

Disagreement within Rational Collective Agents

Javier González de Prado Salas and Xavier de Donato-Rodríguez

**This a draft of a chapter to appear in F. Broncano-Berrocal and A. Carter (eds.), *The Epistemology of Group Disagreement*, Routledge. Please refer to the published version when possible.*

Abstract

The question we want to address is what it takes for groups to constitute rational epistemic agents. The members of a group often hold disagreeing views. How can a group rationally move from such situations of internal disagreement to a unified group attitude? A possible answer is that group attitudes are rational if they result from the application of appropriate judgment aggregation methods. In Section 3.2 we discuss some problematic aspects of this answer, and then, in Section 3.3, we present an alternative proposal, according to which group (epistemic) attitudes are rational insofar as they are formed by responding competently or responsibly to the (epistemic) reasons available to the group as a group (this will require exercises of reasons-responding competences attributable to the group). In Section 3.4, we explore the idea that bare judgment aggregation methods have to be combined with collective deliberation in order for groups to be able to respond competently to reasons. In Section 3.5 we discuss the extent to which collective deliberation can be expected to solve internal group disagreements, paving the way for the adoption of coherent, reasons-responsive group attitudes. We suggest that conciliationist approaches to disagreement offer an optimistic picture of collective deliberation, as a method for bringing about internal group consensus. However, we will also explore possible limitations to the application of the conciliationist picture to realistic instances of group deliberation. Finally, in Section 3.6, we examine the role of dissensus and consensus in epistemically virtuous group deliberations.

3.1 Introduction

Groups are often treated as (collective) rational agents. More specifically, it is common to attribute to groups epistemic states and attitudes, such as knowledge or justified belief and acceptance. Imagine, for example, a speaker saying that a certain company knows that its

products are especially popular among young people. Not only does this way of talking accord with ordinary discourse, but it is also vindicated by several authors in the literature (for instance, List and Pettit 2011; Bird 2014; González de Prado and Zamora-Bonilla 2015; Kallestrup 2016; Hedden 2019; see Tollefsen 2015 for a survey). In this essay, we will assume that groups sometimes count as rational agents.

The question we want to address is what it takes for groups to constitute rational epistemic agents. The members of a group often hold disagreeing views. How can a group rationally move from such situations of internal disagreement to a unified group attitude? A possible answer is that group attitudes are rational if they result from the application of appropriate judgment aggregation methods. In Section 3.2, we discuss some problematic aspects of this answer, and then, in Section 3.3, we present an alternative proposal, according to which group (epistemic) attitudes are rational insofar as they are formed by responding competently or responsibly to the (epistemic) reasons available to the group as a group (this will require exercises of reasons-responding competences attributable to the group). In Section 3.4, we explore the idea that bare judgment aggregation methods have to be combined with collective deliberation in order for groups to be able to respond competently to reasons. In Section 3.5, we discuss the extent to which collective deliberation can be expected to solve internal group disagreements, paving the way for the adoption of coherent, reasons-responsive group attitudes. We suggest that conciliationist approaches to disagreement offer an optimistic picture of collective deliberation, as a method for bringing about internal group consensus. However, we will also explore possible limitations to the application of the conciliationist picture to realistic instances of group deliberation. Finally, in Section 3.6, we examine the role of dissensus and consensus in epistemically virtuous group deliberations.

3.2 Disagreement and Aggregation-Procedures in Groups

Groups that constitute rational collective agents can take part in epistemic practices in a similar way to individual agents. In particular, collective agents can engage in testimonial exchanges and be involved in arguments and disputes with others. Think for instance of a company arguing in court, or a national government disagreeing with the views of some international institution. Thus, once we allow for the possibility of collective agency, it is natural to think that groups can disagree with other agents (whether individual or collective). This type of disagreement would be

with agents that are not acting as members of the group, so let us call it *external collective disagreement*. External collective disagreement would be analogous to standard cases of disagreement among individual agents.

We want to argue, however that a crucial difference between collective and individual agents is that *internal disagreement* is a far more common occurrence in the former than in the latter. By internal collective disagreement we will refer to cases in which there is disagreement *within the group*. So, internal disagreement will take place in a group when its members disagree among themselves while acting as members of the group.¹ At least at certain stages of collective decision-making, some level of internal disagreement is common in most groups. In this way, mechanisms of judgment aggregation such as voting are devised to generate judgments at the collective level starting from situations of internal disagreement among the group's members. By contrast, internal disagreement does not typically happen in individual agents. Or, less controversially, if it happens, it plays a far less general and pervasive role than in groups. Possible situations of internal disagreement in individuals seem interestingly different from the ordinary forms of internal disagreement in groups discussed in the previous paragraph.² An individual may vacillate or be in a situation of uncertainty, but in normal cases there are no disagreements *within* individuals. Individual deliberation does not normally involve disagreements among voices or attitudes internal to the individual, since individuals are not composed of agents with potentially disagreeing stances (for a similar point, see Epstein 2015: 248). The goal of individual deliberation is precisely to make up the individual's mind, by settling what attitudes she ought to adopt, rather than being a process aimed at aggregating preexisting attitudes or judgments of the individual. In the case of groups it is possible to make a clear distinction between the judgments of the group at the collective level and the judgments of its members (in particular, disagreeing judgments at the level of the members can in principle co-exist with a single, unified judgment at the collective level). No such distinction can be made in normal individual agents.

To be sure, individual epistemic deliberation is sometimes seen as a matter of assessing and weighing the evidence available (more generally, a process of balancing the epistemic reasons accessible to the agent). One might think that such weighing deliberations can be modeled as processes with the aim of solving disagreements among different pieces or sources of evidence. More specifically, one can be tempted to characterize individual deliberations by

analogy to judgment aggregation methods, and in particular voting procedures. On this model, the votes in favor and against a certain attitude would be determined by the strength of the evidence for and against it.

However, we should be wary of taking this voting analogy other than in a metaphorical way. First of all, it is not clear whether (individual) epistemic deliberation is adequately characterized by appeal to the picture of the balance of reasons, that is as a process of weighing evidential reasons for and against the relevant attitudes (see Titelbaum 2019; González de Prado 2019a). Anyway, even if we accept this picture, it seems that evidence weighing and amalgamation just does not work like standard processes of voting for different options (or, more generally, like judgment aggregation methods taking as inputs disagreeing judgments). Obviously, pieces or sources of evidence are not agents casting votes expressing their judgments or preferences. Less trivially, evidential reasons interact among themselves in ways that are not easily captured by standard voting methods.

