is (e.g., no matter how many people there are on the committee whose votes align with the same explanation), it remains that the kind of intellectual improvement I might attain by simply learning what a voting pattern is, is going to fall short of understanding. The same, however, does not apply in the case of deliberation.

The argument for this is as follows. The first premise says that expert testimony does not suffice for individual-level understanding-why. This premise (defended in various places by Duncan Pritchard, e.g., 2009; 2014) gains support from the following kinds of cases. Suppose you want to understand why the dinosaurs went extinct, and you ask an expert paleontologist. The paleontologist is in a hurry and simply tells you that they went extinct because of an asteroid. You then come (in the absence of any undefeated defeaters for this expert testimony) to believe the proposition "the dinosaurs went extinct because of an asteroid". While it is uncontentious (regardless of whether one is a reductionist or anti-reductionist in the epistemology of testimony) that you can come to gain *propositional knowledge* on the basis of this kind of testimonial exchange, Pritchard's line is that such testimony isn't enough to secure understand-why given that (i) understanding-why requires a suitable grasp of the how the relevant cause and effect relationship, and (ii) such a grasp is not something one gains *simply* by accepting someone's word, even an expert's.

The second premise of the argument then draws an analogy between testimony from experts and testimony from intragroup members, where the latter is effectively what one gleans by coming to learn that there was a majority voting pattern in favor of one explanation (e.g., arson) over another. The claim is that *if* the former doesn't suffice for understanding, then neither does the latter. That is: if expert testimony to the effect that some causal claim is true doesn't suffice for understanding *why* that claim is true (even if it suffices to furnish justification or even knowledge), then neither will the testimony to the effect that some causal claim is true when the source of that testimony is an aggregation of voting choices by individuals one regards to be competent. Taken together, the two premises imply that mere voting is not going to facilitate group understanding-why—or, at least, not any more than *mere* expert testimony facilitates understanding why.

It's worth noting of course that deliberation is importantly different from mere voting in exactly the kind of respect in which mere voting was shown to (like relying on testimony) be incapable of engendering understanding-why. This is because when a group deliberates about why something, X, is so, the sharing of evidence (and indeed, in some cases, the critical discussion of shared evidence) engages with not *merely* the matter of *what* caused X (e.g., arson caused the fire, an asteroid caused the dinosaurs' extinction, etc.), but *how* it did so. For example, the mayor-appointed committee, upon sharing evidence, will discuss such things as whether the building had enough flammable material to have burnt simply through a burning wire, what an arsonist would have had to do to have brought about the fire in the way it was brought about, etc. Such considerations are, of course, exactly the thing that (*à la* Pritchard) one would have to have some command of if one is to grasp the connection between the relevant cause and effect. And, moreover, by *relying* on one another for such considerations (and not merely for the *verdicts*), group members are grasping an explanation as a group in a way that is (*à la* Boyd) analogous to the way a group might physically grasp something together.

Bringing this all together: we've seen in this section that there are two ways we might plausibly measure the effectiveness of voting as opposed to deliberation in light of the epistemic aim of understanding. The first way is to ask which is more effective if the objective is for the group to reach a settled view about why X only if the group understands why X. The second is to ask which is more

effective if the objective is to promote the group's understanding. In both cases, we've seen (for different reasons) that deliberation outperforms mere voting.

5 Assessing for epistemic justice

Thus far, we've been considering how voting versus deliberation fare with respect to the following kinds of goal-conduciveness: truth, evidence, and understanding. Each of these goals is a traditional epistemic goal.

As recent work in social epistemology has shown, there are important connections between epistemic goals and social power and pressures, connections which can give rise to what Miranda Fricker (2007) terms *epistemic injustice*. Put generally, an injustice to someone is an epistemic injustice if it involves their being wrongfully disadvantaged in their capacity as an epistemic subject (e.g., a potential knower). A central species of epistemic injustice is *testimonial injustice*, for example, when prejudice leads a hearer to give a deflated level of credibility to a speaker's word.26 In a group setting, we might imagine, for instance, a female or minority juror's viewpoint being disregarded on the basis of sexist prejudice or, more subtly, being accepted but being given less weight than the viewpoint of an equally or less competent male juror.

Given the prevalence of these kinds of prejudices and the epistemic harms they lead to, one kind of epistemic value which groups might aspire to in settling the matter is to settle it in a way that mitigates, or is free from, epistemic injustices to individual members of the group. More precisely, let us suppose that a group adopts the following aim: to let *p* or else not-*p* stand as the group's view only if their doing so does not wrong any member specifically in her capacity as an epistemic subject (e.g., as a giver of knowledge, in her capacity for social understanding, and so on) or any other person outside the group in that capacity. Such a group, for short, adopts the aim of *epistemic justice*.

