Evidential collaborations: Epistemic and pragmatic considerations in 'group belief'

Kent W. Staley

Saint Louis University

November 10, 2006

1 Introduction

In the wake of important work by Margaret Gilbert (e.g. 1987, 1994, 2002), a debate has emerged among social epistemologists about the appropriate understanding of statements of collective belief[1]. Both sides of the debate agree that such statements should be understood as expressing a group intention that is not reducible to the intentions of individuals.[2] They disagree, however, about what kind of intentional state such statements express. *Believers* hold that such statements express genuine beliefs. In Gilbert's view, in a case of group belief the group relates to the proposition that they assert in the same way that an individual who believes a proposition relates to that

[1] I will use the expression 'collective belief' to refer to whatever it is that is expressed by collective belief statements ('we believe that ...', 'we hold that ...', 'the committee has concluded that ...', etc.), without committing to the view that these are in fact beliefs properly understood.

[2] For a defense of a strong version of this anti-individualistic stance see Tollefsen 2002.

proposition. *Rejectionists*[3], on the other hand, maintain that statements of collective belief do not in fact express beliefs but some other kind of cognitive attitude, typically labelled *acceptance*. Crucial to the distinction is that one can accept a proposition without believing it (as when one accepts a proposition 'for the sake of argument').

Another distinction that is widely held to separate belief from acceptance is that beliefs are appropriately evaluated on strictly epistemic grounds (beliefs are 'truth-directed'), while accepting a proposition always properly incorporates considerations of utility as well as evidence (acceptance 'aims at success') (Meijers 2003). On a related point, it has been claimed by Wray (2001, 2003) that collectives can only exist, and hence can only act (as when they accept a proposition), on the basis of shared aims, whereas individuals exist independently of their aims, and do not form beliefs for the purpose of achieving an aim.

The purpose of the present paper is to examine the role of evidential considerations in relation to pragmatic concerns in statements of group belief. I will focus in particular on examples from collaborative work in experimental physics. Such examples seem particularly apt for addressing this question, as it is precisely in these cases that one expects evidential considerations to be most prominent. I will not address, therefore, what is typical of statements of group belief, but what characterizes those cases in which the collective exists precisely for the purposes of generating claims about evidence or claims based upon evidence with regard to some particular kind(s) of phenomena.

[3]Gilbert coined the term 'rejectionist' to describe her critics. Tollefsen has introduced the correlative term 'believer'.

Arguments from rejectionists have depicted the features of voluntariness and goal-directedness as indications that collective statements of belief express acceptance rather than genuine belief, and as detriments to the epistemic status of such acceptances. I will argue that these very features, whether or not they are compatible with viewing groups as having beliefs properly speaking, are crucial to the possibility of *enhancing* the epistemic status of group statements, and indeed to the possibility of anything resembling the sciences as we know them. That this point has not come through clearly in previous discussions reflects a tendency to take too narrow a view of epistemic values.

The present paper does not aim to address directly the debate between believers and rejectionists. Indeed, I adopt the perspective that the question of the role of evidence in statements of group belief, though part of a subsidiary argument in that debate, is more important than the debate in which it has emerged. My paper thus has a secondary aim of returning some of the attention of this area of social epistemology to a point that Gilbert herself emphasized in the works that gave rise to the debate between believers and rejectionists: for better or worse, collective belief is a prominent and influential social phenomenon. Knowing how to respond to it and what it means obviously requires that we should consider what role evidential and other epistemic considerations play in the formation of collective beliefs, and likewise for pragmatic considerations.[4] I will argue as well that this problem is of great practical importance for those scientists who work in collaborations or make use of results produced by collaborations (in other words, all scientists). By contrast, although the debate between believers and rejectionists certainly holds theoretical interest, it is unclear what hangs on its outcome.

2 Collaborations as collective epistemic agents

Collaborations can be found throughout the sciences, and different kinds of collaborations are formed for different purposes. Two theorists might collaborate on the basis of complementary expertise. Some collaborations are based on mentor-student relationships. A collaborator might contribute technological or financial resources in addition to her scientific abilities. The particular kind of collaboration I have in mind for the present discussion is one which exists precisely for the purpose of gathering data that can serve as the basis of evidence claims with regard to some type of (natural or social) phenomena.

Such groups vary widely with regard to disciplines, methods, size, and mode of social/institutional organization. What they share is that their activities include the collective assertion of evidence claims not merely as part of their ordinary way of doing things, but in partial fulfillment of the aim

[4] A similar shift away from the debate over the ontological status of group beliefs and toward epistemic considerations is proposed by Mathieson (2006); a further point of significant agreement between Mathieson's argument and mine will be noted in section four.

that constitutes their reason for existing at all: achieving an understanding of some particular aspect of the world.