One first way in which reasons may interact is by rebutting or outweighing each other. Reasons for a given option are outweighed when there are stronger opposing reasons against that option. Outweighing may perhaps be conceived of by analogy with voting procedures – stronger reasons would be associated with more votes, and would therefore defeat weaker reasons bringing in fewer votes. However, reasons can also undercut or attenuate each other (Pollock 1987; Dancy 2004; Schroeder 2007). An attenuator is a consideration that reduces the strength or weight of a certain reason. In the extreme cases of attenuation, the reason ends up having no weight at all and is said to be undercut or disabled. Undercutting is a common phenomenon in epistemic deliberation. For instance, a testimonial report will lose its evidential weight if it is found that the testifier is unreliable or has deceiving intent. Likewise, the measurements made by some thermometer will stop being counted as providing evidential reasons to form beliefs about the environmental temperature if we know that the thermometer is broken.

Undercutting and attenuation are not easily captured by accounts of individual deliberation framed in terms of standard voting methods. In a voting procedure, cases analogous to undercutting defeat would be situations where the votes of certain agents stop being counted because of the votes casted by other voters. That is, certain votes would have the effect of disenfranchising some other voters who were in principle granted suffrage. Note, moreover, that these disenfranchising powers would depend not only on who the voters are (i.e. the content of

the relevant epistemic reasons and defeaters), but also on the issue voted (i.e. the options the agent deliberates about). Furthermore, reasons and defeaters interact in complex, holistic ways: undercutting defeaters can be themselves undercut, and whether a certain consideration has undercutting power may depend in open-ended ways on surrounding features of the issue deliberated about (Dancy 2004). Perhaps it is possible to come up with exotic voting methods that mimic the holistic, occasion-sensitive features of interactions among reasons. Yet, these methods would be so far away from standard voting systems that it would stop making sense to appeal to the voting analogy.

It is well known, on the other hand, that standard judgment aggregation methods sometimes lead to incoherent, or otherwise irrational collective judgments or attitudes despite taking as inputs perfectly rational individual attitudes. An illustrative example is what Pettit calls the *discursive dilemma* (Pettit 2001; List and Pettit 2011). Imagine that we aggregate the judgments of a group of agents on several propositions using a majority voting method on each of them, so that a proposition is accepted if and only if the majority of the group members accept it. Let us apply this method to the propositions p , q , and $p \& q$. It is possible to have a group composed only by members with rational attitudes such that 55% of the group accepts p , 55% accepts q , but only 10% accepts $p \& q$. In this case, the group would end up with incoherent collective attitudes, since it would accept p and it would accept q , but nonetheless it would reject $p \& q$. This problem generalizes to more sophisticated judgment aggregation strategies. As List and Pettit (2011) have shown, for a large class of judgment aggregation methods, there will be cases in which rational attitudes at the level of the members of the group will lead to an irrational aggregated attitude at the collective level. This pessimistic thesis mirrors similar results in social choice theory, such as Arrow's impossibility theorem.

Thus, judgment aggregation methods like voting are not just unsuited as models for individual deliberation; it seems that, in addition, such methods are not always, on their own, good means to form rational, coherent collective attitudes in groups (Buchak and Pettit 2014; Hedden 2019). One could conclude that this reveals a limitation in the rationality of groups: groups can have irrational attitudes even when they are formed by aggregating individual attitudes that are themselves rational, and that were adopted by the members by responding properly to their epistemic reasons. An alternative view is that what makes the collective attitudes of a group rational is not so much whether they were formed following some particular

judgment aggregation method, but rather whether they properly respond to the reasons accessible to the group as a collective agent. We will explore here this type of approach, which has been recently advocated by Hedden (2019; also González de Prado and Zamora-Bonilla forthcoming).

3.3 Group Rationality as Reasons-Responsiveness

Among the different ways of thinking about rationality discussed in the literature, it is possible to distinguish two main trends. On the one hand, there is coherence-based rationality. According to this first approach, rationality is a matter of satisfying coherence requirements, such as avoiding inconsistent beliefs (for instance, Broome 2007a, 2007b). On the other hand, there is reasons-based rationality. On this second approach, being rational amounts to responding properly to one's (apparent) reasons (Schroeder 2007; Parfit 2011; Kiesewetter 2017; Lord 2018).

We will focus here on this second conception of rationality, first because we think that it leads to a more attractive picture of group rationality, and second because coherence-rationality can arguably be derived from a view of rationality as reasons-responsiveness.³ From the perspective of this approach to rationality, the natural thing to say, following Hedden (2019), is that a group is rational insofar as its attitudes are properly sensitive to the reasons available to the group *as a collective agent*. As is customary to do, we will think of normative reasons as considerations that favor or recommend a certain attitude. On our preferred account, a certain attitude is rational for a (collective or individual) agent if and only if such an attitude is sufficiently supported by the agent's *apparent* reasons, that is by the considerations that appear to the agent as reasons for the attitude (Schroeder 2007; Parfit 2011; for different ways of understanding the notion of apparent reason, see Whiting 2014; Sylvan 2015). According to this view, an attitude can be made rational by considerations that merely appear to be reasons to the agent (e.g. false beliefs supported by convincing but ultimately misleading evidence). However, for our purposes in this paper, it is also possible to adopt the view that rationally permissible attitudes are those sufficiently supported by *actual* available reasons (not by merely apparent ones). This type of view has been recently defended by Kiesewetter (2017) and Lord (2018). For the sake of simplicity, we will assume this latter view here. Nothing substantial hangs on this choice.

As pointed out above, a plausible idea is that coherence is an offshoot of reasons-responsiveness, in the sense that responding properly to one's reasons ensures that one's attitudes

are coherent (Kolodny 2007; Kieseewetter 2017; Hedden 2019; Lord 2018). The idea is that incoherent attitudes cannot receive decisive support from the same set of reasons. For instance, the belief that p and the belief that $\neg p$ cannot both receive decisive support from the same body of epistemic reasons, and therefore it can never be rational for an agent to hold beliefs she takes to be inconsistent (it can never appear to the agent that her reasons support a combination of beliefs she knows to be inconsistent). If this idea is on the right track, it will apply generally to attitudes formed by responding properly to some set of reasons, regardless of whether they are individual or collective attitudes. Hedden (2019) has argued that this is the case: as long as a collective agent adopts its attitudes by responding properly to the set of reasons it has access to, the resulting attitudes will avoid the threat of incoherence associated with standard judgment aggregation.

