# Social Process Reliabilism:
# Solving Justification Problems in Collective Epistemology

## Alvin I. Goldman

1. <u>The Epistemology of Collective Subjects</u>

In recent years many philosophers have pursued questions about the nature of groups and group activity, including group action, group intention, group belief, and so forth.  Salient contributions to this literature include Bratman (1999); Gilbert (1989, 1994, 2002); Lackey (this volume; forthcoming), List and Pettit (2005, 2011); Pettit (2003); Quinton (1975/76); Schmitt (1994, 2006); Searle (1995); Tollefsen (2007); Tuomela (1992, 2007); and Wray (2010).  There is widespread agreement that groups -- at least some groups -- are the bearers of propositional attitudes.  Such attitudes are colloquially ascribed to many groups, and most philosophers who write on the subject endorse the aptness of such ascriptions.  Of special interest here is a subclass of propositional attitudes, namely, "<u>doxastic</u>" attitudes, including belief, disbelief, suspension of judgment, and degrees of confidence.  Courts of law, scientific panels, athletic teams, and institutional bodies of many varieties are described in terms of such attitudes.

A few social epistemologists have worked not only on collective belief <u>per se</u> but on its epistemic appraisal.  This topic has been treated by Schmitt (1994, 2006), Mathiesen (2006), Bird (2010), and List and Pettit (2011), among others. Here I focus on one narrow question in collective epistemology: the justification question.   Justification questions certainly arise from time to time in public discourse.  When the Bush Administration claimed that Saddam Hussein possessed weapons of mass destruction, many people questioned whether the Administration was justified in holding this belief (assuming they really believed it).

When it comes to group justification, a natural thought is that such justifiedness must arise from the justified beliefs of its members.  This paper therefore devotes considerable attention to the question of <u>justificational dependence</u> in general, laying the ground for decisions about justificational dependence of groups on their members.  Collective epistemologists generally agree that a group's possession of a belief typically derives (in some fashion) from those of its members.  In this paper, however, the focus is not on group <u>belief</u> but on the justificational status of such beliefs.

It is generally acknowledged that groups come in different varieties, ranging from loose, unstructured assemblages of individuals to highly structured, organized collectivities.  The present discussion focuses on more formally structured groups with some kind of unifying organization.  Whether the theory also applies to looser, merely "summative" groups is discussed briefly in section 10.

## 2.  The Metaphysics and Epistemology of Collective Entities

How the beliefs of collective subjects are aligned with beliefs of their members is one question for the theory of group, or collective, subjects.  List and Pettit suggest that collective subjects have a qualified supervenience relation to their members.  This is a reasonable metaphysical position, by my lights, but it might be preferable to substitute an increasingly popular term of metaphysical art, namely, the grounding relation (Fine 1994, Schaffer 2009).  The claim might then be that the propositional attitudes of collective subjects are grounded in propositional attitudes of their members plus the group's organizational structure.  Under either proposal, group agents are seen as distinct from their members (singly or collectively), a result that List and Pettit, for example, embrace.

A full investigation of the relations between groups and their members, however, must distinguish between metaphysical and epistemological matters.  For present purposes there are two types of questions in this territory I call "metaphysical'.  One type is:  Do group entities exist at all?  What types of groups are there?  A second type of question is psychological.  Do group entities have psychological properties, specifically, propositional attitudes?  If so, how are a group's attitudes related to those of its members?  Both types of questions should be distinguished cleanly from epistemological ones, especially the justification question.  My main aim here is to inquire into the determinants of justificational status (J-status) for collective beliefs.

## 3.  Collective Psychology vs. Collective Epistemology

Before turning to questions about collective justifiedness, however, let us consider the subject of collective belief.  List and Pettit (2011) offer the most systematic treatment of collective belief,[1] and I shall generally follow their treatment.  However, I shall not spell out in detail any particular theory of group belief.  That is a difficult enough problem in its own right, and I'll have enough on my hands with

the theory of justification.  Nonetheless, it helps to have a suitable sample of such a theory before us, and I shall use the List-Pettit model for that purpose.[2]

List and Pettit assume that group beliefs are the outputs of a function that takes profiles of individual members' beliefs as inputs and yields collective beliefs as outputs.  Such a mapping is a judgment aggregation function, or, as I shall call it, a <u>belief aggregation function</u> (BAF).  An example of a BAF is the majoritarian rule according to which a group believes proposition P if and only if a majority of group members believe P.  Another example is a supermajoritarian rule, in which the group believes P just in case a qualified majority of members do so (e.g., two thirds).  Another example is dictatorship, in which the group belief is always the same as that of a fixed member.

The idea of a BAF is initially presented in terms of rules that groups formulate and adhere to when forming beliefs of their own as a function of their members' beliefs.  However, the same concept can be applied to groups that have no officially formulated rules to govern their belief-forming methods.  Instead, interpersonal mechanisms of psychological influence give rise to patterns describable in terms of such a rule.  These types of groups can be subsumed under the BAF framework.  List and Pettit call this a "functionally inexplicit organizational structure" (2011: 60).  "Organizational structure" also refers to ways by which a group enacts and implements a BAF (List and Pettit 2011: 60).  For example, majority voting can be implemented by paper ballots or electronic ones, or by a process of deliberation that leads group members to converge on the attitudes of the majority.

Having introduced the idea of an organizational structure, List and Pettit proceed to explain the relation between group attitudes and member attitudes as follows:

> The things a group agent does are clearly determined by the things its members do;
> they cannot emerge independently.  In particular, no group agent can form intentional attitudes
> without these being determined, in one way or other, by certain contributions of its members, and
> no group agent can act without one or more of its members acting.  (2011: 64).

Despite this, however, they also argue for a degree of autonomy for group agents.  Group beliefs do not supervene on member beliefs in a proposition-wise fashion but instead via a looser relation of <u>holistic</u> supervenience in which the set of group beliefs across propositions is determined by sets of individual attitudes across all of the same propositions (not each proposition taken singly) (2011: 69).  From this they infer that "individual and group attitudes can come apart in surprising ways, thereby establishing a

certain autonomy for the group agent" (2011: 69).  For present purposes, as indicated earlier, it suffices to say that group minds have the metaphysical relation of being grounded in individual minds.  Because grounding is an irreflexive relation, it follows that group minds are non-identical with any individual minds (or sets thereof), and non-identity is all we really need for our epistemological purposes.

4.  Justification Transmission: Intrapersonal vs. Interpersonal

Before turning to group justification, it will help to highlight certain features of justificational factors in individuals.  In the individual domain I shall spotlight some contrasts between intrapersonal and interpersonal justification relations.  It is convenient at this point to formulate the issues in the terminology of "reasons."

Suppose Jones has a good reason, R, for believing P.  For example, she already justifiedly believes certain premises that entail P and she recognizes this entailment.  Does she therefore have at least prima facie (propositional) justification for believing P?  Yes.  Special cases aside, possessing good reason(s) R automatically gives Jones such justification for believing P.[3]  Now consider a different case. Although Jones lacks reason R, Smith possesses R.  Does Smith's having reason R automatically give Jones prima facie justification for believing P?  Of course not!  One person's reason, or justification, never transfers automatically to a second person.  Granted, Jones might acquire good reason to believe that Smith has a good reason to believe P; she might have this higher-order reason without knowing what Smith's reason is.  As Richard Feldman has quipped, "evidence for evidence is evidence."  In this new case Jones' evidence is not R, but the (different) fact that Smith has a good reason to believe P.  The fact that Jones has this higher-order reason does not imply that this reason emerges automatically from Smith's having his first-order reason, namely R.

The foregoing points are perfectly compatible, of course, with the thesis that if Smith has reason R to believe P, he can often easily convey it to Jones via testimony.  Hearing Smith assert P and provide the supporting reason R will often give Jones reason to believe P.  But such conveyance of reasons does not occur automatically; a testimonial act (plus other suitable conditions) is required.  And such a testimonial act will not always suffice, especially if Jones has prior reasons to doubt Smith's credibility.

Granted this point, can a hearer become testimonially justified in believing P if the speaker himself is unjustified in believing P?  In other words, is speaker justifiedness necessary for hearer justifiedness?  Lackey (1999, 2008) argues convincingly to the contrary for both knowledge and

justification. She offers a varied array of examples to demonstrate the point that speaker justifiedness is neither necessary nor sufficient for hearer knowledge or justifiedness (Lackey 2008: 47-71). Even Schmitt (2006), who defends what he calls the "transindividual basing thesis," presents an example of his own to show the non-necessity of a speaker being justified as a condition of a hearer being justified. Schmitt writes: "I can have a justified belief that there are strawberries in the refrigerator on Tina's testimony even if Tina has no good reason to believe this, if I justifiedly though mistakenly believe that Tina has a good reason to believe this" (2006: 193-4).