I will use the term 'evidential collaborations' to denote the kind of group in question. Such collaborations collectively pursue many aims, of course. Quite prominently and necessarily, they seek funding. They also seek to preempt possible discoveries by competing groups. They pursue notoriety in the form of awards, honors, and even headlines. They strive to improve and expand their technological resources. They tend to grow larger over time, and may even seek to do so. For present purposes, I wish to emphasize a pair of significant aims that work in a crucial (but essential) tension with one another: They seek to avoid the embarrassment of making claims that subsequent work reveals to be false; they also seek to achieve prominence and esteem by making novel and significant claims that are upheld by further critical scrutiny.

Should we say that they collectively seek, not merely to avoid embarrassment, but to form only collective belief statements that are true? It seems that the evidence for thinking that such groups have a collective aim of truth would in most cases be about as strong as that for attributing such aims to typical individual scientists. Both individual scientists and evidential collaborations pursue a wide variety of aims simultaneously. Both typically or at least often claim that those aims include some member of a group of truth-related concepts: truth, approximate truth, or empirical adequacy.

It seems reasonable, therefore, to describe these groups as not merely seeking to avoid embarrassment, but as seeking to make collective statements that are true.[5] Such a description, however, leaves out the other member of the pair of aims mentioned above: the drive to make 'novel and significant' claims. Just as the avoidance of embarrassment partly reflects an epistemic aim (truth), the pursuit of prominence and esteem reflects this further epistemic aim. That the novelty and signifance of a group's evidence claims is a properly epistemic aim is a point that has not been adequately noted in discussions of group belief. Furthermore, I will argue, this aspect of collaborative inquiry has interesting philosophical consequences for our understanding of the phenomenon of group belief.[6]

I will make my argument in the context of a specific example: the first claim of evidence for the existence of the top quark by a very large group of particle physicists in 1994.

3 The CDF Collaboration and the evidence for the top quark

The Collider Detector at Fermilab (CDF) Collaboration announced in April 1994 that they had found evidence for the existence of the top quark, an elementary particle postulated by the Standard Model of physics. They issued

[5]So as not to beg the question with regard to scientific realism, all attributions of an aim of 'truth' in this paper should be read weakly, i.e. as expressing a disjunction such as 'truth, approximate truth, or empirical adequacy'.

[6]David Hull (1988) offers an account of scientific development in which such self-regarding and non-truth-related aims as self-promotion play a significant role. My aim here is not to dispute such an approach. Insofar as Hull's account implies that truth-seeking is not the whole story behind the epistemic success of scientific inquiry, his account is compatible with the present discussion.

a paper, titled 'Evidence for top quark production in $\overline{p}p$ collisions at $\sqrt{s} = 1.8$ TeV', that presented a lengthy and detailed discussion of their analysis and cross-checks on their results (Abe et al. 1994). That paper (henceforth 'the Evidence paper'), which bore the names of about 450 physicists as authors, concluded 'The data presented here give evidence for, but do not firmly establish the existence of $t\bar{t}$ production in $\bar{p}p$ collisions at $\sqrt{s} = 1.8$ TeV' (ibid. 3023).

CDF presents this conclusion as being supported by a complicated argument appealing in part to a statistical excess of particle collision events bearing the features that would be expected of events in which top quarks were produced and subsequently decayed ('candidate events'). That statistical excess beyond the prediction from other 'background' processes yielding similar appearances is evaluated quantitatively. They report that the probability of observing as many candidates as they found or more, assuming that only background processes are present, is 2.6×10^{-3}. In addition, they discuss other 'features of the data' such as the number of events in certain categories, the distribution of reconstructed masses of particles involved in the candidate events, and certain qualitatively-discussed kinematic features of the candidate events. Some of these features are claimed to not support the top quark hypothesis, while most of them do support that hypothesis.

I have studied the history of this particular experimental group and especially this particular result in considerable detail (Staley 2004). Here I will discuss a point regarding one of my findings in this case that has particular relevance for the discussion of group belief.

But first, an important if rather obvious observation: there could be no evidence for the top quark (at least none of the kind we presently have) if scientists did not collaborate on a very large scale in order to find such evidence. To explore experimentally questions relating to the top quark at all, scientists must undertake to do so *together*.