If reasons-responsiveness guarantees coherence, it may seem puzzling that judgment aggregation methods can lead to incoherent collective attitudes, despite taking as inputs individual attitudes that are properly responsive to the reasons possessed by the members. We should note, however, that different members of a group may have access to different sets of reasons. And attitudes that are supported by different sets of reasons can be incoherent when taken together (e.g. your evidence may support believing that p , while mine supports suspending judgment). Thus, the individual attitudes of the different members of a group may be jointly incoherent, so that the application of judgment aggregation methods leads to incoherent collective attitudes. By contrast, a single set of reasons available to the group (as a collective agent) will not support incoherent collective attitudes, or at least it will not make it rational for the group to adopt simultaneously incoherent attitudes.⁴

The crucial point is that, on the view of rationality we are exploring here, whether an attitude is rational for a given agent is determined by those reasons the agent has access to, but not by reasons *unavailable* to the agent (see Schroeder 2007; Parfit 2011; Kieseewetter 2017; Lord 2018; González de Prado 2019b). As we have just seen, the reasons accessible to a group at the collective level may differ from the reasons accessible to its members. In this way, it may well be that an attitude that is rational for a member of the group (as an individual agent) is irrational for the group (as a collective agent). For instance, a member of the group can possess evidence that she does not want to share with the other members. As a result, it may happen that such evidence is not available for the group as a collective agent.⁵

What reasons count as available to some group? This will depend on the type of group and its structure. In normal cases, when a reason is available (in the relevant sense) to a group, the group will be in a position to guide its behavior by relying on such a reason, and to appeal properly to that reason in justificatory practices. In turn, the group will become open to challenges and criticism if it does not respond correctly to reasons available to it (that is, the group will be treated as *answerable* to the reasons it has access to).

Arguably, whether a reason is accessible to a given group depends on the epistemic position and competences of the group (Sylvan 2015; Kiesewetter 2017; Lord 2018; González de Prado 2019b). First, a reason is accessible to a (collective or individual) agent only if the consideration constituting it is within the agent's epistemic ken. So, rational doxastic attitudes do not have to respond to evidence constituted by facts the agent has no way of knowing (if you are not in a position to know any fact that constitutes evidence about whether it is raining in Sidney, you may rationally suspend judgment on that issue). Moreover, it is plausible to think that the agent also has to be capable of properly recognizing that the consideration constituting the reason offers support to the relevant attitude (Sylvan 2015; Lord 2018; González de Prado 2019b). For instance, the premises of a sophisticated mathematical deduction may not be accessible to the layperson as reasons to endorse the conclusion of the deduction (it can be perfectly rational for the layperson to suspend her judgment about the truth of that conclusion, even if she knows the premises of the deduction).

On the way of seeing things we will favor here, a reason is available to an (collective or individual) agent only if that agent is in a position to respond competently to the reason – that is, just in case the agent is in a position to respond to the reason in a way that manifests a reliable, virtuous competence to be guided only by good reasons (Sylvan 2015; González de Prado 2019b; also Lord 2018). The manifestation of this type of competence will involve displaying a reliable disposition to be guided by, and only by actual reasons (see Sylvan 2015). In other words, the agent has to be capable of manifesting sufficient sensitivity to the relevant reasons in order to count as answerable to them and, more generally, to count as properly evaluable in terms of rationality (Pettit 2001: 283). In this way, entities that are totally insensitive to reasons, such as inanimate objects, are taken to be *arrational*, this to say, beyond rational evaluation. Only agents that are minimally competent in responding to reasons count as rational at all. The question we

want to address now is what it takes for a group to possess sufficient competence as a follower of reasons.

3.4 Group Deliberation

Kallestrup (2016) has offered an account of the epistemic competences of groups in terms of the competences of their members. We could follow suit here and analyze the reasons-responding competences of a group by reference to the rational competences of its members and the way the group is structured. The idea would be that, in order for a group to behave as a rational agent, its members need to manifest their competence in contributing to the group's aim of responding to reasons reliably (Kallestrup 2016: 13; also Silva 2019).⁶ So, as a result of the competences manifested by its members, a rational group will adopt collective attitudes in ways that reveal reliable dispositions to treat certain considerations as reasons only in case they are (and to refrain from doing so when they are not). The individual competences manifested by the members of the group could involve, for instance, a competence in collecting and pooling evidence, knowing how to engage in team reasoning, the ability to trust instruments and inferential methods only when they are reliable enough, or the capacity to defer competently to expert members in their areas of specialization. While the resulting group-level competences may not be shared by any of its members, they arise from the combination of the different individual competences of the members. Accordingly, as we will see below, the reasons accessible to the group as a collective may not be accessible to any of its members.

With this picture of group epistemic competence in mind, we can ask ourselves what decision-making mechanisms should be implemented in a group in order to make it a competent follower of reasons. Hedden (2019) distinguishes explicitly this practical matter from the question of what attitudes are rationally permissible for a group (and he makes it clear that he is interested in addressing this latter question, and not so much the former practical issue). However, these two questions are not completely unrelated, given that the question about what attitudes are rational for a group presupposes that the group has some competence as a follower of reasons, and therefore is evaluable in terms of rationality. If we are not able to show that groups can function in ways that make them minimally competent followers of reasons, it does not make sense to ask ourselves what reasons are available to some group and what attitudes are

rationally supported by those reasons (in the same way that it makes no sense to ask what reasons are accessible to arrational objects such as chairs and tables).

Our goal here is to examine what types of group organizations and dynamics can ground a collective competence to respond to group-level reasons. The answer to this question is not trivial. Arguably, the members of a group will often disagree about what evidential reasons are available to the group, and about what collective attitudes are supported by such reasons. What decision-making mechanisms should be introduced in the group in order to move from such internal disagreements to collective attitudes that are sufficiently sensitive to the group's reasons? If we try to settle these disputes by applying directly standard judgment aggregation methods, we will find again the problems discussed above: the resulting collective attitudes may be incoherent, even if the attitudes of the members are all rational. There is no guarantee that a group exhibiting this type of attitude-formation mechanisms will be in a position to respond to reasons in a competent, reliable way (more specifically, in a way that reliably avoids incoherence). Thus, there is no guarantee that such a group will count as a rational agent with access to reasons.

It may seem that a possible way of integrating judgment aggregation methods with a reasons-based approach to collective rationality is by resorting to judgment aggregation to fix the group's reasons. In this way, the members of the group could vote to select the set of propositions that is to count as the group's available reasons, in other words, as the premises from which the group's attitudes will be derived. In order to avoid incoherence, the candidate sets of reasons would be composed only by rationally independent propositions (i.e. no subset of the propositions sufficiently supports (dis)believing any proposition in the set not included in that subset). The group's rational attitudes would then be those that are sufficiently supported by the selected set of reasons.