Of course, a group's aiming to issue an epistemic just verdict does not in any way *preclude* aiming at other epistemic goods, and in fact, it would be natural to expect that this aim will generally be paired with other aims. For example, a group might combine this aim with the aim of truth. In which case, the group aims to let *p* or else not-*p* stand as the group's view only if (i) *p* or else not-*p* is true; and (ii) their doing so does not wrong any member specifically in her capacity as an epistemic subject or any other person outside the group in that capacity.

Does voting or deliberation better facilitate epistemic justice? Let's begin by considering the following simple argument for voting: epistemic injustice (at least, of the testimonial variety of epistemic justice we're interested in) depends on testimonial exchange. Voting, but not deliberation, forecloses the possibility of testimonial exchange; so, voting, but not deliberation, forecloses a condition on which testimonial injustice depends. Therefore, voting better facilitates epistemic justice— specifically, by (unlike deliberation) blocking a condition necessary for its manifestation.

If the above argument is sound, then it looks as though voting should be favored to deliberation on epistemic justice grounds, even if it turns out that deliberation beats voting with respect to other epistemic goals.

Put perhaps there is space for the proponent of deliberation to press back along the following lines: even if deliberation is a precondition for epistemic injustice of the testimonial variety to occur, it remains that it is unlikely that testimonial injustice *will* occur in ordinary structured voting groups, e.g., such as juries, where there are norms in place already to give appropriate weight to individual viewpoints. For example, as this line of thought might go, juries are read instructions prior to deliberation that are meant to combat epistemically irresponsible assessment of the evidence, of which testimonial injustice is an instance. Likewise, other groups, particularly those with internal decision procedures that are structured around an office and a charter (Kallestrup 2016; cf., Pettit and Schweikard 2006), specify within that charter the *rules* by which a group will proceed towards its (joint) aims. So

long as such rules demand a fair evaluation of evidence during deliberation, they will *de facto* block epistemic injustice. Or so the thought might go.

The above reply is met with a rather straightforward counterreply. Empirical studies by Waters and Hans (2009) show how deliberation often does engender testimonial injustice in the case of juries (particularly those using unanimity rules), where the rules for evaluating evidence are, and *paradigmatically* so, meant to be impartial ones. What Waters and Hans (2009) found was that (in a study of 3500 jurors in four urban courts) 38% of juries contained at least one juror who succumbed to social pressure by voting along with the rest despite being such that they would have voted differently had they voted privately. (Waters and Hans 2009: 520) And, as Brian Hedden (2016) notes, such pressures are "likely to have a disproportionate impact on "low status" jurors, that is, females, members of minority ethnic groups, jurors with less education, jurors of low socioeconomic status, and the like" (2006: 7). Hedden reaches this conclusion on the basis of studies from, in particular, Christensen and Abbott (2000) and Hastie, Penrod, and Pennington (1983), which report findings that lower status jurors speak less, share less evidence, exert less influence, and are less likely to be elected as a jury foreperson.

These considerations support a presumptive case for thinking that deliberation will be positively correlated with epistemic injustice in a way that mere voting will not. However, even if this is granted, the proponent of deliberation has a card still to play: perhaps even if deliberation leads to epistemic injustice through the kinds of mechanisms described, it also facilitates at the same time epistemic *justice*—viz., perhaps deliberation is on the whole *more* epistemically just than it is unjust, even granting the kinds of considerations Hedden draws attention to.

One attempt to advance this kind of an argument draws from considerations about the *procedural value* of deliberation. According to Fabienne Peter (2013), there is a procedural epistemic value to deliberation which does not simply reduce to the epistemic value deliberation might have in so far as it brings about epistemic values such as truth, knowledge, etc.. The idea is as follows: deliberation (particularly when it involves epistemic peers), brings about relationships of *mutual accountability*, relationships that are characterized by (among other things) a respect for epistemic *equality* among group members. As Peter (2013) puts it:

> 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 [...] (2013: 1264)

Peter emphasizes that the value of this kind of mutual accountability, which is grounded in respect for epistemic equality, along with a "willingness to enter deliberation and to explicate one's beliefs; and [...] uptake" (2013: 1264), is procedural in that it does not "reduce to the value of its result" (2013: 1263). So, for Peter, the procedural value that mutual accountability adds to a correct group stance (e.g., p) is not 'swamped' by the value of that group's correct stance that *p*.

There are two worries for this argument. The first has to do with the swamping claim. It is not clear, without further argument, why the procedural value of Peter-style mutual accountability, in so far as this procedural value is meant to be *epistemic* (as opposed to, say, moral) is not simply swamped by the value of an epistemic end such as accuracy, truth, etc., toward which mutual accountability contributes. Second, the matter of how we *ought* to regard individuals we take to be epistemic peers and how we *in fact* are likely to treat such individuals can come apart. For example, even if juries ought to exhibit mutual accountability, it is a separate question whether they are inclined to meet this normative demand. The kinds of results Hedden draws attention to indicate that this normative demand is, in practice, often not met.

