REASONS, COHERENCE, AND GROUP RATIONALITY

Brian Hedden

University of Sydney

Abstract

If groups can have beliefs and other attitudes of their own, what determines which such attitudes the group rationally ought to have? A widespread presupposition is that group-level beliefs should be a function of the beliefs of the group's members, and similarly for other attitudes. But a host of impossibility theorems show that no such aggregation function can satisfy intuitively attractive constraints while ensuring coherent group-level attitudes. I argue that this presupposition is false. Group-level attitudes should be a function of group-level reasons (evidence, in the epistemic case), not individual-level attitudes. This allows for a theory of group rationality that (i) bypasses a host of pessimistic results in the literature on judgment aggregation and (ii) treats rational individual-level attitudes and rational group-level attitudes in parallel.

1 Introduction

When a group displays the kind of systematic and purposive behavior that is characteristic of agency, it is often natural to talk about the group itself as having beliefs (or credences), preferences, and intentions. Moreover, it is often desireable that such a group's attitudes be rational. But problems of group rationality can seem intractable in light of theoretical results

in the theory of judgment aggregation. Following the seminal work of Arrow (1963/1951), many such results have the following form: An *aggregation* function is a function that takes a set of individual attitudes as input and yields a group-level attitude as output. Such a function, we might think, should satisfy certain conditions, such as the condition that there is no one individual such that the group-level attitude always matches that individual's attitude, the condition that if individuals unanimously hold some attitude, then the group has that attitude as well, and so on. Then, a theorem shows that an aggregation function cannot satisfy all those conditions without yielding group-level attitudes that are incoherent.

A central presupposition in all work on judgment aggregation is that what beliefs a group ought to have is a function of what beliefs the individual members have, and similarly for credences, preferences, and other attitudes. In my view, this presupposition is false. Instead, the attitudes a group ought to have are a function of the group's reasons, rather than a function of individual members' attitudes.[1] In the epistemic case, this means that group-level beliefs should depend not on individual-level beliefs, but on group-level evidence.

My approach to group rationality has its roots in work on individual rationality. A familiar epistemological theory is *evidentialism*, on which what doxastic attitudes one ought to have is determined (to the extent that it is) by one's evidence. A generalization of this view beyond just the epistemic case is the view that what attitudes one ought to have is determined by one's reasons. Plausibly, if such principles of rationality hold for individuals, then they hold for groups as well. Just as an individual ought to have the beliefs and credences that are best supported by the reasons that individual possesses, so a group ought to have the beliefs and credences that are best supported by the reasons that the group possesses, and similarly for other attitudes.[2]

[1] I was inspired in this approach by Buchak and Pettit (2015), who argue that problems of judgment aggregation show that a reasons-based theory of group rationality is impossible and that this casts doubt on whether such a theory is workable even in the individual case. I discuss their argument in Section 3, where I show that the problems they raise stem from ignoring crucial components of the reasons-based theory.

[2] Note that my view is *not* that groups ought to employ a two-step procedure of first sharing or merging evidence and then aggregating members' beliefs; rather a group ought to have the attitudes best supported by the group-level evidence, which may or may not have been arrived at through a process of sharing or merging members' evidence.

One might be suspicious of the very notion of group-level evidence (and reasons more generally). However, if talk of group-level beliefs and other attitudes is legitimate (as I am presupposing in this paper), it is not clear why this would not also be the case with group-level evidence. Moreover, it is worth noting that even the notion of individual-level evidence is not crystal clear. Certainly there is no consensus on what constitutes an individual's evidence. But plausibly, the correct theory of individual-level evidence, whatever it is, can be used to generate an analogous theory of group-level evidence.[3] I consider a few possible conceptions of group-level evidence below.

Importantly, adopting a view of group rationality centered on group-level reasons enables us to bypass pessimistic results on judgment aggregation. Such a view emphasizes what we might call *requirements of reason*, as distinct from *requirements of coherence*.[4] Requirements of coherence govern how one's attitudes hang together and include the requirement that one's beliefs be logically consistent, that one's preferences be transitive, that one's credences be updated by conditionalization, and so on. Requirements of reason, by constrast, govern how one's attitudes should fit, not with each other, but with the reasons for or against those attitudes. Candidate requirements of reason include the requirement that one's beliefs be proportioned to (or supported by) one's evidence, that one desire only what one has sufficient reason to desire, and so on.

As noted, results on judgment aggregation show that it is impossible for a function from individual attitudes to group attitudes to satisfy certain intuitively plausible constraints without yielding group attitudes that violate requirements of coherence. But importantly, we will see that if requirements of reason are sufficiently strict, then by satisfying requirements of reason, one will thereby also satisfy requirements of coherence as a byproduct (Kolodny 2007).

[3] A notable exception is a phenomenal conception of evidence, on which an individual's evidence consists of facts about that individual's phenomenal experience. Assuming that groups necessarily lack phenomenal consciousness, a phenomenal conception of evidence would be unsuitable for groups. But see Schwitzgebel (2015) for an argument that groups can have phenomenal consciousness, and List (forthcoming) for a rebuttal.

[4] For further discussion of the distinction between requirements of reason and requirements of coherence, as its importance, see Broome 1999, 2007; Kolodny 2005, 2007, 2008. We might also put the distinction as one between *formal* and *substantive* requirements of rationality (assuming that requirements of coherence are formally characterizable).

This point holds for both individuals and groups. Thus, on a theory of group rationality centered on group-level reasons, if a group has the attitudes that are supported by its group-level reasons, then those group-level attitudes will *ipso facto* be coherent.

The paper proceeds as follows. Section 2 lays out the structure of a reasons-based theory of individual rationality. Section 3 defends the possibility of a reasons-based theory of rational group beliefs, and Section 4 does the same for rational group credences. Section 5 addresses concerns about whether such a theory can provide sufficiently usable guidance. I argue that we should sharply distinguish between the question of what beliefs or credences a group ought to have and the question of what practical mechanisms (e.g, voting rules) should be adopted as means to the end of arriving at rational group attitudes. My view addresses the first question but leaves open the second. Section 6 briefly considers rational group preferences, conceding that a reasons-based theory of rational group preferences is less plausible than a reasons-based theory of rational group beliefs and credences. Section 7 is a conclusion.

2 Reasons-Based Rationality

If there are requirements of reason (or *substantive* requirements of rationality), and if those requirements are sufficiently strict, then satisfying those requirements of reason is sufficient for satisfying requirements of coherence (Kolodny 2007). Kolodny's central example involves contradictory beliefs. Consider a requirement of coherence that says that one ought not simultaneously believe H and believe $\neg H$. Now suppose that there are requirements of reason. In the epistemic case, it is natural to identify one's (epistemic) reasons with one's evidence. The relevant requirement of reason, then, is that one believe a proposition if and only if it is supported by one's evidence. Suppose further than it is impossible for one's evidence both to support H (to the degree sufficient to warrant belief) and also to support $\neg H$. It follows that if one has violated the requirement of coherence that forbids contradictory beliefs, one must have violated the above requirement of reason. Hence, if one satisfies that requirement of reason, one will *ipso facto* satisfy the requirement of coherence prohibiting contradictory

beliefs.

Assuming that an analogous story can be told for all other requirements of coherence, we get the result that those requirements of coherence are superfluous, at least in characterizing what attitudes one ought to have. Kolodny goes further, arguing that requirements of coherence are superfluous *tout court*. This is a striking claim. One might think that requirements of coherence, even if not required to accurately characterize the optimal case in which one believes as one ought, are still required to accurately rank suboptimal cases. After all, it seems worse to believe H and also $\neg H$ than to believe only H when in fact one's evidence supports $\neg H$. I will not attempt to evaluate Kolodny's defense of the stronger claim that requirements of coherence are wholly superfluous and hence eliminable. For my purposes, the weaker claim is sufficient. Let a *reasons-based* theory of rationality be a theory on which (i) there are requirements of reason, and (ii) if one satisfies all requirements of reason then one thereby satisfies all requirements of coherence.[5] Of course, fully fleshing out such a theory would require giving a theory of reason possession and reasons support; that is, giving a detailed characterization of what it takes to have a given reason, and also of what it takes for a given set of reasons to support a given attitude. I won't be giving such a fully fleshed out theory in this paper, either for the case of individuals or for the case of groups. My concern is with the structural features at the heart of the reasons-based approach to rationality embodied in (i) and (ii).

Any such reasons-based theory of rationality is committed to the following two key theses:

> **Factivity**: A proposition is a reason only if it is true.
>
> **Uniqueness**: For every set of reasons and for every proposition, there is only one attitude of a given type (e.g., belief-like, desire-like, etc) toward that proposition which is supported by that set of reasons.