This method is a version of the premise-based aggregation procedures explored by List and Pettit (2011), and it would ensure that the resulting attitudes are coherent, if it is granted that an agent cannot rationally derive incoherent attitudes from a coherent set of reasons. However, the implementation of this method is not straightforward. First, in principle, the members of a group may disagree about which sets of propositions are permissible inputs to the aggregation procedure, given that the members can disagree about which propositions are rationally related. Moreover, there can also be disagreements about what follows from the set of reasons acting as

the group's premises. This may happen even if all the members behave rationally, insofar as rational individuals can have different inferential capacities and dispositions (disagreements about inferential support will be far from uncommon in cases involving ampliative, non-monotonic inferences). Now, remember that what attitude is rational for an agent depends on what reasons are accessible to the agent. And what reasons are accessible to the agent depends on the agent's inferential capacities (that is, on what reasons the agent can competently rely on). Yet it is not clear what inferential capacities we should attribute to a group, when its members disagree about what the group's reasons support.

How should these disagreements about the inferential implications of the group's reasons be settled? If we just resort again to judgment aggregation methods, the discursive dilemma and related problems will resurface. To see this, note that rational reasoners are not infallible. Arguably, it is possible to make a rational inference from true premises that nonetheless leads to false conclusions (say, because one resorts to a reliable, but defeasible piece of reasoning). If this is so, there can be a group where 55% of the members rationally take p to follow from the set of group reasons $\{E\}$, 55% of the members rationally think that q follows from $\{E\}$ and nevertheless 80% rationally accept that $\neg(p \& q)$ follows from $\{E\}$. This could happen if 40% infer $p \& q$ from $\{E\}$, 40% infer $p \& q$ and only 15% infer $p \& q$ (imagine that reaching these different conclusions involves manifesting different inferential skills and dispositions, not all of them shared by all members).

A natural reaction to the problems associated with bare judgment aggregation methods is to turn one's attention to deliberative mechanisms of collective decision-making (see Miller 1992; Sunstein 1993; Pettit 2001; Dryzek and List 2003; List 2007). In deliberative processes of collective attitude formation, the members of the group discuss among themselves in order to decide what collective attitude should be adopted. Through such deliberations, the members of the group would try to reach an agreement on what the group's reasons are, which conclusions are supported by them, and in what way these conclusions are so supported. Note that collective deliberation is an interpersonal activity, as opposed to individual deliberation, which only involves the deliberating agent. Again, a crucial difference between groups and individual agents is that only in the case of groups there can be internal collective deliberation among the different agents constituting the group.

An attractive possibility is to resort to judgment aggregation methods like voting only after the attitudes of the members of the group have been shaped by a process of inter-subjective deliberation. The hope is that collective deliberation will significantly reduce the initial disagreement and heterogeneity among the members' attitudes, so that a subsequent application of judgment aggregation methods will tend to deliver coherent collective attitudes. In other words, collective deliberation would foster the sorts of conditions that allow for well-behaved applications of judgment aggregation methods (see Dryzek and List 2003; also Bright, Dang and Heesen 2018).

It could even be expected that, after engaging in collective deliberation, groups will reach sufficient internal consensus, at least if certain conditions obtain. In best-case scenarios, collective deliberation will lead to complete consensus among the group members, not only about what attitude ought to be adopted, but also about what reasons there are for adopting it. Of course, it is trivial to apply judgment aggregation methods after this collective consensus has been achieved. The collective attitude of the group will be the attitude endorsed by all members as the group attitude. Similarly, the group's reasons for adopting that attitude will be those considerations agreed by the members to constitute such reasons. The group will count as responding competently to such reasons if, in reaching this inter-subjective agreement, the members of the group contribute to a group-level disposition to form collective attitudes that track reliably those (and only those) considerations that constitute reasons for those attitudes. Remember that in the previous section we characterized collective epistemic competences as arising from the contributions of the members to group-level reliable dispositions to pursue epistemic aims (in this case, the aim of forming collective attitudes that are supported by reasons).

In simple cases, the group will show proper sensitivity to its reasons because all the members are themselves in a position to recognize that the relevant considerations are reasons for the group's attitude. Note, however, that the type of consensus we are considering does not require that all members adopt, as individuals, the same attitude that is to be adopted collectively by the group, or that they have access as individuals to the group's reasons. The relevant form of agreement may come into place by virtue of there being members of the group that agree to defer to the judgment of other members, or to the outcomes of information-processing mechanisms to which they have no direct access. In this way, there can be members of the group who suspend

judgment on the specific content of the resulting group attitude (and also members who cannot directly grasp the support relation between the group's reasons and its attitudes). Imagine, for instance, cases of distributed cognition, where the members of the group contribute to the collective attitude-forming process by providing different inputs for a complex inferential mechanism, even if perhaps many members have no access to the final outcomes of that mechanism (see Bird 2014; also Silva 2019). It may be that each member is an expert in a different domain and just produces a piece of knowledge pertaining to that domain, which is then passed on to other members, eventually leading to a unified group attitude.⁷ Still, all the members can agree that the group attitude will be determined by the outcomes of this complex inferential procedure. Moreover, by so agreeing, the members could be manifesting their competence as a collective in treating such a procedure as responding reliably (only) to epistemic reasons. As suggested above, a key part of this competence would be knowing when to defer to the expertise of other members, and when to trust the outcomes of the relevant information-processing mechanisms (including, for instance, the outcomes of computer calculations).

It might be argued that full internal consensus is needed for rational collective agency.⁸ At any rate, if collective deliberation manages to bring a sufficient level of homogeneity and agreement among the member's attitudes, it seems that the group will be in a good position to adopt rational, coherent collective attitudes. In the next section we discuss to what extent we can expect collective deliberation to result in sufficient group consensus.

3.5 Consensus and Group Deliberation

One may think that the possibility of reaching full group consensus is extremely far-fetched. Arguably, this is the case for practical collective deliberation about what aims are worth pursuing and what actions ought to be performed. On a common view, evaluative questions are non-factual, in the sense that there is no guarantee that there is evidence about what the facts are that will settle an evaluative debate among fully rational agents.⁹ According to this non-factualist view, evaluative disagreements would be underlain at bottom by conflicts of preferences and values that cannot always be settled rationally (at least not by appealing to facts or pieces of information that all parties to the disagreement can recognize). Thus, even if the members of a group get to share the same body of evidence after deliberating among themselves, they may still have different preferences and disagree about evaluative issues. It can be argued that judgment

aggregation methods such as voting offer the only fair (even if imperfect) way of making practical decisions in groups involving recalcitrant internal evaluative disagreements.

The situation is more promising in relation to epistemic collective deliberation. Epistemic deliberation is generally treated as addressing objective matters of fact. Therefore, epistemic disagreement among rational agents will in principle be solvable by appealing to sufficient evidence about what the facts are. Indeed, according to the view known as *uniqueness*, there is a unique attitude rationally permitted as a response to a given body of evidence (see for instance White 2005; Feldman 2007; Greco and Hedden 2016). On this view, two equally rational agents sharing all their evidence will adopt the same attitude. Insofar as collective deliberation allows for the sharing of information among group members, it should promote group consensus, at least if the members respond rationally to such shared evidence.