Our target questions about <u>group</u> justification parallel those we have just been discussing. Are the J-statuses of members' beliefs automatically transmitted to the J-status of a group belief? Does the relation between group and member justifiedness parallel the case of <u>intra</u>personal inference, in which the J-statuses of one's own premise beliefs are automatically transferred to one's conclusion beliefs? Or is it more similar to <u>inter</u>personal cases, in which no such automatic transmission occurs, and instead testimonial acts are normally required? Let us proceed here with the help of an example.

5. <u>Member/Group Justification Transmission? The British Museum Example</u>

What kinds of justification relations obtain between a group's belief in P and its members' beliefs in P? To explore this problem, consider the following example (see Table 1). G is a group whose members consist of 100 guards ($M_1$ - $M_{100}$) at the British Museum. Each of the first 20 guards, $M_1$-$M_{20}$, justifiedly believes that guard Albert is planning an inside theft of a famous painting (= A). By deduction from A, each of them infers the (existential) proposition that there is a guard who is planning such a theft (= T). The remaining 80 guards do not believe and are not justified in believing A. Each of the second 20 guards, $M_{21}$-$M_{40}$, justifiedly believes that Bernard is planning an inside theft (= B), and deductively infers T from B. The other 80 members do not believe B. Each of a third group of 20 members, $M_{41}$-$M_{60}$, justifiedly believes that guard Cecil is planning an inside theft (= C) and deductively infers T from C. The 80 others do not believe C. Thus, 60 members of G (justifiedly) believe T by deduction from some premise he/she justifiedly believes.

What does group G believe (whether justifiedly or unjustifiedly) with respect to these matters? To answer this question we must make some assumptions about G's psychology, especially the BAF by which members' beliefs in various propositions are mapped into G's beliefs. Assume that this BAF is a supermajoritarian operation that proceeds in a proposition-wise fashion. G believes a proposition P if

and only if at least 60% of its members believe P.   Because 60% of G's members believe T, it follows that G itself believes T.  But no premise or set of premises believed by G jointly entail T.  G believes neither A nor B nor C, because each of these propositions is affirmed by only 20% of G's members and rejected by the remainder.  Further assume that G believes no other suitable premises.  So, G believes no premises that deductively imply T; but G does believe T.  The upshot is that G believes T; but it is unclear whether its belief in T is justified.

<div style="text-align:center">Individual-Level Beliefs</div>

| About the premises | About the conclusion |
|---|---|
| 20 members bel (A) (J'dly) | |
| 20 members bel (B) (J'dly) | |
| 20 members bel (C) (J'dly) | 60 members bel (T) (J'dly) |
| 80 members do not bel (A) | 40 members do not bel (T) |
| 80 members do not bel (B) | |
| 80 members do not bel (C) | |

<div style="text-align:center">Group-Level Beliefs</div>

| About the premises | About the conclusion |
|---|---|
| G does not bel (A) | |
| G does not bel (B) | G believes T. |
| G does not bel (C) | |

<div style="text-align:center">Table 1</div>

There are two ways G's justificational status vis-à-vis T might be approached: in light of other beliefs of G and their J-statuses or in light of the members' beliefs (vis-à-vis T) and their J-statuses. What emerges if we consider the J-status of G's belief in T in light of its members' beliefs?  Recall how the case was described.  First, only 20 members of G have justified beliefs in A, 20 have justified beliefs in B, and 20 have justified beliefs in C.  There is no obvious route from these J-statuses to G's being justified in believing A, believing B, or believing C.  But wait!  We also concluded that each of 60 members justifiedly believes T, in virtue of deducing it from either A, B, or C.  Moreover, each of the

premises on which the sixty members based their respective inferences to T justifiedly believed them. So, sixty members justifiedly believe T.

Let us now assume a sample justification aggregation function (JAF), viz., JAF-1:

(JAF-1)  If at least sixty percent of G's members justifiedly believe P, G is also justified in believing P.

JAF-1 is not presented as a correct JAF, or even a prototype of one.[4]  Nonetheless, it is instructive to consider JAF-1 as a sample principle in order to explore the British Museum example and get a feel for member/group aggregation patterns.  It will also help us draw a distinction between "vertical" and "horizontal" methods of group justification acquisition.  Notice that JAF-1 mirrors the BAF previously postulated for G.  There is no necessary connection, however, in the sense that a JAF must always "sanction," or approve of, whatever BAF a given group selects.  On the contrary, a given BAF may be one that a suitable JAF would classify as unsuitable for generating justified group beliefs. In particular, I will later propose principles of justifiedness that are linked to reliability, i.e., truth conduciveness, implying that an acceptable JAF would not authorize just any old BAF.  Proper BAFs (ones that can help generate justified group beliefs) will have to be ones that meet suitable standards of truth-conducivity.  This will become clearer in later sections.[5]  For illustrative purposes, however, it is convenient to begin with a "paired" set of BAF and JAF.  .

As previously indicated, I do not mean to present here any particular account of group belief, or how such beliefs are formed.  On one possible approach groups can select their own BAFs as they see fit.  It might be done by some sort of "constitutional" act, as it were.  Alternatively, BAFs might be artifacts of socio-psychological forces operative within a group's membership, without requiring any voluntary choice or decision on their part. The core question of interest in this section concerns the dependence (or transmission) relation between the J-statuses of a group's beliefs and the J-statuses of its members' beliefs.  Using the British Museum example, we can pinpoint a core analytical problem about J-dependence relations.

The J-status of G's belief in T may be approached from two different perspectives or dimensions: the horizontal dimension and the vertical dimension.  The horizontal dimension addresses the J-status of G's belief in T solely in terms of other beliefs of G, i.e., group-level beliefs.   G's belief in T is unjustified in terms of horizontal J-dependence because G cannot infer T by any legitimate means from

any justified beliefs <u>of its own</u>.  The situation is different, however, when we consider G's belief in T by reference to vertical J-dependence, i.e., G's position vis-à-vis its members.  Consider the members' beliefs in T and the proportion of them that are justified.   As sector Table I shows, 60% of the members justifiedly believe T.  So, given JAF-1, G's belief in T seems to be justified (because 60% of G's members justifiedly believe T).  We confront, then, something of a dilemma.  The verdict delivered by the horizontal dimension and the verdict delivered by the vertical dimensions conflict with one another.  Such a conflict cannot be tolerated in a satisfactory theory of group justifiedness.  Which dimension should take priority: the horizontal one or the vertical one?

I raise this priority question only to set it aside.  It is too large a question to settle in this paper.  What I am going to undertake instead is a more modest task.  I shall concentrate almost entirely on vertical J-dependence and seek to understand it better.  Specifically, I seek to illuminate the crucial interface between member justifiedness and group justifiedness as it arises in the problem of justification <u>aggregation</u>.  Moreover, I shall argue that an existing theory of <u>individual</u> justification -- namely, <u>process reliabilism</u> -- can be borrowed from mainstream individual epistemology and applied fruitfully to the present problem, with only a modest number of tweaks or modifications.

## 6.  <u>Two Models of  J-Dependence between Groups and Their Members</u>

The crucial question here is whether the justification dependence (J-dependence) relation between groups and their members more closely resembles <u>intrapersonal</u> J-dependence, where the J-status of a person's own states are often automatically transmitted to her other states, or more closely resembles <u>interpersonal</u> J-dependence where, lacking such automatic transmission, J-dependence typically requires acts of testimony.  The intrapersonal J-dependence model I have in mind would include (i) automatic (default) transmission of J-status from believed premises to believed conclusion, and (ii) automatic (default) retention of J-status of a memory-preserved belief from one time to another.  What I mean by automatic default transmission is that J-status transfer occurs automatically although its force can be defeated if the subject has beliefs or meta-beliefs that undercut the transmission process.  Absent such defeaters, a memory-preserved belief continues to have (roughly) the same J-status it had originally; and the J-status of an inferred conclusion is shaped (in part) by the J-statuses of the premise beliefs.  (The "in part" qualifier is required by the fact that the strength of the premises/conclusion support relation must also be factored into the equation.)

The second possibility is to interpret the J-dependence relation between groups and their members on the model of the J-dependence relations between testifiers and hearers. This view might be motivated by the List-Pettit view that groups have "minds of their own." This means that group minds are distinct from the minds of any of their members. If so, the intrapersonal model of J-dependence seems dicey and ill-motivated. The only alternative in play, moreover, would seem to be the interpersonal J-dependence relation, one person must communication his reasons to the other and hope that the hearer is suitably positioned (in terms of background beliefs) to accept this testimony justifiedly. In Schmitt's example of Tina illustrates (section 4 above), a hearer's justifiedness in believing the content of testimony depends on his <u>own</u> justified beliefs and not simply on those of the testifier. It does not automatically inherit the source's justificational state. On which model of J-dependence should the member/group justificational relationship be understood?

Here are several reasons to prefer the <u>intra</u>personal model to the <u>inter</u>personal model. First, almost all theorists agree that there is a special, intimate relation between groups and their members. A group is in some sense "composed" of its members, and members' beliefs play some sort of "constitutive" role in determining, or grounding, group beliefs. The testifier/hearer model is intuitively inappropriate for the relation between members and group.

