

---

COLLECTIVE EPISTEMOLOGY<sup>1</sup>

**The traditional focus of philosophical epistemology**

Philosophical epistemology is concerned with knowledge and related phenomena such as belief. In order to have a general label for such phenomena I shall refer to them as *cognitive states*. Differing accounts of a variety of cognitive states have been produced. For instance, according to one venerable – if debated – account of knowledge, it is justified true belief. Belief has been contrasted with acceptance, though the belief-acceptance contrast has been drawn in a variety of different ways.

In nothing I have said so far has the following question come up. Who (or what) is capable of knowledge, belief, or, to use a general label for all of the relevant cognitive states, *cognition*? The question concerns not particulars but kinds. What kinds of entity are capable of cognition?

Most epistemologists would confidently include human beings in this class. One basis for this confidence is simple enough: everyday thought routinely ascribes such states to individual human beings. The detective asks his colleague: 'What does she know?' An informant says, 'I believe he went that way.' As a result the denial that individual human beings are subjects of knowledge and belief is counterintuitive. Such denial has occurred in the history of epistemology, but it needs to be argued.

Epistemologists largely focus on the individual human case, and few spend much time on the question 'Who (or what) knows?' and so on. There is no problem with this focus in itself. It has a consequence, however, that may be less innocuous.

Following an examination of the individual human case epistemologists standardly go on to offer general accounts of knowledge, belief, and so on that are based on that case. Then the question whether, for instance, animals have beliefs is answered in terms of an account of

belief that has been derived from reflection on the individual human case. What justifies this procedure?

If everyday thought routinely ascribes particular cognitive states to entities other than human beings, one might take that to be crucial evidence for a general account of these states. To allow this would not be to deny that the cognitive states of human beings have a distinctive character. On the contrary, one would expect them to have such a character, given the distinctive nature of human beings. It would be to allow that the various cognitive states, as these are conceived of in everyday thought, may not be restricted to human beings.

It is possible that all or some ascriptions of cognitive states to other kinds of entity are metaphorical, or that they are understood to involve some kind of pretence. That they are of this kind, however, cannot simply be assumed. It needs to be argued or shown by reference to examples that make it clear. Sometimes it may be clear enough that those who make them do *not* see them as either metaphorical or as involving some kind of pretence.

Again, it is possible that in all or some cases some kind of mistake is involved. This too will need to be argued. It will not do simply to insist that human beings alone are capable of cognitive states – so that some such mistake must be involved.

**Everyday ascriptions of collective cognitive states**

Everyday life is full of statements such as the following: 'The union believes that management is being unreasonable,' 'In the opinion of the court, this law is unconstitutional,' 'Our discussion group thought it was a great novel,' 'Our family believes in ghosts', 'Bill and Jane have concluded that it would be wrong,' 'We knew we had to stop'.

These appear to be ascriptions of *collective* cognitive states. That is, what is claimed to have a given cognitive state is understood to comprise two or more human beings. It is understood, further, that these people could properly ascribe the cognitive state in question to "us".

Thus construed, the examples given fall into two broad classes. There are ascriptions of cognitive states to two or more people who are understood to constitute an established group of a specific kind such as a union, court, discussion group, family, and so on. There are also ascriptions of cognitive states to two or more people without any presumption that they constitute an already established group.<sup>1</sup>

Though they differ in important ways, both expressions such as 'The union', on the one hand, and expressions such as 'Bill and Jane' on the other, can be thought of as referring to a *population*. Clearly, a population, in the sense in question, need not be an already constituted group, or social group, in the narrow sense in which courts and unions and discussion groups are paradigmatic examples.<sup>2</sup>

Populations may be specified not only by means of expressions such as 'the union' and reference to particular individuals, but also by reference to some feature possessed by all of their members. For example, the phrase 'haemophiliacs' picks out a particular population, comprising all haemophiliacs.

Particular collective cognitive states can be referred to in either of the following ways. With respect to belief, one might say either, 'Population P believes that p', or, equivalently, 'The members of population P (collectively) believe that p.'

I take it that those who make everyday ascriptions of collective cognitive states do not generally regard these ascriptions as metaphorical or as 'pretend' ascriptions. I shall assume that they are innocent of mistake unless and until they are proved guilty.