What happens if the members of a group take themselves to be rational but find that they disagree with each other (despite appearing to share the same evidence)? Certain approaches to epistemic disagreement offer further support to the idea that collective deliberation should result in group consensus. According to *conciliationist* accounts of peer disagreement, rational agents should reduce their initial confidence in their views when they realize that their epistemic peers disagree with them, despite sharing the same evidence and having considered the issue with care (Christensen 2007; Elga 2007; Feldman 2007; Matheson 2009). This is so because the fact that an agent disagrees with her (equally well-informed) peers should be taken by that agent as a reason to think that she may be mistaken, and therefore to adopt a somewhat cautious attitude toward her initial conclusion. In particular, according to Elga's *Equal Weight* view (Elga 2007), if two agents are initially equally reliable about a certain issue, and they come to endorse incompatible conclusions in the face of the same evidence, they should revise their attitudes assigning equal weight to each other's conclusions (that is, they should regard each other as equally likely to have been right). In turn, if your (equally well-informed) disagreeing peer is taken initially to be more reliable than you about the issue discussed, then the possibility that she is right should be assigned higher probability than the possibility that you are. Thus, rational agents sharing the same evidence (including evidence about their respective reliability), and behaving in accordance with Equal Weight, will (ideally) converge toward the same conciliated consensus attitude after realizing that they disagree about some issue.

Collective deliberation would be a way of ensuring that group members recognize the same relevant evidence (including evidence about the reliability of each member). If the members respond rationally to such evidence, they will (again, ideally and according to uniqueness) come to adopt the same consensus attitude. And, on a conciliationist picture, if they find that they keep endorsing different views, they will allow for the possibility that they are mistaken, and will reduce their initial confidence accordingly. In particular, if Equal Weight is right, it can be expected that the members of the group will tend to converge upon a shared consensus attitude after deliberating.

According to a plausible way of understanding conciliationism, agents engaged in peer disagreement should conciliate because they should revise their original reasoning commitments or inferential dispositions (Brössel and Ede 2014; Rosenkranz and Schulz 2015; González de Prado 2019b). The fact that you and your reliable peers fail to draw the same conclusions from the same body of evidence constitutes a reason to doubt the correctness of the inferential dispositions you were relying on. Thus, if you are epistemically humble, you should revise your inferential dispositions, in a way that acknowledges the possibility that your original inference was mistaken and your peers' was right. For instance, you may assign certain degrees of expected reliability to your and your peers' original inferential dispositions (for details, see Brössel and Ede 2014; Rosenkranz and Schulz 2015; also Schoenfield 2018). Insofar as your peers are also humble and revise their original inferential dispositions in analogous ways, you will end up following similar inferential rules, or at least inferential rules that are closer to each other than the original ones. Thus, by conciliating after collective deliberation, the members of a group would not just move toward agreement on what attitudes should be endorsed, but also toward agreement on the inferential rules that lead from the group evidence to those attitudes.

There are reasons, however, to think that this conclusion is too optimistic, at least beyond very idealized scenarios. First, conciliationism, and in particular Equal Weight are not uncontroversial views (see Kelly 2010; Coates 2012; Titelbaum 2015). And, even if Equal Weight is accepted, there is no guarantee that the members of a group will treat each other as peers, or, more generally, that they will agree on their assessments of the expected reliability of each member (as Klein and Sprenger 2015 remind us, recognition of expertise is often deeply problematic). It should also be noted that one could reject uniqueness and argue that members with different norms and values may rationally adopt different doxastic attitudes, despite sharing

the same evidence and having similar inferential competences. For instance, different members may be willing to undertake different degrees of inductive risk (Magnus 2013).

To make things more complicated, members of a group may be motivated by practical, non-epistemic reasons in order to steer the group toward a given collective attitude, even if they do not take that attitude to be supported by the available evidence. An extreme, but not rare, example of this is Sunstein's figure of polarization entrepreneurs or professional polarizers (Sunstein 2000: 97). In general, collective deliberation among the members of a group can be affected by practical interests in a way that individual epistemic deliberation is not (see González de Prado and Zamora-Bonilla forthcoming). For instance, the members of a group may have practical reasons not to share some of their evidence with the rest of the group,¹⁰ or to argue or vote for a view they know not to be supported by the group's evidence. It is important to note that members that manipulate group deliberations in these ways can behave rationally (as individuals), insofar as they act in response to their practical reasons. It may be perfectly rational to believe (as an individual) that p , but to intend that a group you belong to adopts a collective belief in $\neg p$ (say, because this will have positive practical consequences).

Thus, practical interests and ethical considerations can play an important role in processes of collective deliberation by setting the hidden agenda of the members that participate in such deliberations. Take as an example a company board where some members hide or manipulate information moved mainly by economic interests (see Hendriks, Dryzek, and Hunold 2007). Of course, it will often be the case that the members of a group try to convince each other purely on epistemic grounds. Still, there can be other cases in which non-epistemic factors exert a decisive influence in collective deliberation and group attitude-formation processes.

Moreover, phenomena such as group polarization and the endowment-effect can make one question how realistic the prospects of reaching (epistemically virtuous) consensus via collective deliberation are. It has been claimed that processes of group formation in our current societies can be subject to factors that facilitate and stimulate the agglutination of like-biased individuals (Shapiro 1999; Sunstein 2002). According to well-established results, deliberation in ideologically biased groups may result in severe forms of polarization, so that agents will tend to align their views with the most extreme discussants whose biases are similar to their own (Hafer and Landa 2006; see Blumenthal 2012 for a study of the endowment-effect in group deliberation). As a result of such polarization, after deliberating, agents will reinforce their

original attitudes, rather than moderate them in the face of disagreeing stances. This is just the contrary of what one would expect to happen according to a conciliationist picture of rational deliberation. The use of Bayesian simulation models suggests that this polarization phenomenon does not always result from shortcomings in the rationality of group members, but can be a predictable consequence of deliberation among ideally rational agents (Olsson 2013 and this volume).

Still, we should not undervalue the epistemic benefits of resorting to deliberation when trying to solve internal group disagreements, no matter whether consensus is effectively reached. Aikin and Clanton (2010: 410) have argued that the flow of information during the deliberation process generally improves the epistemic position of the discussants in relation to the evidence. A non-biased argumentative exchange also tends to enhance the quality of the basis on which the different positions are shaped. Of course, the success of deliberation processes strongly depends on the fact that the individuals participating in the deliberation exhibit certain deliberative virtues such as honesty, sincerity, temperance, empathy, or truth-aiming (see Aikin and Clanton 2010). If discussants do not exhibit these virtues, or fall into vicious attitudes, it is clear that deliberation will not produce desirable effects.