[5]This also means that I do not need to attempt to rebut standard arguments in favor of requirements of coherence, such as Dutch Book (or money pump) arguments, accuracy arguments (Joyce 1998, Greaves and Wallace 2006), or representation theorem arguments. As I read them, these arguments support only the claim that satisfying the relevant requirements of coherence is a necessary condition for rationality, not that it is a sufficient condition. Hence regarding these arguments as sound is compatible with the reasons-based theory I endorse.

If either of these claims is false, a reasons-based theory is indefensible. First, suppose Factivity is false. Then, one might simultaneously have E and also $\neg E$ as part of one's evidence. Then, belief in E and belief in $\neg E$ would both be supported by one's evidence (assuming that if a proposition is entailed by one's evidence, belief in that proposition is supported by one's evidence). Hence satisfying requirements of reason would be compatible with violating requirements of coherence, *contra* the reasons-based theory.

Second, suppose Uniqueness is false.[6] Suppose, for instance, that evidential support is always relative to a set of evidential standards, or ways of assigning weights to considerations of simplicity, explanatory power, and the like. And suppose further that there are multiple different rationally permissible sets of evidential standards. Such a view is a version of *Permissivism*. Then, it might be that relative to one permissible set of evidential standards, one's evidence supports H but not $\neg H$, while relative to a different permissible set of evidential standards, it supports $\neg H$ but not H. In such a case, if one believes H and also $\neg H$, neither of one's beliefs would be in violation of a requirement of reason. Each would be supported by one's evidence, relative to a permissible set of evidential standards. (One could avoiding commitment to Uniqueness by adopting a requirement that prohibits using one set of evidential standards to with respect to H and another with respect to $\neg H$, but this would constitute a requirement of coherence and thus amount to an abandonment of the reasons-based approach to rationality.)[7]

[6]Kolodny (2007, 248) uses the phrase 'evidence strictness,' which he formulates as: 'Reason either forbids a belief, or it requires it.' This is a statement of Uniqueness, but restricted to the case of binary belief.

[7]One might think that adopting Uniqueness makes the problem of group rationality trivial. If Uniqueness is true, then if all group members are rational, they will have the same prior probabilities. Aumann (1976) argues that if agents have the same prior probabilities and have common knowledge of their posterior probabilities, then those posterior probabilities are the same. And if all group members have the same posterior probabilties, then plausibly those are the probabilities that the group ought to have as well. Three things are worth noting, however. First, even if the group members might not all be rational and do not in fact have the same prior probabilities, it may still be that the group's credences ought to be based on the uniquely rational prior probabilities. Second, group members might not have common knowledge of their posterior probabilities. For one thing, they may not have shared their views in public. For another, it may be that even if they have shared their view in public, this may not suffice for them to achieve common knowledge of their posterior probabilities (cf. Lederman forthcoming). Third, Aumann's result relies on controversial assumptions about knowledge and evidence, in particular positive and negative introspection. Lederman (2015) shows that without these assumptions, there are models in which agents with common priors and common knowledge of their posterior probabilities *do* differ in their posteriors.

Note that it is plausible that if there is a unique permissible set of evidential standards, it is indeterminate which it is (Christensen 2007; Hedden 2015b; Greco and Hedden 2016). Thus, it may be determinately the case that there is a unique rational doxastic state given a body of evidence, but indeterminate what that uniquely rational state is. Still, if one has contradictory beliefs, it may be indeterminate which belief is irrational, but determinately the case that one or the other (or both) is. This Uniqueness-cum-indeterminacy view differs from Permissivism in that it avoids the need for requirements of coherence, and also in that entails that if two individuals employ different sets of epistemic standards, at least one of them must be irrational.

Factivity is widely endorsed. Many theorists of reasons endorse the claim that reasons are facts (Raz 1975; Scanlon 1998; but see Schroeder 2008 for the view that 'subjective reasons' can be false). A false proposition might be one that you believe to be a reason, but it cannot in fact be a reason (though treating that proposition as a reason might be excusable, for instance if one's mistaken belief in its truth is rational). In the epistemic case, where reasons are most naturally construed as evidence, most epistemologists hold that only truths can be part of one's evidence (see e.g., Littlejohn 2013). According to one recently influential theory, one's evidence consists of all and only the propositions that one knows (Williamson 2000), and of course knowledge is factive.

Uniqueness is far more controversial. It has many detractors (Kelly 2013; Schoenfield 2013; Titelbaum and Kopec ms), but also many defenders (White 2005; Greco and Hedden 2016; Horowitz 2013; Dogramaci and Horowitz forthcoming). I will not attempt a defense of Uniqueness here. Rather, I simply note that without it, the reasons-based theory does not get off the ground. My aim is to explore what the reasons-based theory has to say about group rationality.

Finally, while Uniqueness, restricted to doxastic attitudes, has many defenders in epistemology, Uniqueness for practical attitudes such as preferences and intentions is almost universally rejected, including by Kolodny himself (2007, 251), though see Hedden (2015b)

for a tentative defense. This paper is primarily concerned with group epistemology, though I will briefly comment on problems of social choice and aggregation of preferences in Section 6. Henceforth, when I talk about Uniqueness, I will be implicitly restricting it to the type of attitude then under discussion.

3 Group Beliefs and the Discursive Dilemma

We begin with binary beliefs before turning to credences, or probabilistic degrees of belief. The discursive dilemma (Kornhauser and Sager 1986; Pettit 2001) presents a major challenge for the aggregation of individual-level beliefs to arrive at rational group-level beliefs.

To begin, consider a requirement of coherence that says that one ought be such that if one believes P and believes Q, then one does not believe $\neg(P \wedge Q)$.[8] Call this the *conjunction requirement*. Now suppose that all individuals in the group satisfy this requirement, and group beliefs are aggregated by majority rule, so that the group believes a proposition if and only if it is believed by a majority of group members. Unfortunately, under these assumptions a group can fail to satisfy the conjunction requirement.

Let there be three group members. The first believes P, believes $\neg Q$, and hence believes $\neg(P \wedge Q)$; the second believes $\neg P$, believes Q, and hence believes $\neg(P \wedge Q)$; and the third believes P, believes Q, and hence believes $(P \wedge Q)$. Then a majority believes P, a majority believes Q, and a majority believes $\neg(P \wedge Q)$. Therefore, by majority rule, the group believes P, believes Q, and believes $\neg(P \wedge Q)$, violation of the requirement of coherence. This is analogous to the Condorcet paradox whereby pairwise majority voting can yield intransitive group-level preferences even if the individual-level preferences are all transitive (Condorcet 1785).

[8] As Buchak and Pettit (2015, 212) note, one might be skeptical of such a requirement due to the preface and lottery paradoxes. However, as they also observe, this worry can be mitigated by two considerations. First, we can run the paradox with a restricted version of the requirements that applies only when P and Q are atomic, and this is a weaker assumption that what is needed to generate the prefance and lottery paradoxes. Second, 'the discursive dilemma...can be run for modus ponens and other inference rules, in addition to conjunction.'

Further work on the discursive dilemma generalizes the paradox. List and Pettit (2002) give an impossibility theorem stating that no aggregation function can satisfy four plausible conditions (Universal Domain, Collective Rationality, Anonymity, and Systematicity). The details of these conditions and the associated impossibility theorem won't concern us. What is important for present purposes is that Buchak and Pettit (2015) argue that the discursive dilemma and subsequent impossibility results show that the reasons-based theory fails at the group level. That is, they claim that at the group level, it is not the case that if the group satisfies requirements of reason (believing a proposition if and only if it is supported by its evidence) it will thereby satisfy requirements of coherence as a byproduct.

Their argument crucially relies on a certain conception of group evidence. They hold that a group has p as part of its evidence if and only if a majority of group members believe p:

> [J]ust as reasons-sensitivity requires that the beliefs of individuals should be responsive to evidence, so it requires that the beliefs of a group agent should be responsive to the beliefs of its members: they are the evidence-channels that mediate the relationship between the group's attitudes and the world. (213)
>
> Presumably, for such a group to be required by reason to believe p is for the profile of member judgments or opinions to give the group conclusive reason to believe p. Here the profile of individuals' opinions plays the role in the group case that evidence plays in the individual case: group votes in favor of p are evidence for p in the group case (they make it more likely that p), just as perceptual evidence and the like are evidence in the individual case. (214)

But we have already seen that majority voting on each proposition can yield group beliefs in violation of requirements of coherence.

The argument of Buchak and Pettit relies on a conception of evidence that will, and indeed must, be rejected by the defender of a reasons-based theory. For a reasons-based theory requires that reasons, and in particular evidence in the epistemic case, be true. But a majority can *believe* a proposition even if it is false. Hence their conception of group evidence violates Factivity.