## Collective epistemology

To what phenomena are ascriptions of collective cognitive states intended to refer? In other terms, what are *collective* knowledge, *collective* belief, and so on?

Until very recently this question has been

largely been neglected by philosophers. Those who have neglected it include many of those who have focused to some extent on what they refer to as *social epistemology*.<sup>3</sup> In order both to emphasize its distinctiveness and to indicate its character I shall refer to the study of collective cognitive states as *collective epistemology*.<sup>4</sup>

Around 1980 Charles Taylor made some suggestive points relevant to collective epistemology in several publications. In particular he persuasively argued that to say it was common knowledge that such-and-such in a certain context was not to say that anything was yet, in his terms, *entre nous* or, again in his terms, *in public space*.<sup>5</sup>

In my 1987 article 'Modeling Collective Belief' I argued for a relatively full-blown account of collective belief, a closely related version of which is my book *On Social Facts*, 1989.<sup>6</sup> Since then there has been a steady increase in studies in collective epistemology.<sup>7</sup> Nonetheless it has not yet achieved the recognition from epistemologists and other philosophers that it deserves.<sup>8</sup>

One reason that it deserves such recognition has already been indicated. If epistemologists are to offer a plausible account of cognitive states in general, as opposed to the cognitive states of individual human beings, they need to consider in a more than cursory way the nature of collective cognitive states.

Irrespective of this consideration, and whatever precisely they are, collective cognitive states would seem to be a significant aspect of the social world. Their importance was emphasized by one of the founders of sociology, Emile Durkheim, in his classic work *The Rules of Sociological Method*. Though there may be earlier pertinent discussions, Durkheim has some title to be considered the founder of collective epistemology. His suggestive words on the topic, however, badly stand in need of further development.<sup>9</sup>

In the rest of this article I sketch my own approach to collective epistemology, with some reference to related discussions. I hope to encourage those who have not yet engaged with collective epistemology to do so. There are many avenues to pursue. These include the refinement of existing proposals, the development of alternatives, and examination of the

consequences of these proposals for general epistemology.

Philosophers tend to be modest about the relevance or effect of their work in relation to science – the tendency nowadays is to think of philosophy as informed by the sciences rather than the other way round. Many social scientists, meanwhile, are keen to avail themselves of the results of philosophical investigations. Philosophical investigations in collective epistemology, then, are likely to constitute an important contribution to the social sciences.

## Method

My own approach to collective cognitive states has been as follows. Examine the contexts in which everyday ascriptions of such states are generally considered true or false with the aim of arriving at a perspicuous description of the phenomena people mean to refer to when ascribing collective cognitive states.<sup>10</sup> I have largely drawn on my own understanding of when it would or would not be in order to ascribe a particular cognitive state to a given population. My own focus to date has been on collective belief and I shall focus on that here.<sup>11</sup>

## The simple summative account

The first hypothesis about the nature of collective belief that occurs to many people is what I have called the *simple summative account*. This maintains that – according to everyday understandings – the members of a population, *P*, collectively believe that *p* when and only when all or most of its members believe that thing. For instance, to say that the union believes management is being unreasonable is to say that all or most members of the union believe this.<sup>12</sup>

It can be argued, however, that most of a population's members believing that such-and-such is neither necessary nor sufficient for their collectively believing it. I now briefly argue this.<sup>13</sup>

Consider first the question of necessity. That is, consider whether most members of *P* must believe that such-and-such in order that they collectively believe it.

If one looks at examples like those of a union or a court, the answer appears to be negative.

One relevant consideration is this.

Often what is taken to determine the collective belief in such cases is a formal voting procedure where the opinion that receives the most votes is deemed, for that reason, to be the opinion of the court, the union, or whatever. What counts as far as any individual member is concerned, then, is his or her vote – not what he or she personally believes.

Evidently, one's voting in favor of the proposition that Jones was the best candidate does not entail that one personally believes that Jones was the best candidate, whether or not others might be inclined to think one does believe this, given one's vote. In other words, it cannot be argued that because one did not believe that Jones was the best candidate one did not really vote for him. Nor can one generally argue that the absence of correlative personal beliefs invalidates the result of the voting procedure.

Less formal processes can determine a collective's belief. These can be argued to be analogous to the case of voting at least insofar as public expressions as opposed to private thoughts are what counts. It is worth considering one such case in some detail.