The point I wish to emphasize, though, goes beyond this preliminary observation: For practical purposes, the ability to make such claims *requires* a willingness to make compromises between the individual members' beliefs and the belief statement of the group. To substantiate and explain this point I have to say a little bit more about how CDF works in general as a group author of experimental reports, as well as how they worked in the case of the Evidence paper.

As with most particle physics collaborations, CDF publishes their papers under authorship rules that are quite liberal. The collaboration maintains a 'standard author list', inclusion on which is based on having performed service to the collaboration for a minimum of the equivalent one year at full time. This list serves as the default author list for any papers published by the collaboration. One can ask that one's name be removed from the list. In short, the system is an 'opt-out' system.[7]

Thus the decision to publish a result is based on consensus within the collaboration. This consensus need not, however, be perfect. Indeed there are different kinds of imperfection that can be tolerated. The Evidence paper was based on what Bill Rehg and I have elsewhere called 'heterogeneous

[7]CDF member Henry Frisch has argued that this policy has negative methodological consequences that outweigh its advantages. He has argued for replacing the existing rules with an 'opt-in' policy (Frisch 2004).

consensus' (Rehg and Staley, forthcoming). Heterogeneous consensus obtains when members of a collaboration issue a statement of an argument in which the conclusion of the argument may be accepted by all, but there is disagreement as to which of the premises cited actually support that conclusion.

In the case of the evidence paper, the heterogeneous consensus came about in roughly the following way: The CDF collaboration comprises many sub-groups working more or less independently on methods of data analysis intended to yield conclusions about physics questions of interest. Within CDF prior to the Evidence paper, a number of different groups were working on different methods of analysing data in order to look for evidence of the top quark. There were several groups working on 'counting experiments' that looked for the kind of statistical excesses in the numbers of candidate events mentioned above. Different groups looked at different decay paths for top quarks, or used different signatures within a single decay path to separate signal from background. Other groups were looking at kinematic features of events, such as the geometry of decay products or features involving the total energy or momentum of the decay products. Each group was a potential source of evidence for the top quark. When CDF undertook to argue that they had indeed found evidence for the top quark, they faced the decision of which such sources of evidence to include in their principal argument in support of this claim, which sources would figure only in a secondary way as a kind of cross-check on the central argument of the paper, and which if any sources would be excluded entirely.

As it happened, a number of controversies arose within the collaboration over several possible sources of evidence. Some of the counting experiments were suspected by some collaboration members of being subject to potential biases. Others distrusted the results of the kinematic analyses on account of the reliance of the analysis on 'Monte Carlo' computer simulations in the estimates of background. Based on interviews with collaboration members shortly after the publication of the Evidence paper, it is clear that these disagreements had not been resolved in the process of writing and publishing.

Rather, what emerges from those interviews is that, while (nearly) everyone interviewed expressed the view that the Evidence paper was basically correct in its conclusion (more about the vagueness of 'basically correct' shortly), there remained some disagreement over what features of the evidence supported that conclusion. Such a position was facilitated in part by the belief that the overall analysis of the data had tended to be 'conservative', in the sense of tending to draw weaker rather than stronger conclusions from the data. Hence, the consensus regarding the claims put forth in the paper was heterogeneous.

In the discussion to follow, it might be useful if we here draw some distinctions between different statements that possibly express CDF's group beliefs relating to the top quark as reported in the Evidence paper.

I: 'The top quark exists'.

II: 'Our data constitute evidence for the top quark'.

III: 'The data from such-and-such analyses constitute evidence for the top quark'.

IV: ‘Our data (or the data from such-and-such analyses) constitute evidence for the top quark with statistical significance of 2.6×10^{-3}’.

V: *The entire text of the Evidence paper*

The status of any claim about CDF’s group beliefs in this case will depend rather strongly on the specific claim that is attributed to them. As will be seen, the more content that is attributed to their group belief, the more difficulty arises in thinking of the statement as expressing anything like a shared belief in the proposition in question.

4 Pragmatic considerations and epistemic benefits

The discussion of collective belief has been generally carried out with the following assumption. To the extent that pragmatic factors play a role in the formation of group belief, this indicates an epistemic drawback, in the sense that, were the group belief to be based only on non-pragmatic epistemic factors, the belief would be more reliable. The implication is that belief-forming processes governed by purely epistemic considerations are more reliable. This in turn might stem from the view that beliefs formed on the basis of strictly epistemic considerations are those based on the evidence alone, where proper epistemic functioning is achieved when beliefs are ‘forced’ upon us by the evidence and the evidence alone.