At most two of P, Q, and $\neg(P \wedge Q)$ can be true. Hence if evidence is factive, then the group's evidence can include at most two of these propositions. And because any two of these

propositions entails the negation of the third, the group's evidence can only support two of the three propositions. If the group's evidence includes P and Q, say, then the group's evidence does not support $\neg(P \wedge Q)$. If it includes P and $\neg(P \wedge Q)$, then it does not support Q. If it includes Q and $\neg(P \wedge Q)$, then it does not support P. And if it includes only one (or none) of these three propositions, then more needs to be specified in order to determine whether either or both of the other two is supported by the group's evidence, but no matter what, the group's evidence will not support a set of beliefs that violates the conjunction requirement.

Factivity is key, then. Of course, just saying that evidence is factive does not provide a full theory of group evidence. Indeed, as conceded earlier, until we give a fully-developed theory of group-level evidence (and reasons generally), we have not so much a reasons-based theory of group rationality, but a theory schema with a placeholder for one's favored theory of group evidence.

I will not endorse a particular theory of group evidence here, but will instead note a few options. First, the group's evidence may be a function of what individual members' bodies of evidence are. Here are three possibilities:

> **Pooled Evidence**: If at least one individual's evidence includes P, then the group's evidence includes P.
>
> **Majority Evidence**: If a majority of individuals' evidence includes P, then the group's evidence includes P.
>
> **Common-Ground Evidence**: If each individual's evidence includes P, then the group's evidence includes P.[9]

Second, it may be that the group's evidence is not a function of the individuals' evidence. Instead, it may depend both on what evidence the individuals possess (and perhaps their other attitudes) and on how the group is itself structured. Here is one such picture. Suppose that the E=K thesis (that one's evidence consists of all and only the propositions that one knows) holds not only for individuals but also for groups. Thus, the group's evidence consists

[9]The terminology of 'pooled evidence' and 'common-ground evidence' comes from Russell et al (2015). They do not discuss the majority evidence version. See Williamson (2015) for further discussion of functions from individual-level evidence to group-level evidence.

of all and only the propositions that the *group* knows. Suppose also that some broadly functionalist theory of belief is correct, so that what one believes depends in part on how one behaves, or is disposed to behave. How a group behaves, or is disposed to behave, depends not only on its members' attitudes, but also on how it in fact uses them (e.g., whether it uses some particular voting scheme, whether it is hierarchically structured, whether it makes decisions through group deliberation, and so on). Given E=K and a broadly functionalist story about belief, we get the following result: What the group has as evidence depends on what the group knows, which depends in part on what the group believes (since knowledge entails belief), which in turn depends not only on its members' attitudes, but also on how it is structured. Call this a *structural account* of group evidence. Note that a structural account of group evidence need not espouse E=K or functionalism about belief; group evidence might depend on group structure for some other reasons.

I won't take a definitive stand on how we should conceive of a group's evidence, though I will note that I am most sympathetic to a structural account of group evidence. Instead, I want to emphasize two lessons that come from taking a reasons-based theory seriously not only as a view about individual rationality but also about group rationality. First, what a group ought to believe will not be a function of what its individual members believe. At most, it will be a function of what its members' bodies of evidence are (and, if a structural account of group evidence is correct, not even that). This perspective is a marked shift from ordinary presuppositions in the theory of group rationality. Second, if Factivity holds, so that reasons (e.g., evidence propositions) must be true, then a group's evidence will *ipso facto* be consistent, and hence if a group believes all and only what is supported by its evidence, its beliefs will satisfy the conjunction requirement.

I close this section by briefly contrasting this account with a different approach considered by List and Pettit (2011). The discursive dilemma arises only when the propositions being voted on are logically interconnected. One way out, then, is to restrict the set of propositions being voted on to a proper subset of set of propositions under consideration. A group might

adopt a 'premise-based procedure,' in which certain propositions are designated as premises and others as conclusions. The group then arrives at an attitude toward each premise proposition by majority vote or some other rule. Then, the group forms an attitude toward each conclusion proposition based on whether it is entailed (or, perhaps more generally, supported) by the premise propositions it believes. For instance, it might designate P and Q as premises, form an attitude toward these propositions by majority vote, and then form a belief in the conclusion proposition $\neg(P \wedge Q)$ only if either $\neg P$ nor $\neg Q$ was supported by majority vote.

My reasons-based approach shares with their premise-based procedure the feature of enforcing consistent group-level attitudes toward a subset of the propositions under consideration and then determining therefrom what the group-level attitude should be toward remaining propositions. But my approach differs in at least three respects. First, the premise-based procedure enforces consistent group-level attitudes toward a subset of the propositions under consideration by requiring that propositions in that subset have certain logical features such that it is impossible for majority voting to yield inconsistent attitudes toward those propositions. By contrast, the reasons-based theory does so by appeal to Factivity. Majority beliefs in propositions in the premise set might be false, but a group's reasons must be true. Second, on the premise-based procedure, the group-level attitudes toward the premises are a function of the individuals' attitudes toward the premises, whereas the reasons-based approach is compatible with group-level evidence not being a function of individuals' evidence, as on the structural account of group evidence mentioned above. Third, while the premise-based procedure singles out a *proper* subset of the propositions under consideration, the reasons-based approach does not. Certainly, some theories of evidence have it that only certain propositions can be part of anyone's evidence; phenomenal theories of evidence, for instance, have it that only propositions concerning experiences or appearances can be evidence. But other theories of evidence in principle allow any proposition to be part of someone's evidence. For instance, on E=K, any proposition which is in principle knowable can in principle be part of someone's evidence. Thus, the reasons-based approach is compatible with the possibility that for each of

P, Q, and $P \wedge Q$, there is some individual whose evidence includes that proposition, whereas the premise-based approach requires that at most two of these propositions are eligible to be part of anyone's evidence.

.

4 Group Credences and Probability Aggregation

I turn now to group-level credences, or probabilistic degrees of belief. I will not attempt a survey of impossibility theorems in the theory of probability aggregation (see Genest and Zidek 1986; Russell et al 2015; Dietrich and List 2016). Instead, I will simply look at the most obvious and popular probability aggregation rule, note one condition it violates, and then turn to a reasons-based approach to group credences.

Assume that each individual in the group has credences that conform to the axioms of the probability calculus and updates those credences by conditionalization. Is there any adequate aggregation function from a set of individual credence functions to a group credence function?

One aggregation function is (unweighted) linear averaging. With linear averaging, the group credence in any proposition is simply the linear average of the individuals' credences in that proposition. That is, the group credence function, given a set of individual credence functions $\{P_0, P_1, ..., P_n\}$, is the function P_g that assigns to any proposition H the probability $[P_0(H) + P_1(H) + ... + P_n(H)]/n$.

Linear averaging does yield a group credence function that conforms to the axioms of the probability calculus, thus satisfying the synchronic requirement of coherence that one's credences be probabilistically coherent. But it fails to commute with conditionalization, a diachronic requirement of coherence. Conditionalization says that upon learning some new proposition E and nothing stronger, an agent's new credence in any proposition H should equal her old conditional credence in H given E. More precisely, let P_0 be the probability function representing the agent's doxastic state at t_0, and P_1 the probability function representing the agent's doxastic state at t_1. Conditionalization says:

Conditionalization:

It is a requirement of rationality that if between t_0 and t_1 one learns E and nothing stronger, then for all H, $P_1(H) = P_0(H \mid E)$

But linear averaging does not commute with conditionalization.[10] That is, the result of averaging a set of individual credence functions, then conditionalizing the resulting group credence function on E, will not in general equal the result of conditionalizing each individual credence function on E and then averaging the result. To see that linear averaging does not commute with conditionalization, consider the following example. There are two individual credence functions P_1 and P_2, which have the following properties:

$P_1(H \wedge E) = 0.05$
$P_1(E) = 0.25$
$P_2(H \wedge E) = 0.6$
$P_2(E) = 0.75$

If we take the linear average of these two credence functions, the resulting group credence function P_g is such that $P_g(H \mid E) = 0.65$. Thus, averaging the individual credence functions and then conditionalizing the result on E yields a group credence in H of 0.65. But if we first conditionalize the individual credence functions on E and then take the resulting linear average, the post-learning group credence function will be P'_g, where $P'_g(H) = 0.5$.

There are other aggregation functions that do commute with conditionalization, but they also violate some *prima facie* attractive conditions.[11] What does the reasons-based theory say about what credences a group ought to have? Let us begin by looking at what a reasons-based theory says about what credences an individual ought to have. We saw above that the

[10] Weighted linear averaging can commute with conditionalization provided that the weights assigned to individuals' credence functions change in a certain way after the learning event. See Pettigrew (forthcoming).