The epistemic virtues of deliberative consensus have been highlighted as well with respect to truth-conductivity. Using Bayesian models, Hartmann and Rad (2018) show that, under not particularly demanding conditions regarding the reliability of the discussants (including their reliability in assessing other discussants' reliability), reaching a consensus via deliberation is a better truth-tracking strategy than majority voting. In addition, Hartmann and Rad argue that deliberative consensus has the further benefit of ensuring the satisfaction of group members, in contrast to voting procedures, which may leave those members who voted for the losing option unsatisfied.

3.6 Consensus and Dissensus in Groups

Regardless of how easy it is to promote group consensus, we can ask ourselves whether it is always desirable to do so. It may be argued that dissent plays a valuable epistemic role, which risks being overshadowed by an excessive emphasis on group consensus. Take the example of science. Although a great part of the classical philosophy of science was pro-consensus, some later authors, with Feyerabend at the head, have stressed the importance of dissent in scientific

debates.¹¹ It is often argued that, at least in some fields, it is convenient, if not necessary, to work with a diversity of scientific models in order to produce different kinds of predictions, all of which may be relevant. In this way, model pluralism is customarily accepted in climate modeling. Climate scientists often take into consideration different simulation models when investigating climate change, even if such models make incompatible assumptions (Parker 2006). Similarly, it is not uncommon in the history of science to find cases of prominent scientists working in competing research programs at the same time. For instance, Kragh (1999: 199–200) mentions the case of Heisenberg and Dirac, who, during the crisis of quantum electrodynamics in the 1930s, combined a revolutionary and a conservative strategy: at the same time that they were developing new theoretical hypotheses, they were also introducing small corrections in the old theory. This combination of attitudes, at least in certain cases, proved to be a productive and stimulating path for scientific advance.

One of the risks involved in group deliberation is the apparition of groupthink. Groups affected by groupthink tend to rush to judgment or to accept, in a fast and unreflective way, the opinion manifested by the majority of the group. Unfortunately, mental laziness is often a vice that lies behind many cases of agreement. In these situations, dissent is generally excluded acritically, and alternative views are often undervalued. Indeed, taking consensus to be in itself an ultimate aim of deliberation may have the negative effect of fostering a sort of non-rational conformism (Mackie 2006: 285; also Friberg-Fernros and Schaffer 2014; Landemore and Page 2015). Thus, deliberative agreement should not be reached exclusively in response to practical reasons, such as shortening the time of deliberation. Friberg-Fernros and Schaffer (2014) have discussed further epistemic shortcomings of prioritizing consensus for its own sake: for example, the fact that, after agreement, the discussants often cease to develop new arguments, whose consideration could alter that agreement, or tend to forget the already discarded theses.

In order to avoid groupthink, Solomon (2001, 2006) has favored aggregation procedures over group deliberation, particularly in the case of science. But, precisely by examining in detail the way in which scientific discussions are held, authors like Tollefsen (2006), Wylie (2006), and Wray (2014) adopt a more optimistic attitude toward group deliberation. According this optimistic outlook, the dangers of groupthink can be mitigated if every group member behaves as a critical deliberant that carefully examines, under the scrutinizing gaze of the other members, every reasonable option and contrasts it against the background of available evidence (see Janis

1972; Tollefsen 2006; Wray 2014). Tollefsen (2006: 45) adds that dissent is tolerable in group deliberations as long as it is not pervasive and it does not threaten the stability of the group by questioning its most important principles and norms.

Eschewing a simplified opposition between consensus and dissensus, Beatty and Moore (2010) emphasize the importance of dissent in reaching robust forms of consensus beyond mere aggregation. By highlighting the role of dissenting voices, they follow Elster's (1986) considerations about the democratic force of minorities. Elster claims to have more confidence in the outcome of a decision if a minority voted against it than if it was a unanimous decision (Elster 1986/1997: 16, quoted in Beatty and Moore 2010: 198). As strange as it may initially sound, this idea becomes plausible if we take into account that the mere presence of dissenting opinions in a debate is in principle a guarantee of more careful decision processes in which at least more than a single option was considered. Accordingly, Beatty and Moore (2010: 209) vindicate a qualified way of understanding consensus, in which ideas or decisions are accepted after a virtuous process of deliberation with dissenting parties.

Dissent, however, is not always acceptable or epistemically beneficial. When based on unreasonable grounds or against the available evidence, dissent should be avoided or overcome. Actually, it could be argued that an excessive tolerance of dissent will encourage epistemic conformism, insofar as individuals will not try to assess critically the dissenting views of others.¹² Even in contexts where a certain level of dissent within a group is admitted, deliberative process should aim to remove disagreements that have as their source incomplete or mistaken evaluations of the evidence available. This claim is compatible with granting that the value of consensus over disagreement depends on the task deliberators set out to perform (Landmore and Page 2015: 246). In this respect, we should distinguish groups that constitute collective agents and groups that do not. Arguably, the need to reach agreements will be more pressing in the former than in the latter, since a collective agent needs to adopt, at some point, unitary cohesive attitudes, under penalty of dissolving or becoming paralyzed by internal conflicts. This can be clearly seen in the case of political parties, where internal dissenting voices are badly tolerated and tend to end up being obliterated by a uniform view coalescing around the positions endorsed by the leaders of the party.

Thus, whereas dissent may be epistemically fruitful at different stages of collective deliberation, groups constituting collective agents will need to reach eventually some agreement

on what group attitude to adopt, in the face of the total evidence available (see Wray 2014). Accordingly, rational collective agency will often involve deciding in advance the optimal way to respond to internal disagreement. It should be noted, however, that the final agreement reached by the members of the group may reflect the existing internal dissent in the group. The members may consensually agree that there are dissenting voices in the group, so that no strong view can be collectively endorsed. Rather, the group as a collective will adopt a cautious attitude that properly recognizes the plurality of stances within the group.

We cannot expect groups to be always in a position to reach full internal consensus about a certain topic. Sometimes, the best the members of the group can do is to agree that they disagree, and that therefore the group's collective attitude cannot settle conclusively the issue under discussion (for instance, the group may agree to suspend judgment until more substantial consensus is reached). In particular, this will happen when the time available for deliberation is limited and it is clear that agreement is not going to be easily reached. If, in these situations, a practical decision has to be made anyway, the group can resort to bare judgment aggregation methods as helpful heuristic tools, despite their shortcomings – in the same way that individual agents may resort to heuristics and rules of thumb when making quick decisions. At any rate, in order for the group's use of bare judgment aggregation methods to be rational, it has to be subject to top-down deliberative supervision, so that possible incoherences are suitably revised (see Buchak and Pettit 2015).