A problem confronting the testifier/hearer model concerns the communications that would have to occur between members and group. What communication could support the production of justified belief in the group concerning the requisite justified belief states of the members? If we think about the communication as something like a vote, with each member conveying its belief vis-à-vis the proposition in question, this would fail to provide crucial information about the various members' <u>justifiedness.</u> The distribution of members' beliefs alone is uninformative about the J-statuses of those beliefs. If, on the other hand, members must inform the group of their respective <u>reasons</u> or <u>grounds</u> in addition to the beliefs themselves, the quantity of information transmitted may become prohibitive.(for the group to assimilate).

In addition, who or what would be the receiver of these communications? What exactly is the group that would "hear" or otherwise register these communicative contents? Is it the entire membership of the group? Would every member have to receive messages from all of the other members with the propositional contents as sketched in the previous paragraph? That would be a massive amount of communication. And it is hard to see how it would be executed by groups of any

reasonable size of the sorts to which we often ascribe beliefs. In short, the testifier/hearer model of justification propagation seems hopeless as a realistic model of how this might transpire.

There are weighty reasons, therefore, to prefer the first model, in which a group's J-dependence on its members more closely resembles J-dependence relations for intra-personal rather than interpersonal cognition. This is what I am calling the "transmission" model of J-dependence, and it is what I shall assume in the remainder of this paper. However, the notion of (justificational) transmission must be tweaked a bit from the notion commonly found in the standard epistemological literature. To say that member M's J-status with respect to p is automatically "transmitted" to the group is not to say, or imply, that the group's J-status with respect to p automatically becomes the same as the J-status of M's belief. It only means that M's J-status with respect to p is one component of a multi-component vector that (vertically) influences the group's J-status with respect to p. Each of the other members of the group may also have J-statuses with respect to p and these will be additional components in the same vector. What interests us is how the entire vector of members' J-statuses affects the group's J-status.

## 7. How Process Reliabilism Explains Aggregative Collective Justifiedness

To this point I have said nothing about what justification consists in, what it takes -- in "substantive" (non-epistemic) terms -- for an agent's belief to qualify as justified.[6] Next, however, I offer a familiar theory of justification that this author and others have defended in previous writings, i.e., process reliabilism (Goldman 1979, 1986, 2008, 2011a, 2011b, 2012; Kornblith, 1980; Lyons 2009, 2011; Comesana 2002, 2009; Goldberg 2010; Schmitt 1984). Process reliabilism is traditionally advanced as an account of individual justification. With only a few tweakings, however, it can be extended to group justifiedness, as the rest of this paper undertakes to show. The paper does not attempt a full-scale account of group justifiedness. Such a theory would have to cover the "horizontal" dimension of justifiedness, and no full treatment of that dimension will be offered here. The paper continues to focus primarily on the distinctive feature of collective cognition, namely, acquisition of group justifiedness from members' justifiedness (with respect to the same proposition).

The core idea of process reliabilism[7] is that a belief is justified just in case it is produced by a reliable belief-forming process -- or sequence of processes -- in the psychological history of the agent (Goldman, 1979). Exactly how reliable processes must be to confer justification is left vague, in virtue of the vagueness of the term "justified" itself. It is assumed that belief-forming processes (process

types) do not have to be 100% reliable.  More fully, justification-conferring processes can be either underlined unconditionally or underlined conditionally reliable.  A process is unconditionally reliable just in case its ratio of true beliefs to total beliefs (when operating in the relevant domain[8]) is very high. Unconditional reliability is the appropriate standard for belief-underlined independent processes, i.e., processes having no belief inputs.  Belief-independent processes might include perception and introspection.  Belief-dependent processes are ones that take beliefs as inputs to their operation.  Such processes would include inference processes and (some kinds of) memory processes.  For these processes unconditional reliability is an inappropriate standard.  A given inference type might generate a great many false conclusion beliefs because it is often applied to false premise beliefs.  This would not demonstrate any defect in the inference type.  We can define a notion of conditional reliability, however, that supplies a more appropriate criterion.  A process is conditionally reliable only if it has a high ratio of true belief outputs to total belief outputs for those cases in which the inputs to the process are (all) true.

As the phrase "belief-forming process" suggests, the heart of the theory includes the thesis that a belief token's J-status depends on how the token is causally produced, where causal production involves psychological processes.  The idea applies most easily to the underlined doxastic rather than the underlined propositional sense of justifiedness, although it can also be adapted to the latter (Goldman 1979).  In the main, the rest of this paper addresses doxastic justifiedness.  That beliefs are formed (and retained) via psychological processes is uncontroversial for individuals.  Here this idea is generalized to group, or organizational, processes.   We have already assumed that groups have psychologies (section 2.).  Group psychologies differ from those of individuals; but we can live with that difference.  Since we are already tolerating group beliefs, why not profit (theoretically) by tolerating group belief-forming processes?

A further important feature of process reliabilism is the idea that if a belief is generated by a series of intermediate beliefs -- via a sequence of inferences, for example -- the final belief is justified only if it is the product of a succession of conditionally reliable processes eachof which operates over underlined justified input beliefs (see Goldman 1979).  That these inputs must be justified, in order for an output belief to be justified, is an important requirement.  The original theory assumed that output beliefs are justified only if underlined all input beliefs to their processes are justified.  In collective epistemology, however, this requirement may need to be relaxed.  This is an instance of the inevitable need for tweaking when a theory originally designed for one domain is applied to a different domain.

With this overview of individual process reliabilism in hand, we turn now to the task of formulating and defending underlined social (i.e., underlined collective) process reliabilism, especially for the special topic of

justification aggregation.  As previously indicated, we concentrate on the vertical dimension of justification, in which J-statuses of member beliefs play a central role in determining the J-status of an aggregated group belief.  To mesh with process reliabilism, this idea must be re-construed in terms of processes.  The processes in question will take profiles of individual beliefs as inputs and return group beliefs as outputs.  We already encountered belief aggregation functions in section 5; but these aggregation functions should not be thought of as purely abstract functions.  They should be thought of as functions that get implemented by some sort of (explicit or implicit) organizational processes. The J-statuses of group beliefs will then depend on (1) the J-statuses of their member belief inputs and (2) the conditional reliability of the belief aggregation processes used.

For convenience think of the prototypical process as a (secret) balloting process that occurs at a (physical or virtual) meeting of the group.  Each member clicks a box on his/her computer screen to input her personal belief on the proposition in question.  All ballots are transmitted to the group and aggregated automatically in accordance with the pre-selected aggregation procedure.  The group belief so formed is then available for use by the group in ensuing decision-making or actional contexts.  Our question, however, concerns the justificational status of the group's belief so formed.  Although the members input only their beliefs, not the J-statuses of their beliefs (which they themselves may or may not know), each of their beliefs has a J-status, and the distribution of their beliefs plus J-statuses is what determines the J-status of the group belief.  The group belief's J-status is a product of transmission from all of the members' J-statuses.

The process of J-transmission from members to group spans, or crosses, an agential gap, a gap between pairs of agents each consisting of an individual agent and a collective agent.  Is this a tenable position for process reliabilism?  After all, under traditional process reliabilism, all belief-forming processes take place entirely within the head of an individual agent.  How can we now countenance a violation of this principle and still remain within the process-reliability tradition?

An analogous kind of "extension" of process reliabilism across agential gaps has been proposed by Sanford Goldberg (2010) for the case of testimony.  Goldberg thinks of the process of receiving (and reacting to) testimony as a process that begins in the testifier's head, proceeds across the interpersonal gap via an utterance (or inscription), and culminates in cognitive activity of a receiver.  Goldberg argues that this is fine analogue to a memory process in which a subject's earlier belief is retained in the subject's mind over some temporal interval; and as long as no "defeating" cognitive events occur during this interval, the J-status of the belief remains the same over time.  In other words, as I understand him,

it is argued that the J-status of the belief is transmitted along with the belief itself. In extending this picture to testimony between persons, he claims to be adhering to the memory analogue.

Goldberg's proposal for testimony-based justification (and/or knowledge) parallels the present proposal for group justification. In testimony the J-status of interest is the J-status of the hearer's belief. Goldberg is saying that the justifiedness of the hearer's testimony-based belief is automatically transmitted from testifier to receiver. I am not persuaded by this kind of move in the case of testimony among individuals. But I am happy to accept it for collective belief. Although members report their beliefs to the group upon voting, the J-statuses of these beliefs are not similarly reported. Nonetheless, their J-statuses are transmitted (along with J-statuses of their co-members) and automatically influence the J-status of the group's resultant doxastic attitude. In assessing the J-status of the group's attitude, we should take into account not only what is the membership's belief-profile with respect to the proposition in question but what is its justification profile. Why do I accept this transmissional picture in the case of aggregation but not in the case of interpersonal testimony? The crucial difference is that two individuals are completely independent epistemic agents whereas group agents are essentially dependent on their members. The latter dependency is a good reason, by my lights, to acknowledge transmission in the group justification-determining process.