Such assumptions can be seen at work in Christopher McMahon’s discussion of the distinction between epistemic acceptance and discretionary acceptance, where the former is ‘forced by the evidence’ and the latter is not (McMahon 2003). McMahon’s point is that there is an ambiguity in

Gilbert's description of group belief in terms of the group relating to a proposition in the way that an individual who believes that proposition relates to it. In discretionary acceptance, the group members contribute to the group's thus relating by way of actions that are under the control of the will, such as arguing in defense of that proposition. But in cases of epistemic acceptance, individual members contribute to the group acceptance by 'adopting it themselves in the way appropriate to individual belief' (ibid., 353). Such a characterization has the consequence, of course, that in cases of epistemic acceptance, there will be a strong consensus within the group in the sense that each member will individually believe the proposition in question.

I do not deny that such a distinction can be drawn, but I would like to draw attention to the limited range of epistemic concerns that the distinction deploys, as well as to blunt the force of any implication that cases not meeting the standards of McMahon's epistemic acceptance are necessarily worse off in epistemic terms.

There are apparently clear instances of collective judgment that are discretionary in McMahon's sense. He gives as an example the insistence by tobacco companies, over many years, in the face of mounting evidence to the contrary, that cigarette smoking does not cause cancer. To take such a position was certainly in the companies' interest (to a point), and could be understood as a collective decision based on their mutual interests and very much under their collective control.

In other cases, by contrast, all the members of a group may very well come to see the evidence in support of a proposition as so overwhelming that they ‘have no choice’ but to believe it individually, and on that basis may collectively commit to believing it as a group.[8] However, the case of CDF and the evidence for the top quark seems not to fit neatly into either category.

The only way to describe CDF’s group belief with respect to the top quark in a way that meets the standards of epistemic acceptance in McMahon’s sense is to attribute to them the very vague group belief (II). CDF does not even assert (I), and none of the other propositions was at the time believed to be true by each individual who participated in the group’s assertion.[9] But CDF as a group was much more specific in what they asserted than just stating that their data constituted evidence for the top quark. In their paper, in a press release, and in collaboration-approved presentations to professional meetings and academic departments, they addressed as a group the specific features of the data that they regarded as constituting

[8] I am assuming here for the sake of argument that there is nothing inherently problematic in the notion of two individuals sharing a belief. This is not to say that such an assumption cannot be called into question. The historian of science Mara Beller problematizes the notion of consensus, particularly in science, in her study of the Copenhagen school in the history of quantum theory (Beller 1999). Beller’s account, in which even a publication by a single author is regarded as exhibiting a polyphony of views, should give us pause before accepting uncritically that scientific consensus in the sense of shared belief is a straightforward notion. Alban Bouvier has related Beller’s work to Gilbert’s account of collective belief (Bouvier 2004).

[9] This claim requires substantiation: None of the claims III–V was subject to a complete consensus even among the two dozen or so CDF members I interviewed, all of whom appeared as authors on the Evidence paper (and can be regarded as participating in at least that sense), and nearly all of whom made substantive contributions to the analysis presented in that paper (and thus participated in a stronger sense as well).

that evidence, as well as the strength of the evidence (and, for that matter, the weaknesses in the evidence).

Ought we then to treat CDF's collective statements as an expression of a discretionary acceptance? McMahon indicates that the crucial distinction is the extent to which acceptance is 'forced by the evidence'. As a consequence, in cases of epistemic acceptance 'all the members agree that the relevant evidence supports a particular conclusion' (ibid., 353). If CDF did not exhibit this kind of agreement, then their acceptance in this case was discretionary rather than epistemic.

The complication introduced by this kind of case is twofold: First, even supposing that all agree on the conclusion that is supported by the relevant evidence, there need not be agreement on what evidence is relevant or how that evidence supports the conclusion. That is to say, different aspects of the evidence might be (and in this case were) seen as supporting the conclusion by different members of the group. Second, agreement is not all-or-nothing, but is a matter of degrees. That is to say, almost all members of CDF were in agreement at some level with the paper that they collectively published. But as one attempts to specify more precisely what the group belief is, one finds that agreement breaks down, not entirely, but in terms of something like closeness of belief. For example, one might not agree with the specific number given for the statistical significance of the results, but still feel that the data constituted evidence that was statistically significant at some small number. Or one might feel that of the three different counting experiments cited by the group as contributing to the evidence, two were trustworthy while the other should be discounted. It is for this reason that I above described group members I interviewed as agreeing that the Evidence paper was ‘basically correct’.