A proponent of deliberation might lean back, at this point, to a final kind of consideration: mere voting, absent deliberation, is a form of *silencing* (e.g., Barrett forthcoming; cf., Tanesini 2019; Langton 1993), in that one's assertions of one's reasons are *de facto* suppressed via the denial of the opportunity for any kind explanation of one's position. This silencing is injust, epistemically, in so far as it is suppressive; it suppresses one's capacity to justify her view— capacity that might arguably be viewed as a kind of epistemic *right* (e.g., Watson 2018)—as well as the possibility of having any such justification make a difference. The argument then proceeds as follows: the epistemic injustice of silencing through mere voting is a greater epistemic injustice than the kind of testimonial injustice brought about by deliberation.

While there is some intuitive pull to this line of thinking, it's not clear that it ultimately holds up. The reasoning is as follows: even if we grant that mere voting is a form of silencing, it is unclear that silencing is unjust, at least, in so far as it involves the *de facto* suppression of an opportunity to explain one's view. The thesis that a denial of the opportunity to explain or justify one's view (whenever one is permitted to register that view) constitutes unjust silencing overgeneralizes such as to generate the result that almost all standard presidential and political voting in liberal democracies involves unjust silencing. Secondly, beyond the overgeneralization worry, there is a further reason to resist thinking that that mere voting involves any epistemically unjust form of silencing. The argument is that whether or not a given restriction on the extent to which one may justify her view is unjust is context dependent. In the context where a group's views must all be taken into account, a better example of unjust silencing is *disenfranchisement* as opposed to voting in the absence of a capacity to provide additional reasons. And this is so even if the suppression of the opportunity to justify one's view does constitute unjust silencing in the context of a parent-teacher meeting, or a criminal trial.

In summary, we've seen in this section that epistemic justice is an epistemic aim that a group might reasonably adopt, along with any other epistemic aim (or set of epistemic aims), in its endeavor to settle a group view. With respect to this aim, deliberation was shown not only to be a precondition for a central species of epistemic injustice—testimonial injustice—but further, that there are reasons to expect that deliberation will in fact—and regularly does—contribute to this form of injustice in practice (Hedden 2016). In response to this worry for deliberation, we looked at whether the positive procedural value brought about by deliberation might compensate for this epistemic injustice (Peter 2013), and concluded in the negative. Finally, we considered whether there might be an epistemically pernicious form of *silencing* (Barrett *forthcoming*) that is brought about by mere voting that might be in itself a kind of epistemic injustice that is more serious than what is engendered by deliberation, and the argument for this suggestion was ultimately unconvincing. With respect to the aim of epistemic justice, then, it looks as though voting is going to be more efficacious than deliberation.

6 Ways to mitigate the epistemic disadvantages of deliberative intragroup disagreement
In this section, we will address the *deliberation question*—recall: what it would take to overcome or at least mitigate the epistemic disadvantages of solving intragroup disagreement by means of deliberation? Given the pluralist approach to deliberation that we have adopted in this paper, namely that a group's collective endeavor to solve an internal dispute can be aimed at different (albeit not necessarily incompatible) epistemic goals, it comes as no surprise that this question has no unique answer. To put it differently, for each epistemic goal a divided group might aim to when settling on a collective view, there are specific deliberative conditions that can make that view fail to satisfy that goal. Indeed, for some such goals we have concluded that deliberation is *not*, after all, an epistemically appropriate procedure for groups to solve their internal disputes.

This doesn't mean, of course, that there is nothing groups can do to overcome the epistemic shortcomings of deliberative processes. To put things into perspective, let us briefly summarize our main conclusions on deliberation for the four epistemic goals we have considered and, on that basis,

and also on the basis of empirical results, provide some answers to the deliberation question for each of the goals.

6.1 The truth and the evidence goals: the epistemic significance of being divided
Concerning the truth goal, we have seen—drawing on social choice theory—that deliberating *vis-à-vis* the goal of reaching true collective agreements is roughly as epistemically appropriate as voting by majority rule in relatively idealized conditions (this is the result of Hartmann and Rafiee Rad's Bayesian model of deliberation). However, we have also seen—drawing on empirical social psychology—that real-life deliberative situations may involve an array of factors that bear on the reliability of deliberation.

Among those factors, two phenomena that operate at the group level constitute two particularly significant threats to the reliability of deliberation: group polarization and shared information bias. We have also seen, however, that these two reliability-undermining phenomena lose their influence in groups featuring internal disagreements (albeit perhaps not to the point of disappearing).