8. Collective Justifiedness and Belief-Forming Processes

Collective beliefs' J-statuses can potentially be affected by several kinds of events, states, and processes. These can be studied with the helpful guidance of Figure 1 below.
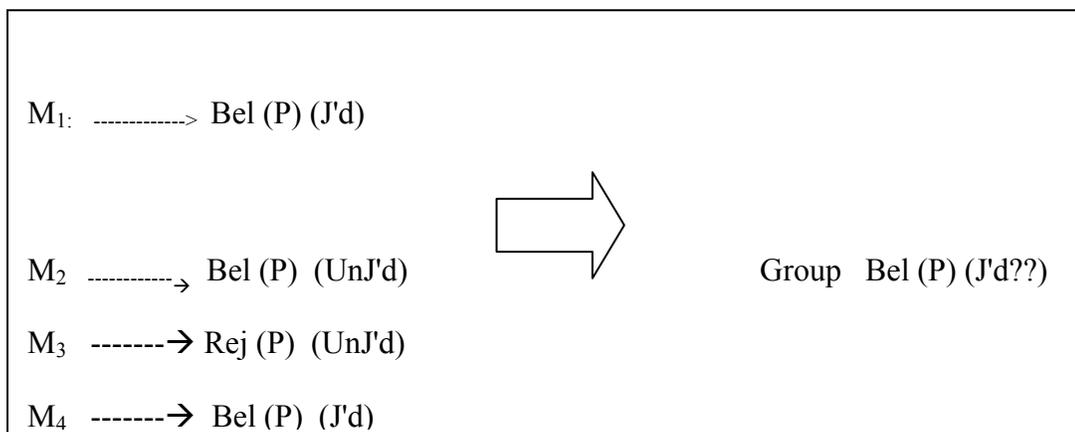
$M_1$: ------------> Bel (P) (J'd)

$M_2$ ----------→ Bel (P) (UnJ'd)

$M_3$ -------→ Rej (P) (UnJ'd)

$M_4$ -------→ Bel (P) (J'd)

Group   Bel (P) (J'd??)

Figure 1

13

The left side of the diagram depicts individual members ($M_1$ through $M_4$) who utilize various belief-forming processes (shown with dashed lines and small arrows) that result in either acceptance or rejection of proposition P.[9] Some of these resulting attitudes are justified, others unjustified. The block arrow in the middle of the diagram depicts a belief-aggregation process that runs from the profile of member attitudes to the group attitude shown on the far right. The J-status of this group belief is partly a function of the members' attitudes and their J-statuses that comprise its inputs and partly a function of the reliability of the process itself, specifically, its conditional reliability.

A key feature of process reliabilism is its historicity. Under individual process reliabilism a justification-relevant history of beliefs consists of psychological events that operate (over variable periods of time) to generate individual beliefs. Traditional epistemology was driven by an attachment to "solipsism of the moment," according to which the justificational status of a subject's believing something at time t is determined exclusively what holds of the subject at time t (especially, his/her mental states at t). By contrast, historical reliabilism says that mental or cognitive events occurring in the subject at times earlier than t can also (indirectly) influence the J-status of a subsequent belief at t. This is why chains of belief-forming processes going back in time can be relevant to a dated belief's J-status. (The historical feature of reliabilism is not fully represented in the diagram. See section 9 below.) When we turn to social process reliabilism, the relevant histories include both the sequence of individual processes that occur in each group member plus the member/group processes that give rise to group doxastic outputs. The full story would also include ("horizontal") processes occurring within a given group; but we agreed to ignore these horizontal processes for most of our discussion.

Our next task is to be more specific about the nature of the aggregation processes that span the member/group "gap" and how to assess their reliability. In the existing literature on belief (or judgment) aggregation, the standard term used is aggregation function. Can we replace the notion of an aggregation function with that of an aggregation process and produce a satisfactory theory within the framework of process reliabilism? "Functions" seem to be rather abstract and mathematical; processes are neither. How can one replace the other?

It is profitable here to re-introduce List and Pettit's notion of an "organizational structure" (section 3 above). Organizational structures, it will be recalled, are rules and procedures that groups use to enact and implement a BAF (List & Pettit, 2011: 60). Construed as linguistic entities, rules are not processes. But there are also exercises or applications of rules and procedures, and these entities might

be excellent examples of group belief-forming processes. They help to constitute an organization's "psychology." They may be conceptualized as computational operations of a group mind, operating (inter alia) on member belief profiles. Whether explicitly or implicitly, they apply procedures by which group BAFs are implemented.

With these points made, I shall now often talk as if a BAF is a belief-forming process (type). Of course, a belief aggregation function will be a belief-dependent kind of belief-forming process because the inputs to BAFs are beliefs (or doxastic states) of group members. A crucial step for epistemologists of collective justification is to establish criteria for justificationally desirable BAFs. Since reliabilist social epistemologists are interested in reliable processes, a precise criterion might be thought to require a fixed truth-ratio threshold of justifiedness. However, given the inherent vagueness of the notion of justifiedness the prospect of designating a specific threshold may be folly. Although we initially presuppose a fixed threshold in this section, a more flexible proposal will be made in section 9, where degrees of justifiedness are introduced. The central aspiration, of course, is to identify the general "shape" of group justifiedness under process reliabilism.

In the original formulation of process reliabilism two key conditions for justifiedness were presented, conditions that loom large for belief-dependent processes in particular. These two conditions can be formulated as follows (cf. Goldman, 1979).
.

> (R1)  If belief B is an output of belief-dependent process $\Pi$, B is justified only if
>     (A) all of the belief inputs to this operation of $\Pi$ are themselves justified, and
>     (B) process-type $\Pi$ is conditionally reliable.
>     (A process type is conditionally reliable IFF for a selected threshold T (> .50), at least
>     T percent of the cases in which all input beliefs to the process are true are also cases
>     in which the output belief is true.)

Conditions (R1) states two necessary conditions for individual justifiedness, applied to outputs of belief-dependent processes. Condition (R1)(A) is a justification transmission condition. It says, in effect, that premise beliefs must "contain justifiedness" in order to transmit justifiedness to the conclusion belief. Otherwise, where does the justificational "juice" come from? Moreover, in the case of inference it is plausible that an epistemic agent's conclusion belief is justified only if all premise beliefs of the inference are justified for the agent. If we introduce a counterpart condition for collective

belief aggregation, we need to consider exactly what that counterpart should be.  Similarly, if we are to find a collective-domain parallel for the conditional reliability requirement of condition (R1)(B), we should specify -- or at least point in the direction of -- a suitable analogue of the conditional reliability requirement in the individual domain.  These tasks are addressed in the next section.

9.  <u>Permuting Process Reliabilism for Collective Epistemology</u>

A straightforward analogue of condition (R1)(A) for cases of group belief based on aggregation, might be $(R1_G)(A)$:

$(R1_G)(A)$   If a group belief B in P is generated by aggregating a profile of member attitudes vis-à-vis P, belief B is justified only if all members' beliefs in P belonging to that profile are themselves <u>justified.</u>.

On reflection, however, should it be required that <u>all</u> beliefs (inP) by the members be justified?  What warrants such a strong requirement?  In the case of inferential justification, every premise used in the inference may be essential to drawing the conclusion, at least to make the conclusion highly likely.  The falsity of a single premise might substantially lower the likelihood of the conclusion's truth.  So we may reasonably hold the conclusion belief to be unjustified if any premise belief lacks justification.  But if a group has a large number of members, a large proportion of whom justifiedly believe the target proposition P, is it really necessary that <u>all</u> such members believe P justifiedly in order for the group to believe P justifiedly?

For example, suppose that group G has 100 members, of which 60 believe P, 30 reject P, and 10 withhold judgment.  Further suppose that under the governing BAF this attitude profile suffices to generate a belief in P by group G.  Now suppose that 58 of the 60 members who believe P do so <u>justifiedly</u>, whereas none of the 30 P-rejections is justified.  How should the <u>group's</u> P-belief be assessed in justificational terms?  According to condition $(R1_G)(A)$, justifiedness is denied to G's P-belief because not all the members' beliefs in P are justified.  This seems grossly over-demanding.  Why couldn't some of them be considered "redundant"?  Why isn't it sufficient for 58 of the 60 believers to be justified in believing P?

Until now we have been assuming a categorical, or binary, notion of justifiedness. That's how the theory of epistemic justification is usually discussed. A belief is either justified or unjustified and the problem is to specify the conditions for each. On reflection, however, maybe this problem is rather artificial. To the extent that we use the notion of justifiedness, or reasonableness, it is much easier to think of it as a matter of degree, a gradable notion rather than a categorical one. And this gradability approach seems more tractable than the categorical one when we are addressing groups with multiple members. Such groups will rarely display unanimity, and there is no obvious way to set a non-arbitrary threshold that would be essential for a categorical notion. Of course, the categorical notion of justification has been popular because of its presumed utility in defining knowledge ("S knows p only if S is justified in believing P"). But recent epistemologists have been attracted by degrees of justification even in their accounts of knowledge. Contextualists, for example, claim that the level of justification required for knowledge attribution varies with the conversational context (DeRose, 1992, 2009; Cohen, 1986). Within this tradition it is highly desirable to have a story about grades of justifiedness. If this holds for individual justifiedness, it might be even more imperative for group justifiedness, where the expected variation in groups inevitably makes categorical notions more difficult to employ.