3.7 Conclusions

As we have seen, engaging in collective deliberation does not always guarantee that the group as an agent will respond rationally to the evidence available, among other reasons because group members may be rationally motivated by practical, non-epistemic considerations when participating in collective deliberation. However, group deliberation in collective agents will tend to facilitate the achievement of internal agreement, not only about what attitude to adopt collectively but also about the reasons for doing so. If consensus among group members is not reached at a first stage, further deliberation about their assumptions, inferential dispositions, evidence, and epistemic values may help the members assess the nature of the underlying disagreement, so that the collective attitude adopted collectively is properly sensitive to the epistemic position of the group (including its uncertainty). Even in cases in which a consensus

has been easily reached from the beginning, deliberation in further stages may reveal new arguments or aspects of the evidence that had been overseen, making the group more responsible and competent in responding to the available reasons.

Bibliography

- Aikin, S.F., & Clanton, J.C. (2010). Developing group-deliberative virtues. *Journal of Applied Philosophy*, 27(4), 409–424.
- Beatty, J., & Moore, A. (2010). Should we aim for consensus? *Episteme*, 7(3), 198–214.
- Bird, A. (2014). When is there a group that knows? Scientific knowledge as social knowledge. In J. Lackey (ed.), *Essays in collective epistemology*. Oxford: Oxford University Press, 42–63.
- Blackburn, S. (1998). *Ruling passions: A theory of practical reason*. Oxford: Oxford University Press.
- Blumenthal, J.A. (2012). Group deliberation and the endowment effect: An experimental study. *Houston Law Review*, 50(1), 41–71.
- Broome, J. (2007a). Wide or narrow scope? *Mind*, 116(462): 359–370.
- Broome, J. (2007b). Does rationality consist in responding correctly to reasons? *Journal of Moral Philosophy*, 4(3), 349–374.
- Brössel, P., & Eder, A.M.A. (2014). How to resolve doxastic disagreement. *Synthese*, 191(11), 2359–2381.
- Buchak, L., & Pettit, P. (2015). Reasons and rationality: The case of group agents.’ In I. Hirose and A. Reisner (eds.), *Weighing and reasoning: Themes from the philosophy of John Broome*. New York: Oxford University Press, 207–231.
- Christensen, D. (2007). Does Murphy’s Law apply in epistemology? Self-doubt and rational ideals. *Oxford Studies in Epistemology*, 2, 3–31.
- Coates, A. (2012). Rational epistemic akrasia. *American Philosophical Quarterly*, 49(2), 113–124.
- Dancy, J. (2004). *Ethics without principles*. Oxford: Oxford University Press.
- Dryzek, J., & List, C. (2003). Social choice theory and deliberative democracy: A reconciliation. *British Journal of Political Science*, 33(1), 1–28.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41(3), 478–502.

- Elster, J. (1986/1997). The market and the forum: Three varieties of political theory. In J. Bohman and W. Rehg (eds.), *Deliberative democracy*. Cambridge, MA: MIT Press, 3–34.
- Epstein, B. (2015). *The ant trap. Rebuilding the foundations of the social sciences*. Oxford: Oxford University Press.
- Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (ed.), *Philosophers without gods*. Oxford: Oxford University Press, 194–214.
- Friberg-Fernros, H., & Schaffer, J.K. (2014). The consensus paradox: Does deliberative agreement impede rational discourse? *Political Studies*, 62(S1), 99–116.
- Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge, MA: Harvard University Press.
- González de Prado, J. (2019a). No reasons to believe the false. *Pacific Philosophical Quarterly*. <https://doi.org/10.1111/papq.12271>
- González de Prado, J. (2019b). Dispossessing defeat. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12593>
- González de Prado, J., & Zamora-Bonilla, J. (2015). Collective actors without collective minds: An inferentialist approach. *Philosophy of the Social Sciences*, 45(1), 3–25.
- González de Prado, J., & Zamora-Bonilla, J. (forthcoming). Rational golems: Collective agents as players in the reasoning game. In L. Koreň, H.B. Schmid, P. Stovall, and L. Townsend (eds.), *Groups, norms and practices. Essays on inferentialism and collective intentionality*. Springer.
- Greco, D., & Hedden, B. (2016). Uniqueness and metaepistemology. *The Journal of Philosophy*, 113(8), 365–395.
- Hafer, C., & Landa, D. (2006). Deliberation and social polarization. *SSRN Electronic Journal*. <https://ssrn.com/abstract=887634>
- Hartmann, S., & Rad, S.R. (2018). Voting, deliberation and truth. *Synthese*, 195, 1273–1293.
- Hedden, B. (2019). Reasons, coherence, and group rationality. *Philosophy and Phenomenological Research*, 99(3), 581–604.
- Hendriks, C.M., Dryzek, J.S., & Hunold, C. (2007). Turning up the heat: Partisanship in deliberative innovation. *Political Studies*, 55(06), 362–383.
- Janis, I. (1972). *Victims of groupthink*. Boston, MA: Houghton Mifflin.
- Kallestrup, J. (2016). Group virtue epistemology. *Synthese*. <https://doi.org/10.1007/s11229-016-1225-7>

- Kelly, T. (2010). Peer disagreement and higher-order evidence. In R. Feldman and T.A. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press, 111–174.
- Kelly, T. (2014). Evidence can be permissive. In M. Steup, J. Turri, and E. Sosa (eds.), *Contemporary debates in epistemology*. Oxford: Wiley-Blackwell, 298–311.
- Kiesewetter, B. (2017). *The normativity of rationality*. Oxford: Oxford University Press.
- Klein & Sprenger (2015). Modelling individual expertise in group judgments. *Economics and Philosophy*, 31(2015), 3–25.
- Kolodny, N. (2007). State or process requirements? *Mind*, 116(462), 371–385.
- Kragh, H. (1999). *Quantum generations. A history of physics in the twentieth century*. Princeton, NJ: Princeton University Press.
- Landmore, H., & Page, S.E. (2015). Deliberation and disagreement: Problem solving, prediction, and positive dissensus. *Politics, Philosophy & Economics*, 14(3), 229–254.
- List, C. (2007). Deliberation and agreement. In S.W. Rosenberg (ed.), *Deliberation, participation and democracy: Can the people govern?* New York: Palgrave, 64–81.
- List, C., & Pettit, P. (2011). *Group agency*. Oxford: Oxford University Press.
- Longino, H. (2002). *The fate of knowledge*. Princeton, NJ: Princeton University Press.
- Lord, E. (2018). *The importance of being rational*. Oxford: Oxford University Press.
- Mackie, G. (2006). Does democratic deliberation change minds? *Politics, Philosophy & Economics*, 5(3), 279–303.
- Magnus, P.D. (2013). What scientists know is not a function of what scientists know. *Philosophy of Science*, 80(5), 840–849.
- Matheson, J. (2009). Conciliatory views of disagreement and higher-order evidence. *Episteme*, 6(3), 269–279.
- Miller, D. (1992). Deliberative democracy and social choice. *Political Studies*, 40(Special Issue), 54–67.
- Olsson, E.J. (2013). A Bayesian Simulation Model of Group Deliberation and polarization. In F. Zenke (ed.), *Bayesian argumentation, Synthese Library* New York: Springer, 113-134.
- Olsson, E.J. (2020). This volume. Parfit, D. (2011). *On what matters*. Oxford: Oxford University Press.
- Parker, W.S. (2006). Understanding pluralism in climate modeling. *Foundations of Science*, 11(4), 349-368.