This results in a rather paradoxical situation: groups whose members would only settle their internal disagreements if the resulting settled views were true can more reliably achieve this goal when their disagreements are pronounced. To put it differently, groups whose members are starkly divided over some issue (e.g., with a 50/50 distribution) might be in a more solid epistemic position to avoid such reliability-undermining phenomena and thus to solve their internal disputes reliably than groups that feature less pronounced disagreements.

Groups can intensify their internal disputes in a number of ways. By way of illustration, a group of 100 members such that, e.g., 99 hold the false proposition p, whereas one holds that not-p might avoid or mitigate reliability-undermining group phenomena such as group polarization or shared information bias (i) by increasing the number of members who defend not-p,[27] (ii) by making some members play the devil's advocate role in defense of the minority view;[28] or (iii) by implementing the previously discussed *advocacy procedure*, tested by Greitemeyer et al. (2006), according to which each group member acts as an advocate for each alternative for some time independently of their individual preferences.

As we have seen in previous sections, increased disagreement within the group (e.g., by means of the just-mentioned advocacy procedure) results in better pooling of the evidence—including the group members' private evidence—, which raises the chances that the collectively accepted view is supported by the best evidence available within the group, which would satisfy the evidence goal.

Accordingly, one answer to the deliberation question could be the following: the very fact of being internally divided makes groups featuring internal disagreements be more protected from group phenomena that make it less likely that they solve their internal disagreements in such a way that they meet the truth and the evidence goals.

Of course, this might not be enough for deliberation to be conducive to such goals and there are certainly more things that groups can do to ensure a better pooling of the evidence and increased accuracy—such as involving team leaders and members with experience at the relevant task in the deliberative process, giving enough time for discussion or having direct access to the evidence (as we have seen in §3). The bottom line, at any rate, is this: what might be considered a disadvantage for many reasons—viz., that a group is internally divided over some issue—turns out to be an epistemic advantage when the group aims to solve such a dispute by means of deliberation with an eye on reaching a true and evidentially well-supported collective agreement.

6.2. The understanding goal: individual competence and the explanatory value of the evidence
We have offered two interpretations of the understanding goal. On the first interpretation, members of a group would only let a view stand as the group's view if the settled view is made *on the basis* of the group's understanding; on the second interpretation, only if it *promotes* the group's understanding. We

have argued, in both cases, that deliberation is superior to voting. However, this doesn't mean that deliberation is always epistemically appropriate.

Consider the first interpretation of the goal. When group members are individually incompetent, it is unlikely that they will grasp the relevant issue in the way required by the goal. Accordingly, one way in which a group can mitigate this shortcoming is by increasing the individual competence of its members. This can be done in several ways. For example, changes can be made to group membership, and incompetent members be excluded from the group, or perhaps, only, from discussion. Another way is to ensure that internal or external experts explain the relevant subject matter to incompetent group members.

Consider now the second goal, for example, in the case of understanding-why. We've argued that one epistemic advantage of deliberation over voting is that the sharing of evidence engages with not *merely* the matter of *what*, e.g., caused X, but *how* it did so, which better contributes to reaching the goal of promoting understanding among group members. However, whether or not deliberation is capable of this does not only depend on the sharing of the evidence, but also on its explanatory value.

Not all evidence is equally explanatory, where this might depend on several factors. For example, sharing with one's group a newspaper article on why SARS-CoV-2 is as infectious as it is is not as explanatory as sharing the specific studies that pin down the particular transmission channels. In some cases, by contrast, e.g. when group members are not particularly competent or lack the relevant expertise, sharing *less* detailed information might facilitate that they grasp the relevant issue more easily than sharing very detailed information they are in no position to understand. Sometimes, it is the *amount* of evidence, not just its quality, that matters for its explanatory value: sharing too much information with group fellows, even if high-quality, can lead to information overload, which rather than promoting group understanding may hinder it.

In this way, two things groups can do to overcome or at least mitigate the epistemic limitations of deliberation *vis-à-vis* the understanding goal is to increase, first, the degree of individual competence of their members and, second, the explanatory value of the evidence, where this will depend, among other things, on increasing the quality of the information shared as well as on controlling the amount of information shared.

One way to improve on both factors (individual competence and the explanatory value of the evidence) and thus to facilitate group understanding is to implement a method akin to the advocacy procedure discussed in §3—to our knowledge, this has not been empirically tested. In the advocacy procedure, each group member acts as an advocate for each alternative for some time independently of their individual preferences. In the method we envisage—call it the *pedagogical procedure*—, the competent members of the group (if any), or all members in groups of competent epistemic peers, act as pedagogues for the rest in the following way: all take care of explaining to fellow members the specifics of the collective views in dispute, in the best of their capacity and drawing on the best of their evidence. In cases where this policy cannot be applied, e.g., because no one in the group is competent, groups can resort to external experts who can give such a pedagogical service to group members.