[11] Note also that even other aggregation functions that do commute with conditionalization, such as (re-normalized) geometric averaging, do so only assuming a fixed set of group members (all of whom update by conditionalization). As group membership changes, there is no guarantee that a group that employs these other rules will update by conditionalization. As we will see, on my theory of rational group credences, a group that always has the credences supported by its evidence will thereby satisfy conditionalization, unless a change in group membership or some other factor results in the group losing evidence. But loss of evidence is a good reason to violate conditionalization, and, assuming loss of evidence need not be irrational, it shows that some violations of conditionalization are rationally permissible.

reasons-based theory is committed to Uniqueness, which in the epistemic case amounts to the claim that given any body of evidence, there is a unique most rational doxastic state to have in response to that evidence. In a Bayesian framework, Uniqueness entails that there is a uniquely rational prior probability function P, and if one has total evidence E, then one's credences ought to equal that prior probability function conditional on E, that is $P(- \mid E)$. It is clear that this means that satisfying requirements of reason entails satisfying the synchronic requirement of probabilistic coherence. Moreover, if one satisfies this requirement of reason at each time, then one will conform to conditionalization as a byproduct (assuming that one never loses evidence). To see this, suppose that at t_1 one has total evidence E_1 and at t_2 one gains evidence E_2, so that one's total evidence is then $E_1 \wedge E_2$. The reasons-based theory says that one's t_1 credences ought to be $P_1(-) = P(- \mid E_1)$ while one's t_2 credences ought to be $P_2(-) = P(- \mid E_1 \wedge E_2)$. But P_2 is the probability function that results from taking P_1 and conditionalizing on E_2. Thus, if one satisfies requirements of reason at each time, then one will satisfy requirements of coherence, in particular conditionalization, as well. Again, it may be that while it is determinately the case that there is a uniquely rational prior, it is indeterminate what it is. Then, if one violates conditionalization, it will determinately be the case that one violated a requirement of reason at some time or other, even if it is indeterminate at which time(s) that is.

It should now be clear what a reasons-based theory will say about what credences a group ought to have. Where P is the uniquely rational prior probability function and E is the total group-level evidence, the group's credences ought to equal $P(- \mid E)$. And if at all times the group satisfies requirements of reason and has the credences it ought to have, it will satisfy requirements of coherence like probabilistic coherence and conditionalization as a byproduct. As before, this is compatible with a range of views about group-level evidence, including some on which the group-level evidence is some function of individuals' evidence, as well as some on which it is not, such as a structural account of group evidence.

If a structural account of group evidence is correct, then what credences the group ought to

have will not be a function of what credences the individuals have. But if group-level evidence is a function of what evidence individuals have, and if moreover individuals assign credence 1 to all and only those propositions that are either part of their evidence or assigned probability 1 by the uniquely rational prior,[12] then what credences the group ought to have will be a function of what credences the individuals have. The reasons-based theory then endorses a version of what Russell et al (2015) call the *Fixed Prior Rule*, which says that the group credences ought to be the result of conditionalizing some fixed prior on a set of propositions, where this set is some function of the individuals' credences (for instance, it could be the set of propositions assigned credence 1 by at least one individual, or by a majority of individuals, or by all individuals, and so on).[13]

The Fixed Prior Rule does obey conditionalization, but it violates other conditions that Russell et al find *prima facie* attractive. I am untroubled by these violations, however. I begin with what might be thought of as the most obvious and undeniable such condition, showing that even it is ill-motivated on a reasons-based approach. I then turn to more controversial conditions and show that they too should be rejected. Let us begin, then, with a seemingly undeniable condition, Unanimity:

> **Unanimity**: If every individual assigns the same credence to a proposition, then the group also assigns that credence to it.

While this condition is certainly intuitive attractive, it lacks a clear motivation on a reasons-based approach to group rationality. To see this, note that two individuals can have the same credence in a proposition, but for different reasons. And once we pay attention to the individuals' reasons for their credences, and the way individuals' reasons relate to the group's reasons, we see good motivations for rejecting Unanimity.[14] Suppose the group's two

[12]This caveat is relevant in the case of uncountable probability spaces, where some propositions must be assigned probability 1 despite their negations being epistemically possible.

[13]Note that the fixed prior need not be interpreted as the uniquely rational one. It could be, say, the average of individuals' priors, which may or may not be rational. But Russell et al do note that it is natural to let the fixed prior 'be some ideal objective prior, if you believe in such a thing. If you want to know how [a group] *should* bet, a natural interpretation of "should" presses in this direction' (1300).

[14]Cf. Lackey (2016) on individuals having different, and even conflicting, grounds for a given belief. Lackey

members have the same priors, and in particular the uniquely rational ones, but they have different bodies of total evidence that nonetheless support the same credence in a proposition. For instance, suppose that the first individual has evidence E_1 and the second has evidence E_2, and moreover $P(H \mid E_1) = P(H \mid E_2) = n$. If we think of the group's evidence as consisting of all evidence that *any* of the individuals has (Pooled Evidence), then the group's evidence is $E_1 \wedge E_2$. And it may be that $P(H \mid E_1) = P(H \mid E_2) \neq P(H \mid E_1 \wedge E_2)$. Alternatively, if we think of the group's evidence as consisting of evidence possessed in common by all the individuals (Common-Ground Evidence), then the group's evidence is $E_1 \vee E_2$, and it may be that $P(H \mid E_1) = P(H \mid E_2) \neq P(H \mid E_1 \vee E_2)$. Of course, if all of the individuals assign the same credence to H *and for the same reason* (i.e. on the basis of the same evidence), then on any plausible view of group evidence on which the group's evidence is a function of the individuals' bodies of evidence, the group will have the same evidence as each of the individuals. Then, by Uniqueness, if the group and all the individuals satisfy requirements of reason, the group will agree with the unanimous judment of the individuals and assign the same credence to H. But we should not require the group to agree with the unanimous credal judgments of individuals when those judgments are based on different reasons.[15]

I now turn to the other conditions that Russell et al advocate and that the Fixed Prior Rule violates:[16]

> **Neutrality**: The way in which the group credence in a proposition depends on individuals' credences is the same for every proposition.
>
> **Irrelevant Alternatives**: The group credence in a proposition depends only on

is considered group doxastic justifcation, however, while my focus is propositional justification. See also Williamson (2015).

[15]Compare the way in which List and Pettit's premise-based procedure, discussed in Section 3, violates a similar unanimity condition. Suppose we classify P, Q, and R as the premise propositions, and $P \wedge Q \wedge R$ as the conclusion proposition. Suppose 1/3 of the population believes P, Q, and $\neg R$, 1/3 believes P, $\neg Q$, and R, and 1/3 believes $\neg P$, Q, and R. Then, a majority believes each of P, Q, and R, and so by the premise-based procedure the group believes the conjunction $P \wedge Q \wedge R$, even though the individuals in the population unanimously believes $\neg(P \wedge Q \wedge R)$. The reason the premise-based procedure overrides this unanimous judgment is because the individuals have this belief regarding the conclusion proposition on the basis of differing beliefs regarding the premise propositions.

[16]See their paper for formal statements of these conditions. They also show that the Fixed Prior Rule violates a condition called *External Bayesianity*, which I do not consider here for reasons of space.

individuals' credences in that proposition, and not on their credences in any other propositions.

Independence Preservation: If propositions A and B are independent relative to each individual's credence function, then they are independent relative to the group credence function as well.

Continuity: The function from sequences of individual credence functions to group credence functions is a continuous function.

Let us take there in turn. Neutrality requires that an aggregation function be blind to the meanings of the propositions in the algebra. If there are finitely many worlds in our probability space, then Neutrality can be achieved by having the prior be uniform, assigning each of the n worlds probability $1/n$. But if there are infinitely many worlds, this is impossible (assuming countable additivity): 'there is no probability function on an infinite σ-algebra which is symmetric under arbitrary measurable permutations' (Russell et al 2015, 1300). But in my view, Neutrality is undermotivated. Neutrality may be a necessary condition for a purely formal characterization of rational group credence, one which cares only about the logical form of propositions but not their content. Reasons-based theorists generally reject the demand for a purely formal characterization of rational individual belief and credence, and with good reason (Goodman 1955; Hedden 2015c). They hold instead that what credence in a proposition is supported by a given body of evidence depends on factors like simplicity, naturalness, and explanatory considerations that presumably cannot be characterized formally. And if there are good reasons to reject purely formal approaches to individual rationality, these reasons plausibly carry over to the case of groups.