I shall not attempt to sketch a full theory of justificational gradability.[10] But a few sample principles are adduced to convey the flavor of what seems appropriate. Here I assume that members' doxastic attitudes have categorical J-status, and this assumption will be used in formulating plausible principles of graded J-status for group belief. Here is one such principle (which introduces two factors):

(GJ) If a group belief in P is aggregated based on a profile of member attitudes toward P, then (ceteris paribus) the greater the proportion of members who justifiedly believe P and the smaller the proportion of members who justifiedly reject P, the greater the group's level, or grade, of justifiedness in believing P.[11]

I further propose that J-status "facts" about the memberships' justificational profile vis-à-vis P are automatically transmitted from that profile vis-à-vis P to the group's J-status level with respect to P. In other words, suppose the group forms a belief in P based on a profile of attitudes held by its members. The group's belief output is a function of its BAF applied to the membership's attitude profile. But the membership's attitude profile per se tells us nothing about the J-statuses of these member attitudes, and nothing about the J-status of the group belief. The J-statuses of the members' attitudes will depend on

the processes by which they severally arrived at their respective attitudes.  Similarly, the J-status of the group belief will depend on the members' attitude profile and on the J-statuses of those attitudes.

Notice that although the J-status of a group belief depends on the J-statuses of its members' beliefs, the group need not be aware of the J-statuses of the several members' attitudes.  The group belief justificationally "profits" or "suffers" from the J-statuses of its members' attitudes because these statuses are transmitted to the group's J-status.[12]  To take an example, the J-status of the Bush Administration's belief in the existence of WMD (in the hands of the Iraqi government) depended substantially on whether the technical experts who contributed to the Administration's opinion themselves had individual justification.[13]

I turn now to the second necessary condition for the justifiedness (or a relatively high degree of justifiedness) of a group belief formed by belief aggregation.  This is the conditional reliability condition appropriate for an aggregation process, which can be stated as follows:

(R2$_G$)(B)  A group belief G that is generated by the operation of a belief-aggregation process
$\Pi$ is justified only if (and to the degree that) $\Pi$ has high conditional reliability.

Conditional reliability as a requirement of justifiedness was originally inspired by two thoughts.  The first is the guiding theme of process reliabilism, viz., that truth-conduciveness of the process(s )used is a key to the justifiedness of a token belief.  For example, beliefs formed by perception (in normal conditions) are usually justified because perceptual processes are generally reliable processes.  By contrast, beliefs formed by wishing thinking or idle speculation are unjustified because these kinds of belief-forming methods are generally unreliable.  However, if truth-conduciveness is the main motif of process reliabilism, why should a theorist retreat to conditional reliability when it comes to belief-dependent processes like inference?  The rationale is straightforward.  When it comes to inferences (like many other things), "garbage in/garbage out" is an apt motto.  If you start with false premises, you cannot expect inference to deliver true conclusions.  In short, one cannot reasonably expect an inferential process to yield true beliefs as a general rule if it frequently starts with false premises.  The only "fair" test of an inferential process is how often it delivers true conclusion beliefs when it is fed true premise beliefs to start with.  This is what a conditional reliability condition requires, in the form it was originally proposed (Goldman 1979).  Why shouldn't the same, or a similar, idea apply here?  The requirement that an inference process map truths into truths -- more precisely, true beliefs into true

beliefs -- looks like a fair, not unduly demanding condition, yet still a substantial one, which can exclude the riff-raff among inferential processes. So it seems natural to apply it to the belief aggregation domain as well.

There are some differences, however, between inferential processes and belief aggregation processes. In the inference case a process would commonly start with a number of premise beliefs and end with a <u>different</u> conclusion belief (of the same agent). In the case of belief aggregation, by contrast, there is a potentially large number of individual agents with different attitudes toward a specified proposition plus a group agent that adopts one of these attitudes toward the <u>same</u> proposition. Achieving true-belief preservation here cannot be a matter of simply conforming to logical connections among propositions. A good BAF must be one that somehow "sees to it" that those group members who believe the truth, or are likely to believe the truth (given the route by which they arrived at their beliefs), have a suitable amount of influence on the group's attitude selection so that the group also believes the truth. This is what will capture the kind of conditional reliability appropriate to the belief aggregation domain. In short, for purposes of a reliabilist theory of justification for belief aggregation, conditional reliability should be interpreted not in terms of a requirement that <u>all</u> input beliefs of the members be true, but rather in terms of (enough) input beliefs having a <u>high probability</u> (or chance) of being true. As it happens, this might neatly fit the template presented by the Condorcet Jury Theorem.

Majoritarian voting is a widely touted belief aggregation process. Although voting is not the same as believing, it is instructive to consider majoritarian belief aggregation in the context of the well-known Condorcet Jury Theorem (CJT). CJT is often seen as a mathematical rationale for majority decision-making. Here I mean to observe that this rationale can also be interpreted in terms of <u>conditional reliability</u>. It nicely illustrates what might qualify as a good belief aggregation process because it approximates what a reasonable conception of conditional reliability might be for the domain of belief aggregation.

Details aside, CJT says that majority voting is a conditionally reliable aggregation procedure in that, <u>conditional</u> on voters' having a uniform competence greater than one half (e.g., prob = 0.52) for believing the truth about a target question, a group belief that emerges from a majoritarian aggregation procedure has an even higher probability of yielding the truth. Under a majoritarian aggregation rule, the group's probability of believing the truth swiftly approaches 1.0 as the group size increases. This is the sense in which a majoritarian process of belief aggregation is conditionally reliable.[14] Where the

condition of uniform member competence is satisfied, the majoritarian process reliably leads to a high probability of the group belief's being true.

The reader might wonder why majoritarian aggregation is being credited only with <u>conditional</u> reliability. Isn't that an affront to its strength? Readers swept away by the putative "wisdom of the crowd" thesis, for example, might contend that a large group of independent individuals are <u>unconditionally</u> disposed to get the truth. Why is this special power of group decision-making being derogated by insistence on the qualifier "conditional"? The reason is simple. Majoritarian belief aggregation is not unconditionally reliable, that is, reliable under all conditions. In fact, it is as much conditionally <u>anti-</u>reliable as it is (positively) conditionally reliable. Specifically, if group members are uniformly <u>in</u>competent, i.e., have a probability less than half of (independently) believing the truth, the probability of a majority-driven group believing what is true rapidly approaches <u>zero</u> as group size increases. This is why it would be excessive to require (as a condition of justifiedness) that a belief aggregation procedure be <u>un</u>conditionally reliable. Perhaps no such procedure is even possible. In any case, conditional reliability seems like an eminently appropriate requirement.

However, several addenda and qualifications must be entered here concerning the CJT and its applicability to our problem. First, in addition to the requirement that all members have a probability greater than half of being right, CJT also requires that the members' competences be independent of one another. Second, the original formulation of CJT is restricted to cases where there are only two options on the table. List and Goodin (2001) generalized the theorem so that it applies to cases involving more than two options, but it is still an open question whether this helps us in our problem area. Why should more than two options be a potential sticking point for our problem? Aren't we discussing cases where the truth or falsity of a single proposition is on the table? Doesn't this already restrict our problem to two options, unlike cases where, for example, voters choose from a list of three or more candidates?

No, the mere fact that we are dealing with true-or-false propositions does not eliminate the problem. Our epistemic agents, being doxastic decision makers, confront at least three possible <u>attitudes</u> to adopt, not two, viz., belief, disbelief, and suspension of judgment (agnosticism). In the CJT literature it is ordinarily assumed that each option on the table has some chance of being "right" or "correct". But no meaning can be attached to the idea of suspension, or agnosticism, being the "right", "correct", or "true" option. These are among the reasons why the CJT is not a perfect illustration of a conditionally reliable belief aggregation process. It is still a helpful illustration, however.

Of course, the conception of conditional reliability it introduces is not exactly the one introduced in connection with individual process reliabilism. In the original notion of conditional reliability, reliability is conditioned on the truth of all premise beliefs. In the CJT case reliability is conditioned on each group member possessing an objective chance (greater than half) of being right (plus these chances being held independently of one another). Although these are slightly different conceptions of conditional reliability, they are close enough to be members of the same family. This strikes me as sufficient to view the CJT approach to epistemic goodness as belonging to the same family as that of the conditional reliability approach to individual process reliabilism.