6.3 The epistemic justice goal: smaller discussion groups and computer-mediated communication
Concerning the aim of epistemic justice, we've argued that deliberation is epistemically inappropriate for it in that deliberation is not only a precondition for testimonial injustice, but also regularly contributes to this form of injustice in practice. As empirical studies on juries have shown, deliberation puts social pressure on jurors to vote differently than if they had voted privately, makes lower status jurors speak less, share less evidence, exert less influence, and be less likely to be elected as a jury foreperson.[29]

One way in which participation can be enhanced during group discussion—thus minimizing the risk of testimonial injustice—is by reducing the size of the group (e.g., by splitting discussion into

smaller subgroups). As it turns out, the bigger the group, the more unequal participation is; and the smaller the group, the more equal it is (Bonito and Hollingshead 1997).

In addition, since group members with lower status participate less and are less influential than those with higher status—which can no doubt lead to testimonial injustices if low status is allocated on non-epistemic grounds (e.g., prejudice, implicit bias)—, then one way to mitigate this epistemic flaw of deliberation is to prevent deliberators from accessing those cues that make them attribute status to other group members, including demographic cues such sex, race or age. Unsurprisingly, face-to-face communication makes such cues more readily accessible to deliberators, which can cause testimonial injustices more easily. By contrast, one way to make participation levels more equal—and thus mitigate testimonial injustice—is computer-mediated communication, which makes the cues that serve to allocate status private, at least in the short term until online group hierarchies emerge (Hollingshead 2001).[30]

Concluding remarks

Deliberation can be assessed from many angles. We have assessed it epistemically. In particular, we have investigated to what extent it is epistemically advantageous and disadvantageous that groups whose members disagree over some issue use deliberation in comparison to voting as a way to reach collective agreements. The way we have approached this question is from a pluralist perspective. We have assumed that a group's collective endeavor to solve an internal dispute can be aimed at different, albeit not necessarily incompatible, epistemic goals, namely the goals of truth, evidence, understanding, and epistemic justice. For the goals of truth and evidence we have explained, drawing on social choice theory, that deliberation and voting are epistemically on a par. But we have also shown how complex it is to give a straightforward answer to the question of how reliable deliberation is as a method for solving intragroup disagreement. This complexity, we have argued, has to do with the interplay between the different kinds of evidence involved in deliberation as well as with several group phenomena widely investigated in empirical social psychology such as group polarization and shared information bias. Concerning the goal of understanding, we have given two interpretations of the goal (the goal of reaching a collective view about why X only if the group understands why X and the goal of reaching a collective view only if it promotes the group's understanding). In both cases, we have argued, deliberation outperforms mere voting. Concerning the epistemic justice goal, however, we have concluded that voting is more efficacious than deliberation. Finally, we have discussed several ways to mitigate the potential epistemic disadvantages of solving intragroup disagreement by means of deliberation in relation to each epistemic goal.[31]

References

Boland, P.J. (1989). 'Majority Systems and the Condorcet Jury Theorem'. *The Statistician* 38 (3): 181–9.

Bonito, J. & Hollingshead, A.B. (1997). 'Participation in Small Groups'. *Annals of the International Communication Association* 20(1): 227-261.

Boyd, K. (2019). 'Group Understanding'. *Synthese*, 1-22. doi: 10.1007/s11229-019-02492-3

Brodbeck, F. C., Kerschreiter, R., Mojzisch, A., Frey, D., & Schulz-Hardt, S. (2002). 'The Dissemination of Critical, Unshared Information in Decision-Making groups: The effects of pre-discussion dissent. *European Journal of Social* Psychology (32): 35–56.

Broncano-Berrocal, F. & Carter, J.A. (Forthcoming). *The Philosophy of Group Polarization.* London: Routledge

Burnstein, E., & Vinokur, A. (1977). 'Persuasive Argumentation and Social Comparison as Determinants of Attitude Polarization'. *Journal of Experimental Social Psychology* 13 (4): 315–32.

Carey, B. and Matheson, J. (2013). 'How Skeptical is the Equal Weight View?' In D. Machuca (ed.), *Disagreement and Skepticism*, New York: Routledge, 131–49.

Chapman, G. B., & Sonnenberg, F. A. (Eds.). (2003). *Decision Making in Health Care: Theory, Psychology, and Applications*. New York: Cambridge University Press.

Christensen, C., & Abbott, A. S. (2003). '10 Team Medical Decision Making' In Chapman , G.B. and Sonnenberg, F.A. (Eds). *Decision Making in Health Care: Theory, Psychology, and Applications*. 267.

Devine, D. J., Clayton, L. D., Dunford, B. B., Seying, R., & Pryce, J. (2001). 'Jury Decision Making: 45 Years of Empirical Research on Deliberating Groups'. *Psychology, Public Policy, and Law*, *7*(3): 622.