The Fixed Prior rule violates Irrelevant Alternatives because it embodies a more holistic approach to group rationality. The group credence in H depends not on individuals' credences in H, but on which propositions they have as their evidence.[17] But again, there is no clear

[17]This can be seen with an example which also demonstrates a violation of Unanimity. Suppose that an object is drawn at random from an urn consisting of black cubes, white cubes, black balls, and white balls in equal proportions. Suppose also we adopt a Pooled Evidence conception of group evidence. In case 1, each of two agents learns that the object drawn is black, and so each of the agents, as well as the group, has 0.5 credence that the object is a black cube. In case 2, one agent learns that the object is black, while the other learns that it is a cube. Hence each individual has credence 0.5 that it is a black cube, but given Pooled Evidence, the group's evidence includes both that the object is black and that it is a cube, and hence ought

motivation for Irrelevant Alternatives, and indeed Russell et al do not give any reason for wanting a 'local,' non-holistic theory of rational group credences.[18]

Independence Preservation is related to Unanimity in that it says that if each individual's credence function has a certain feature, then the group's credence function must also have that feature. There are two things to be said about why Independence Preservation should be rejected. The first is that the same considerations that support rejecting Unanimity also support rejecting Independence Preservation. It may be that all individuals have credence functions which treat H and E as probabilistically independent, but for different reasons, such that on any way of aggregating those reasons, the resulting group-level reasons do not support credences on which they are probabilistically independent.

Second, there are grounds for rejecting Independence Preservation that do not make reference to reasons at all. As Pettigrew (forthcoming) writes,

> It seems to me that the mistake in this objection [to aggregation functions which

to have credence 1 that it is a black cube (violating Unanimity). The two cases are alike with respect to the individuals' credences in the black cube hypothesis, but on the view in question they differ with respect to the credence the group ought to have in that hypothesis, because the two cases differ with respect to what evidence the individuals possess, and hence with respect to what evidence the group possesses, thus violating Irrelevant Alternatives.

[18]Dietrich and List (forthcoming) give no epistemic justification for this principle, which they call 'eventwise independence,' but they give two pragmatic justifications for it: (1) It is easier to implement since it permits a group to just focus on a single proposition (or a few) rather than considering all the propositions on the agenda. (2) It is invulnerable to manipulation of the agenda: 'If the collective opinion about each event depends only on the individual opinions about that event, then an agenda setter who might wish to influence the outcome of the aggregation will not be able to change the collective opinion about any event by adding further events to the agenda or removing others from it' (p. 8).

However, it is important to distinguish epistemic from pragmatic considerations. In this paper, I am concerned with the question of what credences a group epistemically ought to have, not with which voting or other procedures a group should employ in coming to a decision (though we will touch on these issues in Section 5). Pragmatic considerations involving implementability and potential agenda manipulation seem irrelevant to the question of what credences a group rationally ought to have.

Compare an issue concerning Arrow's (1963/1951) Independence of Irrelevant Alternatives condition, which says that the social ranking of A relative to B should depend only on how individuals rank A relative to B, and not on how they rank them vis-à-vis other options. In the context of welfare economics, this serves to bar interpersonal comparisons of utility from determining facts about the social preferences or social good (which is already ruled out by the functional form of social welfare functions but not social welfare functionals; see Sen 1970 for discussion). But in the context of voting theory, where we are concerned with the actual implementation of various voting rules, it is motivated largely by a desire to prevent or mitigate the threat of strategic voting. If we are not skeptical of interpersonal comparisons but are worried about strategic voting, then this condition is undesirable in the context of welfare economics but may be desirable in the context of voting theory.

> violate Independence Preservation] lies in its account of judgments of independence...[I]t is simply not the case that I judge H and E to be independent just in case my credence in H remains unchanged when I condition on E: It is possible to judge that H and E are independent without satisfying this condition; and it is possible to satisfy this condition without judging them independent. (12)

On Pettigrew's view, the judgment that H and E are independent is a propositional attitude whose content concerns the objective chances of H and E (or perhaps their evidential probabilities; the same points apply in either case). Judging that H and E are independent is a matter of believing that the objective chances of H and E are independent. And one can believe that the objective chances of H and E are independent without having a credence function relative to which H and E are independent, and *vice versa*. For instance, if you are uncertain about the bias of a coin, you will judge the events *Heads on first toss* and *Heads on second toss* to have independent objective chances even though they're not independent relative to your credence function (since the first event is evidence for the coin's being biased toward heads, which in turn is evidence for the second event). With these points in mind, we can distinguish between Independence Preservation, which states that a group must have credences relative to which H and E are independent if all members do, and a condition stating that a group must judge H and E to be independent (in the sense of believing their objective chances are independent) if all members do. While the latter has a clear intuitive motivation (though, as we have seen, it is one a reasons-based theory would reject, if members can believe a proposition for different, and non-decisive, reasons), the former does not. Absent some further argument in favor of Independence Preservation, we should reject that condition as unmotivated.

Finally, consider Continuity. Here is the complaint of Russell et al:

> The Fixed Prior rule is also *discontinuous*. An arbitrarily small difference in individual credences can make a large difference to the group credence. In particular, consider a series of cases where one individual's credence in w approaches zero, while everyone else assigns some constant positive credence. In the pooled evidence version, the group credence has the same positive value in each case—but if that individual's credence reaches zero, the group jumps down to join her. A

> similar point holds for the common ground evidence version, where instead each of the other individuals assign zero to w in each case. So either version of the rule has cases involving a discontinuous jump around zero. (1301)

Clearly, sharp jumps can also occur with a Majority Evidence version of the Fixed Prior rule. But is discontinuity such a bad thing? To begin with, let us keep in mind that our concern is with the reasons, or evidence, possessed by individuals and by the group. When an individual (rationally) comes to assign credence 1 to a proposition, that must be due to her coming to have that proposition as part of her evidence (or gaining some other evidence that entails it). Thus there is an important difference between an individual assigning credence 0.99 to a proposition and her assigning it credence 1, and it is appropriate that the group's credences respond strongly to that difference. Relatedly, if Uniqueness is true and the individuals themselves have rational credences, then 'a series of cases where ones individual's credence in w approaches zero' must also be a series of cases which differ from each other in what evidence that individual possesses, and these differences in the individual's evidence will sometimes (but not always) entail differences in the group's evidence. So while Russell et al are correct that 'the Fixed Prior rule is not sensitive to any differences in non-zero credences' (ibid), it is important to keep in mind that Uniqueness demands that differences in an individual's non-zero credences be themselves driven by differences in the propositions they assign credence 0, that is, differences in their evidence, and that these differences will sometimes entail differences in the group's evidence. Because the difference between extreme credences (0 or 1) and non-extreme credences corresponds to a difference between having a reason and not, a difference which is important by the lights of the reasons-based theorist, I do not find Continuity a compelling condition.

Let us take stock. As in the case of group-level binary beliefs, my proposal is that what credences a group ought to have is a function of what reasons, or evidence, the group has. Where P is the rational prior probability function and E is the total group-level evidence, the group's credences ought to equal $P(- \mid E)$, in parallel to what the reasons-based theorist says about rational individual credences. And as in the case of group binary beliefs, if a

group satisfies requirements of reason, it will satisfy requirements of coherence (in this case conditionalization) as a byproduct. This proposal (if group evidence is some function of individuals' evidence) violates a number of other constraints that might be imposed on an aggregation function, but these further conditions, while *prima facie* attractive, do not stand up to scrutiny and should be rejected.

5 Operationalizability and Guidance

I expect that many will find the proposed views about rational group belief and rational group credence implausible. The main objection, I take it, will be that the reasons-based theory of rational group attitudes does not provide an operationalizable (or practical, or implementable) guide to what beliefs or credences a group ought to have. It does not provide adequate guidance.

The reasons-based theory may fail to be operationalizable or guidance-providing for two reasons. First, it will sometimes be impossible to know what the group's evidence is. Second, it will sometimes be impossible to know what beliefs or credences are supported by that evidence.

Start with the former. If evidence propositions (and reasons more generally) must be true, then it may be difficult or impossible to know what individuals' or the group's evidence is. This is one reason why Buchak and Pettit propose that a group's evidence is a function of individuals' *beliefs*, so that e.g., a group's evidence includes E just in case a majority of members believe H.[19] It may be easy to determine what individuals believe, but it is far

[19]They oppose conceiving of a group's evidence as a function of individuals' credences for usability reasons: 'in the vast majority of cases members cannot put fine-grained credences at the disposal of the group, offering them as inputs to a suitable aggreation function' because 'individuals do not have access to their degress of belief at any fine level of grain' (2015, 222). They also criticize a proposal somewhat related to my own on the grounds that it is not operationalizable. They consider a view on which the group ought to track members' evidence and 'votes are meant to be an indication of what the evidence is, not the evidence itself' (2015, 223). But they argue that on this view, discovering that majoritarian voting on each proposition yields inconsistent group-level beliefs could not play the epistemic role envisaged by Kolodny in the individual case. There, Kolodny holds that discovering that one has inconsistent beliefs indicates that at least one of one's beliefs is not supported by one's evidence, and thus provides a second-order reason to re-assess and eliminate the inconsistency. But Buchak and Pettit think that this move will not work in the group case:

harder (the objection goes) to determine what their evidence is (or what the group's evidence is, if group evidence is not a function of individuals' evidence). But of course beliefs need not be true, and so their proposal violates Factivity.