Following McMahon, it seems that we ought perhaps to say that CDF formed an epistemic acceptance of the very vague statement (II), while all the rest of their group statements express discretionary acceptance. This seems like a very odd thing to say about what is probably a very pervasive type of scientific activity, and the very basis for much of what we regard as established scientific knowledge. Viewing most of CDF’s collective statements as discretionary acceptance seems strikingly odd in light of McMahon’s further comments on this kind of acceptance. In such cases, he writes, ‘The thinking underlying a group’s decision concerning what its position will be need not be self-consciously calculating; it may involve some form of collective self-deception. But the will enters into the adoption of the group’s view in a way that is incompatible with proper epistemic functioning’ (ibid., 351).

Clearly in this case, some – perhaps most – individual CDF members decided to defer to the judgment of the group in signing onto a paper that asserted some things with which they were not in perfect agreement. Insofar as such a decision is under the control of the individual will, then the case looks like an instance of discretionary acceptance. It is not at all clear in this case, however, that this contribution of the will interferes with proper epistemic functioning. To the contrary, assuming that we wish for groups such as CDF to say *something* about their findings, the distributed nature of expertise in such a group constitutes a reason to believe that such deference to the collective judgment enhances the group’s epistemic functioning.

The obvious objection to this last claim would be to point out that I am here, by assuming that we do want to hear from CDF about their evaluation of their data, bringing in a practical aim that is in an obvious tension with epistemic concerns. Our epistemic concern is not, after all, to hear from CDF no matter what. Our epistemic concern is to learn what is true about the phenomena that the group is investigating, and we are more likely to hear the truth from them if they report to us only those statements that every member of the collaboration agrees is compelled as a belief by the available evidence. If they would confine themselves, in other words, to reporting their epistemic acceptances, they would be less likely to mislead us and themselves.

In reply, I do not deny that greater reliability would be achieved if no evidential collaboration ever made collective statements that were not subject to completely uniform – and in McMahon's sense epistemic – agreement. I do deny that such a policy would confer only epistemic benefits and no epistemic harms. There are other epistemic values besides reliability, and evidential considerations do not exhaust a group's epistemic concerns. Evidential collaborations operate at the frontiers of accepted knowledge and some degree of uncertainty attends nearly every important step they take.

I would like to suggest that if one were to ask that groups should ideally issue statements of group belief only when there is complete uniformity in what each individual member is 'compelled' to believe based on the evidence, one would in fact be saying that ideally such groups would issue almost no statements of any interest, and thus that there simply would be very little interesting empirical science. Plausibly, results in science that are of interest tend to rest on heterogeneous consensus and homogenous consensus is generally a feature of retrospective assessment of past scientific work. It is not only reliability that is of interest to evidential collaborations. What is also of interest is the discovery of novel and significant phenomena, an achievement that both requires collaboration and that tends to generate only imperfect consensus. [10]

Consider the billions of dollars that made possible the work of CDF and other groups at Fermilab. No one in their right mind would pay that kind of money simply in order to find truth, which can be had at virtually no cost ('if the rest mass of the top quark is greater than five kilograms, then it is greater than four kilograms'). If any epistemic aim justifies such an expenditure of resources, or indeed any expenditure at all, on scientific research, then it is not truth per se, but *informative* answers to *significant* questions. This is an epistemic concern insofar as it is strictly a matter of the pursuit of knowledge that can be distinguished (in principle, at least) from the pursuit of fame or funding. Yet it is a concern that can only be met by precisely the kind of pragmatically-guided collaborative belief-formation that pervades the production of research reports by evidential collaborations. In other words, the incorporation of pragmatic factors *enhances* the epistemic credentials of the claims that such groups put forward.

This point can be viewed in another way. Would the proclamations of evidential collaborations be on stronger evidential ground if they directly reflected the beliefs that the evidence 'forces upon' each individual member?

[10] Mathieson (2006, 166–68) presents further reasons why a group may, while still "aiming at the truth," adopt a view that is not shared by all of its members.

In one sense, they probably would. A collaboration could wait until the evidence became so clearly univocal that every group member would come to see it in the same way, and only then issue a statement. It is likely that such statements would exhibit a lower error rate than the results such groups presently publish. It is also likely that these statements would be concerned with largely mundane issues of little interest to anyone, or else with scientific questions that, though they may have once held interest, have long since come to be regarded as uninteresting. Assuming, on the contrary, that we value not only truth itself, but the *kind* of information that is currently made available by the activities of evidential collaborations, the strongest evidential ground is in fact the judgment of the group as a whole. There is a recognition within such groups that expertise is distributed throughout the group, making the group as a whole 'smarter than' its members individually (Giere 2002).