- Pettit, P. (2001). Deliberative democracy and the discursive dilemma. *Philosophical Issues*, 11, 268–299.
- Pollock, J.L. (1987). Defeasible reasoning. *Cognitive science*, 11(4), 481–518.
- Rollin, K. (2020). This volume.
- Rosenkranz, S., & Schulz, M. (2015). Peer disagreement: A call for the revision of prior probabilities. *Dialectica*, 69(4), 551–586.
- Schoenfield, M. (2014). Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, 48, 198–218.
- Schoenfield, M. (2018). An accuracy based approach to higher order evidence. *Philosophy and Phenomenological Research*, 96(3), 690–715.
- Schroeder, M. (2007). *Slaves of the passions*. Oxford: Oxford University Press.
- Shapiro, I. (1999). Enough of deliberation: Politics is about interests and power. In S. Macedo (ed.), *Deliberative politics*. Oxford: Oxford University Press, 28–38.
- Silva, P. (2019). Justified group belief is evidentially responsible group belief. *Episteme*, 16(3), 262–281..
- Solomon, M. (2001). *Social empiricism*. Cambridge, MA: MIT Press.
- Solomon, M. (2006). Groupthink vs. the wisdom of the crowds: The social epistemology of deliberation and dissent. *Southern Journal of Philosophy*, 44, 28–42.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge*. Oxford: Oxford University Press.
- 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.
- Stasser, G., & Titus, W. (2003). Hidden profiles: A brief history. *Psychological Inquiry*, 14(3/4), 304–313.
- Stevenson, C. (1944). *Ethics and language*. New Haven, CT: Yale University Press.
- Sunstein, C.R. (1993). *The partial constitution*. Cambridge, MA: Harvard University Press.
- Sunstein, C.R. (2000). Deliberative trouble? Why groups go to extremes. *The Yale Law Journal*, 110(1), 71–119.
- Sunstein, C.R. (2002). The law of polarization. *The Journal of Political Philosophy*, 10(2): 175–195.

- Sylvan, K. (2015). What apparent reasons appear to be. *Philosophical Studies*, 172(3), 587–606.
- Titelbaum, M. (2015). Rationality's fixed point (or: In defense of right reason). *Oxford Studies in Epistemology*, 5, 253–294.
- Titelbaum, M. (2019). Reason without reasons for. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics (Vol. 14)*. Oxford: Oxford University Press, 189–215..
- Tollefsen, D. (2006). Group deliberation, social cohesion and scientific teamwork: Is there room for dissent? *Episteme*, 3(2), 37–51.
- Tollefsen, D. (2015). *Groups and agents*. Cambridge: Polity.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19(1), 445–459.
- Whiting, D. (2014). Keep things in perspective: Reasons, rationality and the a priori. *Journal of Ethics & Social Philosophy*, 8(1), 1–22.
- Worsnip, A. (2018). The conflict of evidence and coherence. *Philosophy and Phenomenological Research*, 96(1), 3–44.
- Wray, K.B. (2014). Collaborative research, deliberation, and innovation. *Episteme*, 11(3), 291–303.
- Wylie, A. (2006). Socially naturalized norms of epistemic rationality: Aggregation and deliberation. *Southern Journal of Philosophy*, 44, 43–48.

Notes

¹ Note that, in general, individuals may have different attitudes when acting as members of a group (say, when engaging in the group's decision-making) and when acting as individual agents in pursuit of their own goals.

² Of course, an agent may disagree with her past views. A rational agent may also hold views that are disagreeing in the sense that they are inconsistent. In particular, in non-transparent situations a rational agent may be unaware that she is actually believing inconsistent contents (think of Kripke's belief puzzles). Likewise, an agent may believe inconsistent contents when appreciating their inconsistency requires sophisticated logical skills (although, arguably, in these cases she will not be perfectly rational). However, it is harder to find cases where an individual agent believes simultaneously contents that she recognizes as inconsistent.

³ This last point is argued for, among others, by Kolodny (2007), Kiesewetter (2017), and Lord (2018). For dissenting views, see for instance Broome (2007a, 2007b) and Worsnip (2018).

⁴ The last part of this sentence is intended to leave room for Permissivism, that is the view that a set of reasons may support either of a set of incompatible attitudes (Kelly 2014; Schoenfield 2014). We will express our sympathies for Uniqueness, the negation of Permissivism (see White 2005; Feldman 2007), but for the time being, we just need the claim that a coherent set of reasons cannot make it rational for an agent to adopt at the same time two attitudes that are incoherent. This is something compatible with Permissivism.

⁵ As we will see below, the evidence available to the group may but in principle need not be evidence available to all group members.

⁶ Kallestrup's view relies on Sosa's AAA account of performance assessment (Sosa 2007).

⁷ Maybe the last step is to introduce all the information gathered in a computer, which then produces automatically the result that is to become the group's attitude. In this case, it could be that all members ignore the content of the final attitude adopted collectively by the group, even if they agree that it will be whatever result is delivered by the computer.

⁸ Wray (2014) defends this view in relation to scientific co-authorship. For a critical discussion of this idea, see Solomon (2006). Bright, Dang, and Heesen (2018) distinguish between collective belief and collective assertion, in particular, the type of (rational) collective assertion involved in co-authored scientific papers. They argue that consensus is not necessary for the latter.

⁹ In general, this non-factualism or non-cognitivism has been defended for evaluative or normative discourse. The position has been defended in moral philosophy and metaethics as well as in aesthetics. Examples of non-cognitivism about morality are Blackburn's (1998) metaethical quasi-realism, Gibbard's (1990) norm-expressivism, or Stevenson's (1944) moral emotivism.

¹⁰ In relation to this, we have the well-known phenomenon of hidden profiles (see, for example, Stasser and Titus 1985, 2003), where, in addition to a common body of information shared by everyone in the group, some of the members possess further unshared pieces of information. This often leads to suboptimal group decisions, especially in groups manifesting shared information bias, in which deliberations tend to revolve around information already shared by the members of the group.

¹¹ Other notable proponents of this view include Longino (2002) and Solomon (2001, 2006).

¹² See Rolin (this volume) for an extended discussion of this issue.