A final tweaking of the process reliabilist theory needs to be undertaken in the context of justified group belief based on aggregation. The need for this tweaking is also reflected in the CJT. One of the most important qualifications included in the CJT (mentioned earlier) is that members' beliefs (or judgments) be independent of one another. For example, if they all form their respective beliefs based on the same evidence, their belief-forming processes would not be independent. Hence, one cannot expect the group's likelihood of being right to be "boosted" in the same fashion as it would if all have entirely separate bodies of evidence. This is why CJT is correct only if it is subject to the independence requirement. This factor cannot be ignored in the present process reliabilist approach either.

Isn't this a serious threat, however, to the viability of process reliabilism in this terrain? Isn't it well known that reliabilism and evidentialism are entirely distinct -- indeed, opposed -- epistemological approaches? How can we contemplate incorporation of an evidentialist component in our theory and still label it "process reliabilism"? In fact, however, a synthesis of reliabilism and evidentialism is not unprecedented. I offered such a synthesis -- and a rationale for it -- in a recent paper (Goldman 2011b). I also pointed out that viewing states of mind as evidence (or states of evidence possession) is not foreign to the spirit of reliabilism, although it was never mentioned explicitly in early statements of reliabilism.[15] When one considers that belief-forming processes need mental states as inputs, and those input states might be viewed as items of evidence or evidence possession, the notion that there is a fundamental gulf or opposition between process-elements and evidential elements evaporates. They can be quite complementary. However, fusing evidential factors and process factors into a unified theory of justifiedness is no simple task, and it will not be undertaken here. I flag the importance of evidential elements, however, because they cannot be ignored even within a predominantly process-oriented approach. For a more systematic treatment of evidence in the context of group justifiedness, see Lackey (forthcoming).

It is time to complete the general contours of (aggregation-based) group justifiedness by saying more about the left-most portion of Figure 1. This part of the diagram is only suggestive of the larger story. It depicts group members as forming beliefs in P and it labels some of those beliefs as justified and others as unjustified. What determines their J-statuses? In each case there will be a story about the cognitive route by which an individual arrives at his/her belief, some of them a series of reliable belief-forming processes and others a series of processes with mixed reliability. It could be that all or some of the members formed his/her beliefs "individualistically," without interacting with any group members. Other members may have had extended conversations or debates about P with other members before arriving at their current doxastic attitudes. Such activities should be considered parts of a <u>deliberative</u> process (among the members), but not parts of the group aggregation procedure that generates a <u>collective</u> doxastic attitude. These conversations and individual belief-forming processes are all parts of <u>pre-aggregation</u> processes. Thus, the full history of justificationally-relevant activities that culminate in a collective belief (with a particular J-status) is not exhausted either by the process depicted by the block arrow at the center of Figure 1 nor by the individual processes depicted at the left of the diagram. The full history would go back even further, requiring more space than the page allows. The full history would be a leftward spreading tree, with (finitely) many additional nodes and branches. This provides a glimpse of the reason that I call process reliabilism a <u>historical</u> approach to justifiedness. .

10. <u>Other Kinds of Reliable Aggregation Process Types</u>

Let us now depart a bit from purely majoritarian aggregation processes and consider related belief aggregation processes discussed by philosophers and political theorists. Many of these approaches also view reliability, or truth conduciveness, as the fundamental epistemic aim. And although they rarely if ever mention <u>justification</u> as their targets of analysis, they strike me as kindred spirits to process reliabilism, at least with respect to belief aggregation.

Among the methods proposed for optimal reliable belief aggregation functions are ones involving the deployment of <u>experts,</u> who are given more than the customary amount of influence or weight in the aggregation process. In the usual majoritarian schemes, group members are all assigned equal weight in influencing a group choice. In other schemes, however, the system makes special use of experts to boost truth-getting properties of the group.[16] Experts are understood (by definition) to have greater than average truth-getting competence, at least in their areas of expertise (Goldman, 2001). This

gives rise to systems in which the membership possesses <u>heterogeneous</u> levels of competence as opposed to a <u>homogeneous</u> level.

Bradley and Thompson (2012: 68-69) lay out several possible schemes in which even heterogeneous levels of competence across a membership can have epistemically superior properties for a group. What they call "oligarchic" schemes are voting rules that count only the votes of the most reliable members. A second category include schemes in which everyone's vote is counted but weighting rules assign different weights to different voters according to reliability. The oligarchic category includes expert majority rule, whereby a proposition is accepted by the group if it is accepted by a majority of its most reliable members. The weighting-rules category assigns weights to voters as a function of their respective competences or the log of their competences (Grofman et al., 1983). Then the weighted sum of all of the votes determines the group belief. Under certain assumptions, e.g., independence, any of these procedures yields a higher probability of a group belief's being true than a (rival) majority rule and/or an equal-weight system.

Returning to the issue of justifiedness, here is a key question. Suppose that one of these expertise-sensitive systems is adopted by a committee or political unit, and the designated experts really are more reliable than their co-members. Does the use of such a reliability-enhancing system generate greater (prima facie) <u>justifiedness</u> of a group belief than a less reliable system? I would say that it does generate greater justifiedness. In the preceding cases, certain facts pertaining to the membership's J-status profile -- including the kinds of evidence possessed by assorted members -- automatically influences the J-status of the group belief in a proposition P. Analogously, the facts that "ground" the reliability property of the expert-oriented system contribute positively -- and automatically -- to a positive J-status for the group's belief in P. It is automatic in the sense that the group itself need not have any belief about the correctness of the choice of this or that individual as an expert. Such higher-order beliefs are not needed for the system to be highly reliable and only high reliability is required to influence the J-status of the group's belief in P. This is quite consonant with the <u>externalist</u> character of process reliabilism.

Of course, if the group has "defeater" beliefs to the effect that some choices of people as (putative) experts are misguided, then such beliefs might lower the J-status of the group's belief in P. But it would not cancel the positive and automatic contribution of the de facto reliability of the weighted voting system. The defeater beliefs just compete with the de facto reliability of the system in influencing the <u>ultima facie</u> justifiedness of the group's P-belief. Finally, in saying that a group would

(ceteris paribus) have greater justifiedness in believing P based on a weighted voting scheme than on a (less reliable) equal-weight voting scheme, this pertains only to epistemic justifiedness, not political (or moral) justifiedness. There may be grounds to object to an "epistocracy" on moral or political grounds, but the present question concerns epistemic character only.


11.   Some Residual Issues


I conclude this essay with remarks on further issues that cannot be addressed in depth but should not be wholly neglected. The general burden of this section is to support the claim that process reliabilism is a fruitful approach to collective epistemology despite certain challenges that need to be met.

The first challenge is the question of whether process reliabilism is applicable to all kinds of groups to which collective beliefs are ascribed.   How does it apply, for example, to loose assemblages of individuals that might also be considered "plural subjects"? Are there causal processes that link doxastic profiles of member beliefs to collective beliefs, processes that possess a determinate level of conditional reliability?[17]  Consider all the people currently situated in Times Square, New York. It's a hot day and they are all thinking it is hot. So we might describe the assemblage as having a collective belief, "It is hot." Is there an aggregation procedure that <u>causes </u>such a plural-subject belief to be held? This seems doubtful.

In this case, however, it is equally doubtful that the Times Square assemblage qualifies as a genuine group, or the subject of a collective belief "It is hot". This is what has been called a merely "summative" ascription to a group (Quinton 1975-76), not a genuine collective attitude. Since there is no group attitude, the question of group justifiedness does not arise.

But suppose the Times Square assemblage "deliberates" with one another, now on another subject. Suppose that a number of them spot a flying entity high up in the sky, and several cry out, "It's an alien spaceship." Others are soon prompted to repeat the cry, "Look, an alien spaceship," and by a process of contagion a large majority of them form the belief that an alien spaceship is approaching. Here the assemblage acquires the character of a unified (or semi-unified) group, and more plausibly has a genuine group belief. But is a <u>belief aggregation function </u>used to produce the belief? By assumption this group formulates no formal belief aggregation rule, so how does our theory apply to it?

This kind of problem was already anticipated in section 3, where we broadened the notion of a BAF to include "functionally inexplicit organizational structures." As indicated there a BAF does not require a <u>formula</u> or formally instituted <u>procedure</u> that a group follows. It can instead be a dynamic pattern of interpersonal influence by which beliefs spread from some members to others. This might be considered a variant of the member-group aggregation process that constitutes a BAF for purposes of justification appraisal. And there are certainly causal processes at work there.

Alternatively, the pattern of dynamic interpersonal influence may be conceptualized as a <u>pre-aggregation</u> step on the road to group belief. In Figure 1 the left side depicts members forming individual beliefs by means of causal processes. These processes might not be wholly internal to the members. There might also be interpersonal processes involving contagion. There still needs to be a group aggregation process, even under this option, but the contagion processes would not comprise the aggregation process. The aggregation process would be some species of metaphysical event -- composition or constitution -- that generates a group belief without causing it.