The point I have been emphasizing can be seen as a particular application of a point made by Alvin Goldman in *Epistemology and Cognition*: epistemic appraisal can be made with respect to a variety of standards. Reliability is one important standard, but another is what Goldman calls 'power', where the power of a cognitive system is roughly speaking some kind of function of the proportion of questions the system seeks to answer that it can answer correctly (Goldman 1986, 122–41).[11] To this I would

[11] Power in this sense is thus suggestive of, but clearly distinct from, power in the statistical sense. Goldman has subsequently abandoned the terms 'reliability' and 'power' in favor of a unified measure of veritistic value (Goldman 1999, esp. 90). However, the epistemic phenomena relevant to reliability and power continue to be reflected in veritistic value. Goldman thus continues to attribute epistemic value to the ability to answer

in fact add another epistemic value of considerable importance to cognitive systems that operate within scientific contexts: the ability to generate *interesting new questions*. No sooner had CDF published the Evidence paper than physicists within CDF and elsewhere began to see within the details of their data suggestions of new possibilities in physics that could not be established without the collection of more and different data (see Staley 2004, Epilogue). The episode serves as an illustration of the connection between information that can serve as the basis for answering a question and the information that can be used to formulate an interesting new question.

5 Goal-directedness and the will

As Goldman notes in his discussion of reliability and power, one important difference between the two is that in evaluating a cognitive system's power, one is concerned with evaluating it with respect to goal-responsiveness, whereas reliability is not related to goal-responsiveness (Goldman 1986, 124). The dual demand that beliefs should be forced upon us by the evidence and that they not be based on considerations of utility amounts to an insistence upon seeing justified belief as passively acquired. However, by focusing on contexts dealing with questions that are very hard to answer, we see that one needs to work to ascertain both what the evidence is and what the appropriate conclusion is to derive from it. The effort to determine what evidence is relevant and what the relevant evidence indicates are highly goal-directed activities. In these contexts, waiting passively for justi-

questions of interest, and not just to the avoidance of error.

fied belief to come would result in mere ignorance, or else knowledge of only trivial truths.

This brings us to another assumption apparently at work in the debate between believers and rejectionists: involuntarism with regard to belief. This is the thesis that, although one can choose to accept a proposition, what one believes is not under the control of the will. Of course this is not quite the same as the thesis that beliefs are forced upon one by the evidence, since, as Wray points out, there are other causes of belief besides evidence (Wray 2003, 369). According to involuntarists, these other causes, however, are also not under the control of the will of the epistemic agent.

Deborah Tollefsen has noted that, even granting involuntarism to be true of individuals, it might fail to hold at the group level (Tollefsen 2003, 398). She argues that since involuntarism about belief is, if true, a contingently true thesis about human beings, it cannot be assumed to be true of groups. Even if individuals cannot choose by an act of will what they will believe, it is possible that groups might.

Brad Wray defends his rejectionist position not by appealing directly to voluntarism, but rather by insisting that the important distinction between beliefs and mere acceptance is that 'agents accept views in light of their goals, whereas beliefs are not acquired in light of goals' (Wray 2003, 369). Since groups only accept propositions in light of their goals, groups do not believe the propositions they accept. Tollefsen disputes the claim that belief is not goal-directed, arguing that 'there are lots of goals that inform and direct us in our acquisition of beliefs' (Tollefsen 2003, 400). Tollefsen's point seems to be that the kinds of activities that result in our having beliefs are typically goal-directed activities, a point reflected in Goldman's characterization of assessments of cognitive power as resting in part on considerations of goal-responsiveness.

Tollefsen's point is certainly sound, but does not seem directly to address Wray's argument. Wray appears to be claiming, not that the activities that result in our having beliefs are not goal-directed, but rather that we do not choose our beliefs on the grounds that they serve our purposes. In other words, suppose my engaging in activity X is directed at some aim Y, and doing X results in my believing p. Wray wishes to claim that my desire to achieve Y might explain why I did X, but not why I believe p. Indeed, believing p might not at all serve my aim of achieving Y. By contrast, he seems to be arguing, in deciding to *accept* a proposition p, I make my decision precisely in light of my desire to achieve some particular aim.

Insofar as the dispute between believers and rejectionists turns on the debate over voluntarism vs. involuntarism, we are faced with an attempt to resolve one thorny and methodologically troubled debate by appeal to another debate that suffers exactly the same difficulties. I am therefore pleased to see in the exchange between Wray and Tollefsen a move away from that issue.