Dietrich, F. (2006). 'General Representation of Epistemically Optimal Procedures', *Social Choice and Welfare*, 26(2): 263–283.

Dietrich, F. & List C. (2004). 'A Model of Jury Decisions Where All Jurors Have the Same Evidence'. *Synthese* 142 (2): 175–202.

—— (2007). 'Judgment Aggregation by Quota Rules: Majority voting generalized'. *Journal of Theoretical Politics* 19(4): 391-424.

Dietrich, F. & Spiekermann K. (2013) 'Epistemic Democracy with Defensible Premises'. *Economics and Philosophy* 29 (1): 87–120.

—— (2020). 'Jury Theorems'. In M. Fricker (ed.), *The Routledge Handbook of Social Epistemology*. New York and Abingdon:

Elga, A. (2010). 'How to Disagree About How to Disagree'. In R. Feldman and T. Warfield (eds.), *Disagreement*. New York: Oxford University Press.

Emerick, B. (forthcoming). 'The Violence of Silencing'. In J. Kling (ed*.), Pacifism, Politics, and Feminism: Intersections and Innovations*. Brill.

Fallis, D. (2007). 'Collective Epistemic Goals'. *Social Epistemology* 21: 267–280.

Fallis, D. & Mathiesen, K. (2013). 'Veritistic Epistemology and the Epistemic Goals of Groups: a Reply to Vähämaa'. *Social Epistemology* (27): 21–25.

Fallis, D. (2007). 'Collective Epistemic Goals'. *Social Epistemology* 21: 267–280.

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.

—— (2010). 'Can There Be Institutional Virtues?' In T. Szabo Gendler & J. Hawthorne (eds.), *Oxford Studies in Epistemology*. Oxford: Oxford University Press.

Goethals, G. R., & Zanna, M. P. (1979). 'The Role of Social Comparison in Choice Shifts'. *Journal of Personality and Social Psychology* 37: 1469–1476.

Goldman, A.I. (1999). *Knowledge in a Social World*. New York: Oxford University Press.

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

Gradstein, M., & Nitzan, S. (1986). 'Performance Evaluation of Some Special Classes of Weighted Majority Rules'. *Mathematical Social Science* 12: 31–46.

Grofman, B., Owen, G., and Feld, S. (1983). 'Thirteen Theorems in Search of the Truth'. *Theory and Decision* 15 (3): .261–78

Hastie, R., S. Penrod, & N. Pennington. (1983). *Inside the Jury*. Cambridge, MA: Harvard University Press.

Hedden, B. (2017). 'Should Juries Deliberate?' *Social Epistemology*, DOI: 10.1080/02691728.2016.1270364

Hightower, R., & Sayeed, L. (1996). 'Effects of Communication Mode and Prediscussion Information Distribution Characteristics on Information Exchange in Groups'. *Information Systems Research* 7: 451–465.

Hollingshead, A. B. (2001). 'Communication Technologies, the Internet, and Group Research'. In M. A. Hogg & R. S. Tindale (Eds.), *Blackwell Handbook of Social Psychology: Group Processes* (pp. 557–573). Malden, MA: Blackwell.

Isenberg, D. J. (1986). 'Group Polarization: A Critical Review and Meta-Analysis'. *Journal of Personality and Social Psychology* 50(6): 1141.

Kallestrup, J. (2016). 'Group Virtue Epistemology'. *Synthese*, 1-19.

Kelp, C. (2014). 'Two for the Knowledge goal of Inquiry'. *American Philosophical Quarterly*, 51(3): 227-232.

Kvanvig, J. L. (2013). 'Curiosity and the Response-Dependent Special Value of Understanding'. In T. Henning and D. Schweikard (eds.), *Knowledge, Virtue and Action: Putting Epistemic Virtues to Work*, Routledge: 151-74.

Lackey, J. (2013). 'Disagreement and Belief Dependence: Why Numbers Matter'. In D. Christensen and J. Lackey (eds.), *The Epistemology of Disagreement: New Essays*, Oxford: Oxford University Press: 243–68.

—— (2016). 'What is Justified Group Belief'. *Philosophical Review* 125: 341–396.

Ladha, K. (1992) 'The Condorcet Jury Theorem, Free Speech and Correlated Votes'. *American Journal of Political Science* 36 (3): 617–34.

Larson, J. R., Foster-Fishman, P. G., & Keys, C. B. (1994). 'Discussion of Shared and Unshared Information in Decision-Making Groups'. *Journal of Personality and Social Psychology* 67(3): 446–461.

Levine, J. M., & Moreland, R. L. (Eds.). (2006). *Small Groups: Key Readings.* Psychology Press.

List, C. & R. E. Goodin. (2001). 'Epistemic Democracy: Generalizing the Condorcet Jury Theorem'. *Journal of Political Philosophy* 9: 277–306.