Can process reliabilism live with the suggestion that some <u>links</u> or <u>phases</u> in the total process of group formation are <u>non-causal</u> phases? Why not? Why can't process reliabilism concede that in the alien spaceship example, individuals in Times Square communicate with one another and causally persuade a majority of them to believe that an alien spaceship approaches. At this juncture, with the group having formed a suitable "unity" (however that happens), a metaphysical step of <u>composition</u> or <u>constitution</u> occurs in which the numerous individual beliefs generate a collective belief. Process reliabilism can accept the presence of a non-causal aggregatone link because it is only one link in a series of processes most of which are causal. This leaves ample room to appraise the collective belief by means of the reliability of the historical process that leads to it. Given these various options, process reliabilism need not fret unduly over cases involving collective beliefs by loose assemblages of people. In any event they are only borderline or limiting cases of collective beliefs that are candidates for justificational appraisal. If they are not a perfect fit with the mold of the best candidates for such appraisal, this should be neither a surprise nor a serious concern.

I turn next to a different issue concerning the application of process reliabilism to collective-level epistemology. The focus thus far has been the vertical dimension of group belief formation. What about the horizontal dimension, in which inference or reasoning occurs within a group mind? Presumably, social process reliabilism implies that the J-statuses of collective inferential beliefs are a function of the belief-forming processes that produce them. Is this applicable at the level of collective

belief?  One question is whether a causal-process model of collective propositional attitudes is really appropriate.  A second question is whether -- assuming such a model is appropriate -- it is the causal processes that are specifically relevant to the J-statuses of the collective beliefs they generate.  The second of these questions will be briefly addressed here.

I argue that the J-status of an inferential belief does depend on the properties of the causal process responsible for generating it.[18]  I shall use an example of an individual cognizer, but the point readily generalizes to groups -- at any rate, groups capable of inference or reasoning.  Suppose that Chad justifiably believes two propositions of the form "P or Q" and "not-Q".  Suppose further that he uses these beliefs to infer a belief in P.  However, he doesn't use the "process" of disjunctive syllogism to draw this inference.  Instead, he uses a process that overgeneralizes disjunctive syllogism.  It is a psychological propensity of his to view as valid any inference from two premises of the form "P $ Q" and "not-Q" to a conclusion of the form "P", where "$" is replaceable by any binary truth-functional connective  (disjunction, conjunction, material conditional, etc.).  If this is the process he uses, intuitively his conclusion belief is not doxastically justified (although somebody else could have used the same premises to arrive at the same conclusion justifiedly).  Chad's problem is his use of a flawed belief-forming process.  His "dollar-sign" inferential process is not conditionally reliable.  Applied to a range of easily constructible examples (with true inputs), its set of outputs would have a rather low truth-ratio.  Clearly, it is this conditionally unreliable process that explains the unjustifiedness of his belief in P.  What else might be driving our intuition here?

Could a group epistemic agent exemplify the same case?  Why not?  Of course, there must be machinery, whether human, silicon, or what have you, to ground the group's hypothetical mimicking of Chad.  But as long as we understand the way it works (which computations are being executed, if you will), we will deem the inference type to be a bad one.  And conditional unreliability of the inference pattern is an obvious rationale for dismissing it as defective.

I turn next to memory.  Assume that epistemologists like me are right to think that a transmission principle is commonsensically used to impute a J-status to a belief at time $t_n$ that duplicates the content of a belief held earlier, at time $t_o$.  The question is: Is there an underlying process assumption that underlies this memory transmission principle?  If we were given a case in which no retention process is at work, would we respond differently?  And might this be detectable in a group agent example?

A good test case would be one in which there are accidental content replacements in the same "container".  In a group agent case, a relevant scenario would be one in which individuals who represent

the group (in important ways) are continually leaving or joining the group in such a fashion that the group never ceases to have a certain belief content represented but there are no causal links between their being so represented; it's just a matter of chance. Intuitively, I suggest, this would not be regarded as memory. What would be missing is an operative process that reliably maintains the content "in place" from moment to moment. This is how we conceptualize memory, I suspect, and the notion of process is central to it.

Here is a final thought in a different vein, concerning the feature of historicity distinctive to process reliabilism. A critic might complain that the historical approach to justification embedded in process reliabilism is unsatisfactory because it includes the cognitive biographies of individual epistemic agents, distinct from the group agent whose J-properties are in question. It is fine to track the J-status of a belief backwards as long as one stays within the target agent whose belief is in question. Going any further back, however, would be going too far! Yet that is precisely what we are proposing for the J-assessment of collective beliefs. We are proposing (in Figure 1) that a collective belief's J-status depends in part on the history that takes place within the members but outside the group. This is surely wrong, contends the critic.

I can understand and appreciate this sentiment; occasionally I feel it myself. But "classical" process reliabilism already made room for precisely this ecumenical sentiment. "What Is Justified Belief?" already distinguished between two forms of reliabilism: underline{historical }and underline{terminal-phase} reliabilism (Goldman, 2012: 42-43). As the label suggests, terminal-phase reliabilism makes a J-status assessment depend on how well the agent performs given the (evidential) resources available at the underline{final phase }of belief formation. Historical reliabilism, by contrast, highlights how well the agent proceeds at underline{all }of the decision points underline{leading up to }as well as including the final phase. Each approach has some claim to being a satisfactory explication of "justified," where justifiedness is construed broadly as a state that results when one proceeds appropriately given one's evidence, perspective, or starting point. An assessor's choice of a starting point can be flexible, leaving room for both the historical account and the terminal-phase account.[19]

## References

Bird, A. (2010).  "Social Knowing: The Social Sense of 'Scientific Knowledge'."  Philosophical Perspectives 24: 23-56.

Bradley, R. and Thompson, C. (2012).  "A (Mainly) Epistemic Case for Multiple-Vote Majority Rule."  Episteme 9(1): 63-80.

Bratman, M.E. (1999).  Faces of Intention.  Cambridge: Cambridge University Press.

Cohen, S. (1986).  "Knowledge and Context."  Journal of Philosophy 83: 574-583.

Comesana, J. (2002).  "The Diagonal and the Demon."  Philosophical Studies 110: 249-266.

Comesana, J. (2009).  "What Lottery Problem for Reliabilism?"  Pacific Philosophical Quarterly 90(1): 1-20.

Comesana, J. (2010).  "Evidentialist Reliabilism."  Nous 94(4): 571-600.

DeRose, K. (1992).  "Contextualism and Knowledge Attributions." Philosophy and Phenomenological Research 52: 913-929.

DeRose, K. (2009).  The Case for Contextualism.  Oxford: Oxford University Press.

Fine, K. (1994).  "Ontological Dependence."  Proceedings of the Aristotelian Society 95: 269-290.

Gilbert, M. (1989).  On Social Facts.  Princeton: Princeton University Press.

Gilbert, M. (1994).  "Durkheim and Social Facts."  In W. Pickering and H. Martins (eds), Debating Durkheim (86-109).  London: Routledge.

Gilbert, M. (2002).  "Belief and Acceptance as Features of Groups," Protosociology 16: 35-69.

Goldberg, S.C. (2010).  Relying on Others.  New York: Oxford University Press.

Goldman, A.I. (1979).  "What Is Justified Belief?"  In G. Pappas (ed.), Justification and Knowledge. Dordrecht: Reidel.  Reprinted in Goldman (2012).  All page citations for this article refer to the 2012 reprinting.

Goldman, A.I. (1986).  Epistemology and Cognition.  Cambridge, MA: Harvard University Press.

Goldman, A.I. (2008).  "Immediate Justification and Process Reliabilism."  In Q. Smith (ed.), Epistemology: New Essays (63-82).  New York: Oxford University Press.  Reprinted in Goldman (2012).

Goldman, A.I. (2011a).  "Reliabilism."  The Stanford Encyclopedia of Philosophy (Spring 2011 edition), E.N. Zalta (ed.), http://plato/stanford.edu/entries/reliabilism/.  Reprinted in Goldman (2012).

Goldman, A.I. (2011b).  "Toward a Synthesis of Reliabilism and Evidentialism."  In T. Dougherty (ed.), Evidentialism and Its Discontents (254-280).  New York: Oxford University Press.  Reprinted in Goldman (2012)

Goldman, A.I. (2012).  Reliabilism and Contemporary Epistemology. New York: Oxford University Press.

Grofman, B., Owen, G. and Feld, S. L. (1983).  "Thirteen Theorems in Search of the Truth," Theory and Decision 15: 261-278.

Huemer, M. (1999).  The problem of memory knowledge.  Pacific Philosophical Quarterly 80: 346-357.

Kornblith, H. (1980).  "Beyond Foundationalism and the Coherence Theory." Journal of Philosophy 72: 597-612.

Lackey, J. (1999).  "Testimonial Knowledge and Transmission."  Philosophical Quarterly 49: 471-490.

Lackey, J. (2008).  Learning from Words: Testimony as a Source of Knowledge.  Oxford: Oxford University Press.

Lackey, J. (this volume).  "A Deflationary Account of Group Testimony," in J. Lackey (ed.), Essays in Collective Epistemology.  Oxford University Press.