Now turn to the second issue, that it may be impossible to know what beliefs or credences are supported by a body of evidence. This is a common objection to Uniqueness. It is less pressing in the case of binary beliefs than in the case of credences (Kelly 2013). But with credences, it seems that Uniqueness may make it impossible to know what credences one ought to have. After all, how are we to know what the uniquely rational prior probability function is? The problem is even worse if, as I suggest, it is indeterminate what the uniquely rational prior is. Thus, even if we could come to know what the group's evidence is, the reasons-based theory would still fail to provide adequate guidance.

However, these considerations about operationalizability and guidance do not show that the reasons-based theory is false. First, note that they are not unique to the case of groups. Insofar as they cast doubt on the reasons-based theory of group rationality, they also cast doubt on the analogous theory of individual rationality. But Factivity is almost universally accepted, and Uniqueness is widely defended though undoubtedly controversial. Of course, this might just mean that the reasons-based theory is false in both the group and the individual cases.

But second, and more importantly, many epistemologists now reject the demand that their theories be operationalizable, on the grounds that *no* theory of rationality can meet this demand. One reason stems from Williamson's (2000) anti-luminosity argument, which

> [T]he proposed procedure would work for sure only in the special case where the evidence at the disposal of members is the same. Let that assumption be lifted and it will no longer be clear that an incoherence in majority voting indicates that a majority are responding improperly to their evidence. The incoherence may simply be the result of the fact that different members confront different bodies of evidence; and yet it may be that those bodies of evidence, considered as a whole, would resolve the conflict. Second, it is highly optimistic to think that even in the simplified case where all members confront the same evidence, there will be a way for the group agent to determine the particular proposition on which a majority of members are failing to respond appropriately. In both cases, the group will not have evidence sufficient to conclude that it believes one of the three conflicting propositions without sufficient reason. (223)

purports to show that no (non-trivial) condition is luminous, in the sense that whenever one is in that condition, one is in a position to know that one is in it. Hence conditions of the form *being such that one ought to have such-and-such attitudes* likewise fail to be luminous. Thus, if we understand the demand for operationalizability as the demand that one always be in a position to know what one ought to do, and the anti-luminosity argument is sound, then no theory is operationalizable.

Of course, the anti-luminosity argument is controversial, and even if it is sound one might reject a knowledge-based formulation of the operationalizability criterion. But even setting aside knowledge-based formulations of operationalizability, it is important to note that our powers of introspection are far from perfect. Thus, while it may be difficult to determine whether one *knows* a proposition (and thus has it as part of one's evidence, given E=K), it may also be difficult to determine whether one believes it. In the group case, this means that while members' votes may be an imperfect guide to what they know, they are also an imperfect guide to what they believe. See Srinivasan (2015) for extensive discussion of reasons to reject an operationalizability criterion in theories of normativity (see also Williamson 2008).

What we have, I think, is a gradient. Some theories may be more operationalizable, or provide more practical guidance, than others. But it is not clear that we should always prefer more operationalizable theories over less operationalizable ones; operationalizability is surely not the only criterion by which to judge a theory of rationality. Nor is there likely to be any non-arbitrary threshold level of operationalizability which is necessary for a theory of rationality to be true.

Nevertheless, the fact that no theory is fully operationalizable, and the fact that (in my view) the correct theory of rationality has the result that one often *ought* to be uncertain what one's beliefs and credences ought to be, does mean that we should sharply distinguish two questions: First, what beliefs and credences ought a group have? Second, what mechanisms or procedures are reasonably reliable ways of yielding group-level beliefs and credences that are

as close as possible to those that the group ought to have?[20] (Note that a similar distinction is relevant in the case of individuals, where we should distinguish the question of what beliefs and credences an individual ought to have, given certain evidence, and the question of what belief-forming mechanisms (habits, patterns of reasoning, etc) the individual ought to employ in attempting to arrive at rational beliefs and credences.)

An analogous distinction is important in the theory of social choice. Arrow (1963/1951) considered functions from a set of individual rankings of alternatives to a social ranking of alternatives and proved that there is no such function that satisfies each of a number of *prima facie* attractive conditions. But the rankings can be understood in different ways. In the context of welfare economics, the social ranking is interpreted as representing social betterness (or perhaps what preferences society as a whole ought to have), and the individual rankings are interpreted as either each individual's preferences or rankings of betterness for that individual. In the context of voting theory, by contrast, the individual rankings are interpreted as the rankings of candidates or policies expressed on a ballot, and the social ranking is interpreted as the outcome of the election. It is important to keep this distinction in mind. Among other things, different conditions on functions from individual rankings to a social ranking require different motivations in the context of welfare economics and in the context of voting theory. And some conditions may have considerably more plausibility in one context than in the other (see fn. 18 above).

In addition, provided that the notion of social betterness makes sense, it is natural to

[20] Alternatively, one could think the of the ultimate (epistemic) goal as having *accurate* beliefs and credences rather than *rational* ones, in which case one might want to evaluate attitude-forming mechanisms and procedures for how reliably they tend to yield accurate group-level attitudes. I have no quarrel with this thought, so long as, again, we do not conflate the question of what attitude-forming mechanisms and procedures ought to be adopted with the question of what attitudes are rational. Note that this is not to take a stand on the attractiveness of an 'accuracy-first' approach to epistemic rationality, along the lines of Pettigrew (2016). Accuracy-first theorists typically to use accuracy dominance and expected accuracy considerations to justify epistemic principles. But my claim is simply that we should not conflate the *actual* reliability of procedures to yield accurate attitudes with what attitudes are rational. Even most accuracy-first theorists will have no objection to this insistence. Of course, whether an accuracy-first approach is compatible with other theoretical commitments of the reasons-based approach, and in particular with Uniqueness, is a separate question and one I will not address here. In what follows, I will stick with the assumption that mechanisms and procedures are to be evaluated for their tendency to produce rational attitudes, but the same points will apply if one adopts an accuracy-centric approach instead.

think that different voting rules should be judged in large part by how reliably they yield approximations of the social betterness ranking. Voting systems can be viewed as fallible means of detecting what the social good consists in (cf. Coleman and Ferejohn 1986). By the same token, we can view practical procedures for arriving at group beliefs or credences on the basis of individual beliefs or credences as fallible but hopefully reasonably accurate ways of detecting approximately what beliefs or credences the group ought to have.

One might worry that it is impossible to evaluate voting systems for how effective they are as a means to the end of rational group beliefs and credences, since it is sometimes impossible to know what beliefs and credences the group ought to have. If it is impossible to specify the end, it is likewise impossible to evaluate means for effectiveness in yielding that end. But while it may not *always* be possible to know what beliefs and credences a group ought to have, this does not mean that such knowledge is never, or even rarely, possible. It also does not mean that it is always, or even often, impossible to know *approximately* what beliefs and credences a group ought to have. Compare the case of individuals. Even though we may not always be in a position to know exactly what beliefs and credences an individual ought to have, nevertheless we very often have a fairly good idea. In arriving at such judgments, we employ considerations of simplicity, explanatory power, intuitive plausibility, and the like. The same considerations can be used to arrive at judgments of approximately what beliefs and credences a group ought to have. And so long as such judgments are available enough of the time, we can use them (along with other information) in evaluating the effectiveness of voting systems and other mechanisms for reaching group beliefs and credences.

Importantly, voting systems, and practical procedures more generally, cannot be evaluated solely on the basis of *a priori* theoretical considerations. After all, they are means to an end, and empirical evidence is essential to judging how effective those means are. For instance, how reliable a voting system is in approximating the social betterness ranking, or what beliefs and credences a group ought to have, depends on factors like what profiles of individual attitudes are likely (rather than just which ones are possible) and what the actual risks are of strategic

voting and agenda manipulation.