List, C. & Spiekermann, K. (2016). 'The Condorcet Jury Theorem and Voter-Specific Truth'. In H. Kornblith & B. McLaughlin (eds.). *Alvin Goldman and his Critics*. Oxford: Wiley Blackwell: 219-231.

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

Matheson, J. (2015). *The Epistemic Significance of Disagreement*'. Dorcrecht: Springer.

Miller, B. (2013). 'When is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement'. *Synthese* 190: 1293–1316.

Nemeth, C., Brown, K. and Rogers, J. (2001). 'Devil's Advocate Versus Authentic Dissent: Stimulating Quantity and Quality'. *European Journal of Social Psychology* 31: 707-720.

Nitzan, S., & Paroush, J. (1982). 'Optimal Decision Rules in Uncertain Dichotomous Choice Situations'. *International Economic Review* 23(2): 289–297.

Pacuit, E. 2019. Voting methods. *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2019/entries/voting-methods/>.

Peter, F. (2013). 'The Procedural Epistemic Value of Deliberation'. *Synthese* 190: 1253–1266.

Pettit, P., & Schweikard, D. (2006). 'Joint Actions and Group Agents'. *Philosophy of the Social Sciences* 36(1): 18–39.

Sliwa, P. (2015). 'IV—Understanding and Knowing'. *Proceedings of the Aristotelian Society* (Vol. 115(1): 57-74.

Stasser, G., & Stewart, D. (1992). 'Discovery of Hidden Profiles by Decision-Making Groups: Solving a Problem Versus Making a Judgment'. *Journal of Personality and Social Psychology* 63(3): 426–434.

Stasser, G., & Titus, W. (1985). 'Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion'. *Journal of Personality and Social Psychology* 48(6): 1467–1478

Stoner, J.A.F. (1961). 'A Comparison of Individual and Group Decision Involving Risk'. Massachusetts Institute of Technology.

Tanesini, A. (2019). 'Silencing and Assertion'. In Sanford Goldberg (ed.), *The Oxford Handbook of Assertion*. Oxford, U.K.: Oxford University Press. pp. 749-769.

Tucker, A. (2003). 'The Epistemic Significance of Consensus'. *Inquiry* 46: 501–521.

Waters, N., & V. Hans. 2009. 'A Jury of One: Opinion Formation, Conformity, and Dissent on Juries'. *Cornell Law Faculty Publications*, Paper 114.

Weatherson, B. (2013). 'Disagreements, Philosophical and Otherwise'. In J. Lackey and D. Christensen (eds.), *The Epistemology of Disagreement: New Essays*, Oxford University Press.

Watson, L. (2018). 'Systematic Epistemic Rights Violations in the Media: A Brexit Case Study'. *Social Epistemology* 32(2): 88-102.

Wray, K. B. (2001). 'Collective Belief and Acceptance'. *Synthese* 129: 319–333.

NOTES

1 Collective agreement is a gradable notion: the more members of a group *G* agree on *p*, the broader the collective agreement. When a collective agreement is complete (all members of *G* agree on *p*) or else when it is broad (many members of *G* agree on *p*), we call it 'consensus' (see Tucker 2003: 509-510 for the former view; Miller 2013 for the latter). For our purposes, the kind of collective agreement we are interested in need not be consensual, but sufficiently broad for the relevant view to be considered the group's view.

2 That members of a group agree to take a certain course of action, φ, as a group is a special case of letting a view stand as the group's view, namely the view that the group will or ought to φ.

3 *Operative members* are, according to Lackey, those who "have authority or power to determine certain outcomes for the group as a whole" (Lackey 2016: 350). In contrast, *passengers* are group members with no or little authority or power that simply go along with the resolutions and decisions of operative members (cf. Fricker 2010).

4 For simplicity, we will not consider the third kind of view members of deliberative and non-deliberative groups might agree on to feature as the group's view: *suspension of judgment*, i.e., neither endorsing *p,* nor not-*p*.

5 In §3, we will consider borderline cases of groups that are motivated by non-epistemic reasons but that nevertheless pursue epistemic goals.

6 For discussion on how groups can aim at epistemic goals *qua* groups, see Fallis (2007).

7 For a similar methodological approach to a different question (whether or not deliberation has procedural in addition to instrumental value), see Peter (2013: fn. 12)

8 Epistemic justice, as a goal, is not purely epistemic: epistemic *in*justice, as Fricker (2007) conceives it, involves both epistemic and moral harm. Still, given the significance of this notion in recent epistemological theorizing, it is interesting to assess what kind of method to solve disagreement fares better with respect to it: deliberation or voting.

9 Other epistemic goals might be the following: the collectively accepted view (CAV) is known by the group as a whole (*non-summativist knowledge*); CAV is epistemically justified (*justification*); CAV coheres better with other collective views of the group than the opposite view (*coherence*). CAV is a view upon which individual group members would justifiedly place her trust (*epistemic trust*).