Lackey, J. (forthcoming).  "What Is Justified Group Belief?"

List, C. and Goodin, R.E. (2001).  "Epistemic Democracy: Generalizing the Condorcet Jury Theorem," Journal of Political Philosophy 9: 277-306.

List, C. and Pettit, P. (2005).  "On the Many as One," Philosophy and Public Affairs 33(4): 377-390.

List, C. and Pettit, P. (2011).  Group Agency.  Oxford: Oxford University Press.

Lyons, J.C. (2009).  Perception and Basic Beliefs.  New York: Oxford University Press.

Lyons, J.C. (2011).  "Should Reliabilists Be Worried About Demon Worlds?" Philosophy and Phenomenological Research.

Mathiesen, K. (2006).  "The Epistemic Features of Group Belief." Episteme 2(3): 161-176.

Pettit, P. (2003).  "Groups with Minds of their Own."  In F.F. Schmitt (ed.), Socializing Metaphysics (467-493).  New York: Rowman and Littlefield.

Quinton, A. (1975-76).  "Social Objects," Proceedings of the Aristotelian Society  75: 1-27.

Schaffer, J. (2009).  "On What Grounds What."  In D. Chalmers, D. Manley, and R. Wasserman (eds.), Metametaphysics (347-383).  New York: Oxford University Press.

Schmitt, F.F. (1984).  "Reliability, Objectivity, and the Background of Justification." Australasian Journal of Philosophy 62: 1-15.

Schmitt, F.F. (1994). "The Justification of Group Beliefs." In F.F. Schmitt (ed.), <u>Socializing Epistemology: The Social Dimension of Knowledge</u> (257-287). Lanham, MD: Rowman & Littlefield.

Schmitt, F.F. (2006). "Testimonial Justification and Transindividual Reasons." In J. Lackey and E. Sosa (eds.), <u>The Epistemology of Testimony</u> (193-224). Oxford: Oxford University Press.

Searle, J. (1995). <u>The Construction of Social Reality.</u> New York: Free Press.

Sylvan, K.L. (2012). "How to Be a Redundant Realist." <u>Episteme</u> 9(3): 271-282.

Tollefsen, D. (2007). "Group Testimony." <u>Social Epistemology</u> 21: 299-311.

Tuomela, R. (1992) "Group Beliefs." <u>Synthese</u> 91: 295-318.

Tuomela, R. (2007). <u>The Philosophy of Sociality: The Shared Point of View.</u> Oxford: Oxford University Press.

Wray, K. B. (ed.) (2010). <u>Collective Knowledge and Science,</u> special issue of <u>Episteme</u> 7(3): 182-283.

Wright, C. (2002). "Anti-Skeptics Simple and Subtle: Moore and McDowell." <u>Philosophy and Phenomenological Research</u> 65: 330-348.

Wright, C. (2003). "Some Reflections on the Acquisition of Warrant by Inference," in S. Nuccetelli (ed.), <u>New Essays on Semantic Externalism, Skepticism, and Self-Knowledge.</u> Cambridge, MA: MIT Press.

[1] List and Pettit (2011) use the term "judgment" rather than "belief." The trouble with "judgment," however, is that its referent is ambiguous as between a speech act and a psychological state. Although speech acts can be evaluated in epistemic terms, as are beliefs, the criteria for speech act evaluation seem quite different than those for belief evaluation. So I shall generally avoid the term "judgment."

[2] I will not be assuming that a group belief in a proposition requires that all members of the group believe it. Thus, my examples will not fit with Gilbert's (1989) model. Indeed, although Gilbert uses the language of "group belief", her "joint acceptance" view is best seen as a theory of group action rather than group belief, because "acceptances" in her treatment are verbal expressions (see 1989, p. 306).

[3] I set aside here problems about "transmission failure" of the sorts that preoccupy Crispin Wright (2002, 2003) and others.

[4] For one thing it is possible to produce examples in which JAF-1 would assign positive J-status for group beliefs in each of an inconsistent set of propositions. This might well preclude its acceptability as a correct JAF. Here is such an example provided by David Christensen (personal communication). Continuing with the British Museum example, suppose that each group of guards has done some elementary reasoning and are justified in the results. Assume that the 40 guards in the remaining group (who don't suspect A, B, or C) justifiedly believe that no guard is planning to steal. As before, let proposition T = "At least one guard is planning to steal. Since T is justifiedly believed by each member of the first three groups of guards, this is a total of 60 guards, so the group justifiedly believes it. Next consider another proposition, "T $\rightarrow$ (B v C)" (where the arrow represents the material conditional). This proposition is justifiedly believed the 40 guards who suspect B and C. It is also justifiedly believed by the 40 guards who don't believe that anyone is planning a theft. So it is justified believed by 80 guards. So the group justifiedly believes this as well. Finally, consider a third proposition, "~(B v C)". This is justified believed by the 20 guards who suspect A and by the 40 guards who believe that no one will make a theft. This makes sixty guards, so the group justifiedly believes this third proposition as well. Hence, the group justifiedly believes an inconsistent set of proposition, when JAF-1 is used. This result may well demonstrate the unacceptability of JAF-1. (However, in light of the many impossibility results associated with judgment aggregation principles -- see List and Pettit, 2011 -- we should not reject such principles too hastily, even ones with unattractive properties).

[5] What this implies is that what are called "justification aggregation functions" in the middle sections of the paper are really best viewed as conditional justification aggregation functions. They indicate when a group belief is justified conditional on the J-statuses of the input members beliefs. The aggregation function or process at the "terminal phase" of a group process of belief formation does not all by itself determine

the J-status of the output group belief.  A J-status for the group belief is partly a function of the JAF used but also partly a function of the J-statuses of the inputted member beliefs.

[6]  Throughout I confine the discussion to "doxastic" justifiedness rather than "propositional" justifiedness,.  Presumably the theory on which I shall focus has decent ways to explain propositional justifiedness in terms of doxasstic justifiedness.  At least so I proposed in Goldman (1979).

[7]  Like almost all theories of justifiedness, the distinctive parts of process reliablism must be supplemented with a component that takes account of defeaters.  We will turn to defeaters in due course.  To get started, however, the need to handle them will be set aside.

[8]  Which is the relevant domain is a question that confronts reliabilism.  It might be the actual world (past, present and future), it might be a selected portion of environments in the actual world, or it might be some class of possible worlds more inclusive than the actual world.  I don't try to resolve this question here.

[9]  Although the processes are "individual," this does not preclude their being related to other people.   Indeed, the processes might take as inputs the reception of testimonial acts of others, so they could be embedded in larger, deliberative processes.  The important point is that these processes, represented by dashed arrows, are not "vertical" processes, which support a mapping from individual doxastic attitudes to group doxastic attitudes.  Only the block arrow in Figure 1 represents a vertical process.

[10]  Actually, degrees of justification are very easily handled in process reliabilist terms.  Ceteris paribus, the higher the truth-ratio of a belief-forming process, the greater the justifiedness of the beliefs it generates.   However, this approach presupposes the correctness of process reliabilism, so doubters of this approach might not be cheered by this account.

[11]  This principle and the one that follows both use a categorical notion of justifiedness (with respect to members' beliefs) and a graded notion of justifiedness (with respect to group belief).  It is arguable that this is objectionable, and if so new principles would have to be devised that employ graded justifiedness throughout.  I leave this task for another day, or for others.

[12]  Of course, if a group does happen to have a higher-order belief about these features of the membership's justificational profile, such a higher-order belief would influence the J-status of its belief in P.  But this does not conflict with the transmission of the membership's J-profile.  Transmission does not imply total determination, just default determination.

[13]  My treatment of the WMD example assumes throughout that the Bush Administration genuinely believed that the Iraqis had WMD.  An alternative possibility, of course, is that

they merely pretended to believe it and made public pronouncements accordingly. While historically quite possible, this scenario would eliminate a collective belief for us to evaluate; so I ignore it.

[14] Moreover, while Condorcet's original theorem assumed a uniform level of competence greater than one-half, this condition can be relaxed. It is sufficient that the <u>mean</u> probability of being right for each individual be above one half. See Grofman, Owen, and Feld (1983).

[15] More generally, belief states of a cognizer may include justified beliefs that <u>represent</u> external states of affairs and thereby constitute states of <u>evidence possession,</u> where the items of evidence possessed are the represented external facts. Comesana (2010 ) is another proposal by a reliabilist to synthesize reliabilism and evidentialism, although Comesana's rationale for such a synthesis is quite different from mine.

[16] List and Pettit discuss the deployment of experts under the heading of "epistemic gains from decentralization" (2011: 95-97). Suppose that a group seeks to form a belief on a certain conclusion based on k premises. Instead of inviting each group member to make a judgment about each premise, the group might be partitioned into k sub-groups, one for each premise, whose members specialize on that premise and make a judgment on it alone.

[17] This problem was pressed upon me by Holly Smith and Frederick Schmitt.

[18] The example is drawn from Goldman 2012: 7.