It is also worth noting that such empirical evidence can paint a very different picture about the prospects for good voting systems than the theoretical work alone. Mackie (2003, especially Ch. 3) surveys evidence that given actual preference profiles, majoritarian cycles are rare and popular voting rules (other than the much-derided plurality rule) yield roughly the same results the vast majority of the time.[21] Moreover, he argues that actual threats of cycles from pairwise majority voting and of harmful strategic voting, while no doubt real, are far less serious than theoretical work on paradoxes of social choice (Arrow 1963/1951), strategic voting (Gibbard 1973; Satterthwaite 1975), and multidimensional chaos (McKelvey 1976; Schofield 1978) might be taken to suggest. Thus, while theoretical work may suggest that there is no 'perfect' voting system, it seems likely that a number of voting systems are fairly good ways of accurately approximately the social betterness ranking.[22]

Similarly, it may be that given what profiles of individual beliefs and credences are likely to arise, a number of realistic, feasible voting systems or other aggregation procedures are fairly good ways of accurately approximating what the group's beliefs and credences actually ought to be. Of course, whether this is in fact the case depends on what the empirical facts are. Unfortunately, while there is a great deal of empirical work on preference aggregation, there is far less on belief and credence aggregation. So we know only that proposition-wise

[21] See especially the studies by Chamberlin et al (1984), Feld and Grofman (1992), Levin and Nalebuff (1995), Felsenthal and Machover (1995), and Regenwetter et al (2006).

[22] Some theoretical work also points in this optimistic direction. For instance, Nurmi (1992) ran computer simulations which involved deviations from the so-called *impartial culture* assumption, on which each individual has an equal, and independent, probability of having any of the possible preference rankings. Nurmi found that under even small deviations from impartial culture, popular voting rules other than plurality agreed with each other on the winner the vast majority of the time.

List and Goodin (2001) give a model of voting over more than two alternatives in which in which there is a best candidate or policy and each voter is more likely to rank the best candidate or policy first than she is to rank each of the suboptimal alternatives first. Their findings on the accuracy of a set of standard voting rules (plurarity, Condorcet, Borda, Hare, and Coombs), given the assumptions of their model, lead them to conclude that 'Any of these standard decision rules is pretty much as good as any other, on epistemic grounds' and that 'All of these standard decision rules have great epistemic merits, at least whenever the electorate is reasonably large' (p. 294).

Finally, Dryzek and List (2003) argue that deliberation can in principle help engender some of the conditions (e.g., single-peakedness of preferences) necessary to avoid the negative results of Arrow, Gibbard, and Satterthwaite, and they cite empirical evidence that deliberation not only can, but is likely to, have these results.

majority voting *may* yield inconsistent group-level beliefs and that linear averaging *may* yield violations of conditionalization at the group level, but do not know whether these violations of coherence requirements are likely and, when they occur, how egregious they are likely to be.

In any event, regardless of whether the empirical evidence paints an optimistic or a pessimistic picture, it remains the case that empirical evidence is necessary to determine which voting rules or other group-attitude-forming mechanisms should be adopted. Theoretical work on voting theory and judgment aggregation certainly has an important role to play here (among other things, in investigating what follows given different assumptions about how things stand empirically, e.g., with respect to the distribution of individual attitudes). But such theoretical work alone cannot be the whole story.

I have been emphasizing the evaluation of practically implementable group attitude-forming procedures in terms of their reliability in outputting attitudes that approximate those that the group rationally ought to have. But in closing this section, I want to flag that it is generally sensible to evaluate such attitude-forming procedures on other grounds as well. First, we may care about how costly a given procedure is in terms of time, complexity, and so on. Second, we may care about fairness, inclusivity, and how satisfied group members are with the procedures adopted (see Steele et al 2007). Third, there may be reasons to care about whether a procedure is liable to yield incoherent group attitudes, over and above the fact that incoherence entails that those attitudes violate some requirement of reason or other. For one thing, incoherence itself can bring practical costs in terms of predictable exploitability, and so it can be worse to have a set of incoherent attitudes than a set of coherent ones, even if the former are closer than the latter to the attitudes the group rationally ought to have. For another, it is arguably a necessary condition of agency that one's attitudes not display too much incoherence, and that one has a disposition to correct incoherence among one's attitudes upon detecting it. (Note, though, that a reasons-based theorist might respond that agency is tied to a disposition to have (approximately) substantively rational attitudes rather

than to a disposition just to have formally coherent ones. On an interpretationist theory of mental states, this would involve the relevant principle of charity being cashed out in terms of requirements of reason rather than coherence.) If so, and if it is desirable that a given group qualify as an agent, then it will be desirable to implement attitude-forming mechanisms that won't generate egregiously incoherent attitudes, or at least mechanisms that will detect and correct incoherence when it arises.

6 Group Epistemology vs. Social Choice

So far, I have been talking almost exclusively about group epistemology. Can a similar reasons-based story be told about group *practical* rationality? There is one major obstacle. As we saw, in both the individual and the group cases, satisfying requirements of reason with respect to some type of attitude entails satisfying requirements of coherence as a byproduct only if a Uniqueness thesis holds for that type of attitude; that is, only if one's reasons uniquely fix what attitudes of that type one ought to have. Consider the coherence requirement that one's preferences be transitive, and suppose that one violates this requirement by prefering A over B, B over C, and C over A. If there were a unique set of preferences supported by one's reasons, it would be the case that at least one of one's preferences is contrary to reason. For instance, if rational preferences must correspond to facts about betterness (perhaps betterness for that agent), and if betterness is transitive, then having intransitive preferences means that for at least one pair x and y, one prefers x over y despite x's not being better than y.

But while Uniqueness for beliefs and credences has its fair share of adherents, Uniqueness for preferences and other practical attitudes is perhaps almost universally rejected, including by Kolodny himself (2007, 251).[23] Humeans about preference hold that the ends that one sets oneself cannot be irrational, and even non-Humeans typically hold that there is no single set of ends that one is rationally required to adopt. Uniqueness for intentions is even less

[23]Kolodny sketches an alternative route to the claim that satisfying requirements of reason yields satisfying requirements of coherence as a byproduct, one which does not rely on Uniqueness. I will not consider this proposal here.

plausible. Plausibly, in Buridan's ass cases, it is permissible to intend to go left or to intend to go right.

Now, Uniqueness theses for practical attitudes are less implausible than it might initially seem (Hedden 2015b; Greco and Hedden 2016). While no doubt different preferences are rationally permissible, it is far from obvious that different preferences are rationally permissible even in light of the same reasons. For example, while different tastes in ice cream or music may be rationally permissible, it may be that one's reasons (e.g., one's tastes, or perhaps one's evidence regarding what one's tastes are) uniquely fix what one's taste-based preferences ought to be. Likewise, it is not intuitively obvious that different values and hence value-based preferences can be equally rational; after all, values are the sorts of things we argue about and aim to provide reasons for. And even if different people can rationally have different value-based preferences, it may be that these differences must stem from differences in their reasons, including their evidence about the normative facts. Perhaps two people with exactly the same reasons cannot rationally differ in their value-based preferences. (Even Uniqueness for intentions has something to be said for it, though I won't go into details here.)

But if Uniqueness for preferences fails even for groups, then there can be no reasons-based theory of group practical rationality. This would not, I think, cast doubt on the plausibility of a reasons-based theory of group epistemic rationality. For there may be a story about why Uniqueness holds for beliefs and credences but not for preferences. For instance, it may be that there are objective standards of correctness for beliefs and credences (e.g., truth or accuracy) but not for preferences (e.g., there may be no such thing as an objective betterness ordering). Only if no such explanation can be given would the failure of a reasons-based theory of practical rationality cast doubt on the reasons-based theory of epistemic rationality.

7 Conclusion

Reasons, and in particular evidence, play a central role in the theory of individual rationality. I have argued that they should play a similarly central role in the theory of group rationality.

Much previous work on group rationality focuses on functions that aggregate individuals' beliefs and credences without attention to what reasons and evidence those attitudes are based on and what reasons are possessed by the group. On my view, a notion of group-level reasons is needed to determine which group-level attitudes are rational. What reasons the group has may be a function of what reasons are possessed by each individual member, or it may depend also on how the group is in fact structured. Either way, the group ought to have the attitudes that are best supported by its reasons, just as individuals ought to have the attitudes that are best supported by *their* reasons. In this way, my view treats group rationality in parallel with individual rationality. Moreover, it has the result that if the group has the doxastic attitudes that it ought to have, it will thereby also satisfy requirements of coherence as a byproduct.

On this approach, it is important to distinguish between (i) what attitudes a group rationally ought to have, and (ii) what attitude-forming procedures ought to be adopted as a means of approximating the rationally required group attitudes (or, perhaps, as a means of approximating the *objectively correct* group attitudes). Making this distinction has three important upshots for research on group rationality. First, we should be suspicious of conditions (or axioms) in the theory of judgment aggregation that lack a clear motivation in a reasons-based framework. For instance, we saw above that even Unanimity is ill-motivated, given that members can have different reasons for having that credence, and the group's reasons might not be identical to those of any particular member.