10 For example, if veritism (see fn. xi) is correct and truth is the primary epistemic goal so that the epistemic value of other goals is merely instrumental to it, then if, let's suppose, taking a vote is more conducive to truth than deliberation, we should certainly prefer the former.

11 The thesis that truth is the fundamental epistemic goal is often referred to as *veritism* (e.g., Goldman 1999). Cf., Kelp (2014) and Kvanvig (2013) for criticism.

12 See Lackey (2013) and Carey and Matheson (2013) for the epistemic significance of numbers in responding to disagreement.

13 In non-deliberative cases, social evidence can be part of the group members' private evidence (e.g., evidence from polls). However, social evidence is not always available in the case of voting as it is in the case of deliberation.

14 Of course, the kind of admittedly idealized deliberative cases we are considering are cases in which group members aim to reach a collective agreement pursuing the epistemic goal of truth. This plausibly suggests that all group members are sincere in their assertions. Still, even in these idealized cases, other group members can still *incorrectly* judge that fellow members are not being sincere, which can potentially lead to unreliable assessments about the distribution of opinions within the group.

15 Exposing group members to the positions of other group members in the absence of any other information has been observed to make them shift their own positions in a direction that is consistent with the average position of the group, where this is explained by social comparison mechanisms (Goethals *&* Zanna 1979).

16 For a comprehensive critical review of the different positions in the epistemology of disagreement, see Matheson (2015). For how the implications of conciliationism for intragroup disagreement, see Sheff (this volume) and Skipper Rasmussen & Steglich-Petersen (this volume).

17 See Dietrich and Spiekermann (2020) for an explanation of why this sort of conditionalization is needed.

18 As List & Spiekermann (2016) helpfully explain the two conditions in an informal fashion: "[The competence condition] says that each voter is better than random at identifying the truth, and that different voters are equally reliable. Furthermore, this holds irrespective of whether the truth is [*p* or not-*p*]. (...) "[The independence condition] says that once we hold the truth in question fixed—thereby conditionalizing either on [*p* or not-*p*]—learning the votes of some voters does not give us any information about the votes of others." (List & Spiekermann 2016: 221)

19 To remedy this, Dietrich and Spiekermann (2013) distinguish causal from probabilistic independence, and formulate a version of CJT that conditionalizes on any common cause.

20 As Dietrich & Spiekermann (2020) helpfully explain, one thing that many jury theorems aim to establish is what they call the *growing reliability thesis*, namely that larger groups are better truth-trackers, in the sense that they are more likely to select the correct alternative (by majority) than smaller groups or single individuals.

21 Note, however, that in the case of large groups an increase in individual reliability doesn't necessarily translate into an increase in collective reliability. As Goodin and Spiekermann (2018: 138) point out "For a large group, any mean individual competence level appreciably above random will have *almost* 'maxed out' the 'wisdom of crowds' effect already. Increasing mean individual competence further does not, therefore, have much effect for large groups."

22 For a book-length of the philosophical aspects of group polarization, see Broncano-Berrocal & Carter (*forthcoming*). See also Olsson (this volume).

23 These cases differ with respect to the cases we've been considering so far in that the propositions in dispute are both false, while the previous examples are such that the relevant disagreement is between a proposition and its negation.

24 This kind of position is not far from internalism in traditional individualistic epistemology: after all, many internalists rank higher a false justified belief (e.g., the perceptual beliefs of brains in vats) than a true unjustified belief (e.g., a lucky guess in normal circumstances).

25 For defenses of the thesis that understanding-why is a species of propositional knowledge, see for example Grimm (2006), Sliwa (2015), Lipton (2004), and Woodward (2002). See Kelp (2015; 2017) and Sliwa (2014) for a defense of the view that objectual understanding reduces to propositional knowledge.

26 A separate strand of epistemic injustice, which we'll set aside for the present purposes, is called *hermeneutical injustice*, which occurs when individuals are deprived in unjust ways of the opportunity to conceptualize their own experiences.

27 Concerning group polarization more specifically, there is an extensive empirical literature on why groups polarize (see Broncano-Berrocal and Carter *forthcoming*, Ch. 2., for a review), but not so many studies on group depolarization—for some exceptions see Abrams et al. (1990) and Vinokur & Burnstein (1978). Yet, group polarization has been observed in homogeneous groups, not in diverse ones, which gives reason to think that increasing group diversity translates into less polarization.

28 But see Nemeth et al. (2001), who demonstrate that authentic dissent leads to better information processing and decision making that the devil's advocate strategy.

29 See Devine et al. (2000) for a comprehensive review of the empirical research on jury deliberation.

30 As Levine & Moreland (2006) point out, a particular technology for which there is strong evidence of equalizing participation and priming informational over normative influence is email.