Nonetheless, I believe that goal-directedness, no less than voluntarism, is ultimately irrelevant to the questions highlighted here. I claim that when we ask about the role of evidence in, and the evidential support for, statements of group belief, we are likely to be misled if we focus on the issues of will or of goal-directedness. Rather, what we should be asking in such situations is whether the group's statements are *constrained* by the group's own judgments of the relevant evidence, and whether the group functions so as to produce *unbiased* collective assessments of the evidence.

To put it another way, the evidential status of claims put forth by collaborations turns not on the question of whether the propositions thus asserted are accepted in the light of aims, but rather on what aims the group is pursuing in accepting such propositions, and how effectively the group pursues those aims. In particular, we want to know whether the group is collectively pursuing epistemic aims in the acceptance of a proposition, and whether their method of pursuing those aims is likely to succeed.

It might be helpful here to consider the issues from two perspectives: that of a recipient of group judgments and that of a contributor to group judgments. The considerations are different but closely-related in these two cases. For neither the recipient nor the producer of group judgments does the question of goal-directedness play a central role in evidential assessment.

When confronting a claim presented by a collaboration, the recipient's interest is not typically to determine whether the group itself believes or merely accepts the proposition that has been claimed. Rather, the recipient will want to know, 'Should *I* believe this?' Simply knowing that the proposition has been accepted by the group to serve some aim does not by itself help to resolve this question affirmatively or negatively. Knowing what aims guided the group's decision, however, might be very relevant. Especially important will be to know the relative weight given to aims that are in some sense in tension. Was getting a mention in the newspapers valued over retaining the respect of researchers in neighboring disciplines? Was priority valued over thoroughness? Also relevant for the recipient's assessment is the process followed by the group in determining the content of their statement. Here relevant questions might include the mechanisms within the group for resolving disagreement, the extent to which procedures are instituted that facilitate criticism to eliminate biases, the extent to which the group in a particular case followed their own procedural guidelines, etc.[12] Learning that the collective statement is not the subject of complete agreement within the group is certainly relevant and might weaken one's inclination to believe the statement in question. Such information should, however, be evaluated in light of the nature and extent of the disagreement, and with an awareness of the need to balance the twin epistemic concerns of reliability and informativeness with respect to questions of interest.

Contributors to group judgments face related questions. Of course, in any given deliberation, members of the group face the sometimes difficult question of when the consensus statement is sufficiently at odds with their own individual beliefs to warrant withholding their participation. In addition, members of evidential collaborations, which aim for both reliability and relevance, must give considerable thought to the group's structure and procedures, while also consistently keeping an eye on the processes through which those procedures are implemented. Such a group is a kind of organism of its own that exists for the purpose of producing reliable and relevant scientific information. The members of the group are responsible for the proper epistemic functioning of this organism, and they can only fulfill this responsibility by a vigilant exertion of the will at both the individual and collective

[12] Rehg and I have dubbed such assessment *probation by process* (Rehg and Staley forthcoming).

level. Brad Wray maintains that 'plural subjects' such as evidential collaborations are 'constituted by their goals' (Wray 2003, 369). I do not dispute this claim, but add that *these* plural subjects are constituted by epistemic goals, and they are, or should be, mindful of the need to self-organize in a way that promotes those goals.[13]

Collaboration physicists with whom I have talked often use the term 'sociology' to refer to any aspect of their work pertaining to the interpersonal dynamics of group decision-making. The term is often used with a certain amount of distaste, and is meant to contrast with what they regard as the much more important and valuable 'physics' that is their primary concern. Nonetheless, collaboration scientists ignore such 'sociology' at their own peril. CDF members' reputations as physicists depend in part on their collective ability to solve the problem of 'sociology'. Success as a knowledge-producing group requires that they, like any evidential collaboration, be able to solve problems in applied social epistemology.

[13]The growing literature on problems of judgment aggregation and the 'discursive dilemma' indicates that groups seeking to aggregate individual judgments on a set of logically related propositions must, to put it roughly, choose between (a) using a uniform method to aggregate individual judgments on each proposition, and (b) ensuring that the judgments reached in the aggregate meet the minimal rationality standard of logical consistency (see, e.g. Pettit 2001, List and Pettit 2002, Pauly and van Hees 2006, List 2006). This problem alone suggests that the problem of group commitment to a proposition is one in which evidential and pragmatic problems are not so easily separated.