Second, even the possibility of yielding incoherent group attitudes is not a decisive strike against an attitude-forming procedure. After all, while the set of group attitudes that a group rationally ought to have will always be coherent, some incoherent sets of attitudes are closer to rationally ideal attitudes than some coherent sets of attitudes (the beliefs of a perfectly coherent conspiracy theorist are further from the rational ideal than some slightly inconsistent but far more sensible beliefs). Given that no practically implementable attitude-forming procedure will always yield the rationally required group attitudes, a procedure that

never yields incoherent attitudes may not be superior to a procedure that sometimes does so. Whether this is so will depend on factors like how frequently the latter yields incoherent group attitudes, and how egregious the incoherence is, as these factors indicate how reliable it is in approximating the rationally required group attitudes.

Third, while what attitudes a group rationally ought to have given a body of reasons is a largely, if not wholly, *a priori* matter, what attitude-forming procedures ought to be adopted is an empirical one. While theoretical work focusing on impossibility theorems and computer simulations certainly has a crucial role to play in this investigation, it must be supplemented by empirical investigation in order to determine what procedures are actually most reliable in yielding attitudes that best approximate those that a group rationally ought to have.[24]

[24]Thanks to Kenny Easwaran, Daniel Greco, Caspar Hare, Matthew Kopec, Philip Pettit, and Daniel Wodak, as well as audiences at MIT, the University of Sydney, and the Society for Exact Philosophy, for helpful feedback on this work.

References

Arrow, K. 1963 (1951). *Social Choice and Individual Values*. 2nd Ed. New Haven: Yale University Press.

Aumann, R. 1976. 'Agreeing to disagree.' *The Annals of Statistics*, 4(6): 1236–1239.

Broome, J. 1999. 'Normative Requirements.' *Ratio* 12 (4):398–419.

Broome, J. 2007. 'Wide or Narrow Scope?' *Mind* 116 (462):359–370.

Buchak, L. and Pettit, P. 2015. 'Reasons and Rationality: The Case of Group Agents.' In Hirose and Reisner (eds), *Weighing and Reasoning: Themes from the Philosophy of John Broome*. New York: Oxford University Press, pp. 207-31.

Coleman. J. and Ferejohn, J. 1986. 'Democracy and Social Choice.' *Ethics* 97: 26–38.

Condorcet, Marquis de. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: Imprimerie Royale. (Reprinted New York: Chelsea Publishing, 1972).

Chamberlin, J., Cohen, J., and Coombs, C. 1984. 'Social Choice Observed: Five Presidential Elections of the American Psychological Association.' *Journal of Politics* 46: 479–502.

Christensen, D. 2007. 'Epistemology of Disagreement: The Good News.' *Philosophical Review* 116: 187–217.

Dietrich, F. and List, C. 2016. 'Probabilistic Opinion Pooling.' In Hàjek and Hitchcock (eds), *Oxford Handbook of Philosophy and Probability*. Oxford: Oxford University Press.

Dogramaci, S. and Horowitz, S. Forthcoming. 'An Argument for Uniqueness about Evidential Support.' *Philosophical Issues* 26: Knowledge and Mind.

Dryzek, J. and List, C. 2003. 'Social Choice Theory and Deliberative Democracy: A Reconciliation.' *British Journal of Political Science* 33(1): 1–28.

Feld, S. and Grofman, B. 1992. 'Who's Afraid of the Big Bad Cycle? Evidence from 36 Elections.' *Journal of Theoretical Politics* 4: 231–7.

Felsenthal, D. and Machover, M. 1995. 'Who Ought to be Elected and Who is Elected? An Empirical Investigation of 92 Elections Under Three Procedures.' *Electoral Studies* 14:143–69.

Genest, C. and Zidek, J.V. 1986. 'Combining Probability Distributions: A Critique and an Annotated Bibliography.' *Statistical Science* 1(1): 114–35.

Gibbard, A. 1973. 'Manipulation of Voting Schemes: A General Result.' *Econometrica* 41: 587–601.

Goodman, N. 1955. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Greaves, H. and Wallace, D. 2006. 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.' *Mind* 115 (459): 607–632.

Greco, D. and Hedden, B. 2016. 'Uniqueness and Metaepistemology.' *Journal of Philosophy* 113 (8): 365–395.

Hedden, B. 2015a. 'Time-Slice Rationality.' *Mind* 124 (494): 449–91.

Hedden, B. 2015b. *Reasons without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.

Hedden, B. 2015c. 'A Defense of Objectivism about Evidential Support.' *Canadian Journal of Philosophy* 45 (5): 716–43.

Horowitz, S. 2013. 'Immoderately Rational.' *Philosophical Studies* 167 (1): 1–16.

Joyce, J. 1998. 'A Nonpragmatic Vindication of Probabilism.' *Philosophy of Science* 65: 575–603.

Kelly. T. 2013. 'Evidence can be Permissive.' In Steup, Turri, and Sosa (eds.), *Contemporary Debates in Epistemology*. Wiley-Blackwell, pp 298–311.

Kolodny, N. 2005. 'Why Be Rational?' *Mind* 114 (455):509–563.

Kolodny, N. 2007. 'How Does Coherence Matter?' *Proceedings of the Aristotelian Society* 107: 229–63.

Kolodny, N. 2008. 'The Myth of Practical Consistency.' *European Journal of Philosophy* 16 (3):366–402.

Kornhauser, L.A. and Sager, L.G. 1986. 'Unpacking the Court.' *The Yale Law Journal* 96 (1): 82–117.

Lackey, J. 2016. 'What is Justified Group Belief?' *Philosophical Review* 125 (3): 341–96.

Lederman, H. 2015. 'People with Common Priors Can Agree to Disagree.' *The Review of Symbolic Logic* 8.1: 11–45.

Lederman, H. 'Uncommon Knowledge.' Forthcoming in *Mind*

Levin, J. and Nalebuff, B. 1995. 'An Introduction to Vote-Counting Schemes.' *Journal of Economic Perspectives* 9: 3–26.

Littlejohn, C. 2013. 'No Evidence is False.' *Acta Analytica* 28 (2):145–59.

List, C. 'What is it Like to be a Group Agent?' Forthcoming in *Noûs*.

List, C. and Goodin, R. 2001. 'Epistemic Democracy: Generalizing the Condorcet Jury Theorem.' *Journal of Political Philosophy* 9 (3):277–306.

List, C. and Pettit, P. 2002. 'Aggregating sets of judgments: An impossibility result.' *Economics and Philosophy* 18: 89–110.

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

McKelvey, R. 1976. 'Intransitivities in Multi-Dimensional Voting Models and Some Implications for Agenda Control.' *Journal of Economic Theory* 12: 472–482.

Nurmi, H. 1992. 'An Assessment of Voting Systems Simulations.' *Public Choice* 73: 459–87.

Pettigrew, R. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Pettigrew, R. 'On the Accuracy of Group Credences.' Forthcoming in Szabó Gendler and Hawthorne (eds), *Oxford Studies in Epistemology*, volume 6.

Pettit, P. 2001. 'Deliberative Democracy and the Discursive Dilemma.' *Philosophical Issues* 11: 268–99.

Raz, J. 1975. *Practical Reasoning and Norms*. London: Hutchinson & Co.

Regenwetter, M., Grofman, B., Marley, A., and Tsetlin, I. 2006. *Behavioral Social Choice*. Cambridge: Cambridge University Press.

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

Satterthwaite, M. 1975. 'Strategy Proofness and Arrow's Conditions.' *Journal of Economic Theory* 10: 187–217.

Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University Press.

Schoenfield, M. 2013. 'Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief.' *Noûs* 47 (1):193–218.

Schofield, N. 1978. 'Instability of Simple Dynamic Games.' *Review of Economic Studies* 45: 575–594.

Schroeder, M. 2008. 'Having Reasons.' *Philosophical Studies* 139 (1):57–71.

Schwitzgebel, E. 2015. 'If Materialism is True, the United States is Probably Conscious.' *hilosophical Studies* 172 (7):1697–1721.

Sen, A. 1970. *Collective Choice and Social Welfare*. Amsterdam: Elsevier.

Srinivasan, A. 2015. 'Normativity Without Cartesian Privilege.' *Philosophical Issues* 25 (1): 273–99.

Steele, K., Regan, H., Colyvan, M., and Burgman, M. 2007. 'Right Decisions or Happy Decision-Makers?' *Social Epistemology* 21 (4): 349–68.

Titelbaum, M. and Kopec, M. 'Plausible Permissivism.' Manuscript.

White, R. 2005. 'Epistemic Permissiveness.' *Philosophical Perspectives* 19 (1): 445–459.

Williamson, J. 2015. 'Deliberation, Judgment and the Nature of Evidence.' *Economics and Philosophy* 31: 27–65.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, T. 2008. 'Why Epistemology cannot be Operationalized.' In Smith (ed), *Epistemology: New Essays*. New York: Oxford University Press. pp. 277–300.