6 Conclusion

To help clarify the stance I have taken here, and by way of conclusion, it may help for me to contrast my own view with that of Brad Wray, who has also approached the issue of group belief with a particular interest in its relevance for understanding scientific inquiry. In a paper defending the claim that what Gilbert calls group belief is best understood as group acceptance, Wray notes a distinction between what he terms 'epistemological and social' considerations, noting that scientists working in groups are motivated by both (Wray 2001, 329). He makes this point in service of a critique of such 'social constructivist' accounts of science as that offered by Latour and Woolgar (1986). To the extent that Wray's 'social' considerations can be regarded as falling under what I am calling 'pragmatic' considerations, our views are in agreement. We both recognize that evidential collaborations are moved by both pragmatic and epistemic aims. To defend his rejectionist position, Wray must make the further assumption that in every group commitment to a proposition, some pragmatic considerations enter, and that this disqualifies such commitments from counting as a belief. I take no stand on this point in the present essay, but I do wish to make the further claim, not addressed in Wray's account, that the epistemic standing of statements by evidential collaborations can, and in general should, be enhanced by those pragmatic considerations that do play a role.

In focusing on such issues, I do not mean to deny the philosophical interest or importance of debates concerning propositional attitudes such as belief or acceptance. Rather, I wish to highlight an opportunity for social epistemologists, evidence theorists, methodologists, and working scientists to cooperate in an inquiry of considerable practical importance: what are the appropriate aims of evidential collaborations and what kinds of organizations, methods of judgment aggregation, and discursive practices will best enable such groups to meet those aims? Here I do not pretend to have contributed to answering such questions, but rather have sought to emphasize some of the *epistemic* considerations that such an inquiry will need to incorporate.

REFERENCES

Abe, F., M. G. Albrow, et al. [CDF]. 1994. Evidence for top quark production in $\overline{p}p$ collisions at $\sqrt{s} = 1.8$ TeV. *Physical Review D* 50:2966–3026.

Beller, Mara. 1999. *Quantum dialogue: The making of a revolution.* Chicago: University of Chicago Press.

Bouvier, Alban. 2004. Individual beliefs and collective beliefs in sciences and philosophy: The plural subject and the polyphonic subject accounts. *Philosophy of the social sciences* 34:382–407.

Frisch, Henry. 2004. The twin questions of authorship and the reproducibility of results in large scientific collaborations. Paper presented at Philosophy of Science Association 2004 meeting, Austin, Texas, November 18, 2004.

Giere, Ronald. 2002. Scientific cognition as distributed cognition. In *The cognitive basis of science*, edited by Peter Carruthers, Stephen Stich, and Michael Siegal. Cambridge: Cambridge University Press.

Gilbert, Margaret. 1987. Modelling collective belief. *Synthese* 73:185–204.

———. 1994. Remarks on collective belief. In *Socializing epistemology: The social dimensions of knowledge*, edited by Frederick Schmitt. Lanham, Marlyand: Rowman and Littlefield.

———. 2002. Belief and acceptance as features of groups. *Protosociology* 16:35–69.

Goldman, Alvin. 1986. *Epistemology and cognition.* Cambridge, Mass.: Harvard University Press.

———. 1999. *Knowledge in a social world.* New York: Oxford University Press.

Hull, David. 1988. *Science as a process: An evolutionary account of the social and conceptual development of science.* Chicago: University of Chicago Press.

Latour, Bruno and Steve Woolgar. 1986. *Laboratory life: The construction of scientific facts*, second edition. Princeton: Princeton University Press.

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

List, Christian. 2006. The discursive dilemma and public reason. *Ethics* 116:362–402.

Mathieson, Kay. 2006. The epistemic features of group belief. *Episteme* 2:161–75.

McMahon, Christopher. 2003. Two modes of collective belief. *Protosociology* 18–19:347–62.

Meijers, Anthonie. 2003. Why accept collective beliefs? A reply to Gilbert. *Protosociology* 18–19:377–88.

Pauly, Marc and Martin van Hees. 2006. Logical constraints on judgment aggregation. *Journal of Philosophical Logic* 35:569–85.

Pettit, Philip. 2001. Deliberative democracy and the discursive dilemma. *Philosophical Issues* 11:268–99.

Rehg, William and Kent Staley. Forthcoming. The CDF collaboration and argumentation theory: The role of process in objective knowledge. *Perspectives on Science.*

Staley, Kent. 2004. *The evidence for the top quark: Objectivity and bias in collaborative experimentation.* New York: Cambridge University

Press.

Tollefsen, Deborah. 2002. Challenging epistemic individualism. *Protosociology* 16:86–117.

———. 2003. Rejecting rejectionism. *Protosociology* 18–19:389–405.

Wray, K. Brad. 2001. Collective belief and acceptance. *Synthese* 129:319–33.

———. 2003. What really divides Gilbert and the rejectionists? *Protosociology* 18–19:363–76.