Suppose that the members of a poetry discussion group have been focusing on Sylvia Plath's famous poem, 'Daddy'. After a lengthy discussion of its content and structure has taken place someone asserts that it is a very powerful poem. The other members respond in a positive way. One says "Yes, indeed", another "That's so true", and so on. On some such basis a group member might reasonably say, 'Our discussion group thinks this is a powerful poem'.<sup>14</sup>

Why would this collective belief ascription be in order? One thing to consider here is what the positive responses in question amount to.

Suppose you are in this group and don't particularly admire the poem in question. However, you don't want to get into an argument, so you refrain from expressing your own view. Indeed, at the crucial juncture you nod your head in an approving way. There is perhaps some likelihood that people will take you to think the poem powerful if you do this. In that case you are doing something that is likely to mislead them. You may in no way intend to mislead anyone, however.

Your primary intention may best be described along the following lines. You intend to express your readiness to see the belief that the poem is a powerful one established as the belief of the group. Your having an intention of this kind will be entirely reasonable insofar as the aim of the discussion as a whole is primarily a matter of developing a collective perspective on the poem.<sup>15</sup>

I return now to the question: why would the collective belief ascription at issue here be in order? My conjecture is roughly as follows. It is in order insofar as the members of the poetry discussion group have all indicated their readiness to let the belief in question be established as the group's belief. I say more later by way of explanation of precisely what is expressed by the participants. For now the main point is this. In this informal type of case, also, there may be a collective belief that *p* without all or most – or indeed any – members of the population in question believing that *p*. This stark conclusion follows from my understanding of the process of collective belief formation, and seems true to the facts. Precisely which members of a population believe that *p* when the members collectively believe that *p* will need to be determined empirically.<sup>16</sup>

To go back briefly to the case of voting, though this has special features, it too can be argued to involve concordant expressions of readiness with respect to a group's belief. Participating in the voting process, whichever way one votes, is a way of indicating one's readiness to accept as the belief of the group that belief picked out by the voting process.

I turn now to the question of sufficiency: is the fact that all or most members of a population believe that such-and-such sufficient for the members collectively to believe it? It is easy enough to argue for a negative answer.

Consider a court. A certain matter may not yet have come before it. It would then seem right to say that, as yet, the court has no opinion on the matter. The individual justices may, at the same time, have definite personal opinions about it. What they now think, however, is not relevant to the question of what the court now thinks.

Similarly, every single member of a poetry discussion group may personally think a certain

poem is a powerful one, but the group may not have discussed it yet. It would seem, therefore, that there is as yet nothing that the group believes about this poem.<sup>17</sup>

To further the point, note that if discussion does take place, the resulting 'opinion of the court' or 'judgment of the poetry group' – based on the relevant expressions – could be opposed to that of the majority of its members both before and after the formation of the collective belief. Were the contrary personal views of the members to come out into the open after the formation of the collective belief it might cause some degree of shock or at least surprise. It is unlikely, though, to be seen as a rebuttal of the claim about what the group believes or to show that the group is 'of two minds'.

The intuitive disjunction between personal and collective beliefs may be further brought out by considering an example in which two (or more) distinct groups have the same members. Suppose the members of a certain court also constitute a poetry discussion group. (The justices decided it would be a good thing for them to get to know each other better outside the courtroom and since all liked poetry they decided to form such a group.) One day the court has to decide the merits of a certain poem, in conjunction with a certain legal action. While in session it brings various expert witnesses before it, and so on. At a certain point it may be reasonable to judge that the court believes that the poem has great merit. The justices' poetry discussion group, meanwhile, has not yet discussed the poem. Failing special circumstances, then, it may *not* be reasonable to judge that the poetry discussion group believes the poem has great merit. (If it ever does come to discuss the poem, it might come to a different conclusion. That would not necessarily affect the opinion of the court.) In short, one group believes that such-and-such, another doesn't, and the members of both groups are the same. That these members personally believe that such-and-such, if they do, cannot be what determines the beliefs of either group, logically speaking. This example is instructive in various ways but for now I use it only to emphasize the intuitive disjunction at issue.

My discussion so far argues for the rejection of the simple summative account of collective belief statements, according to which these are in place

when all or most members of the relevant group have a given belief.

This conclusion accords with central statements of Durkheim in the *Rules*. He says there that the generality of a given belief – that the group members generally have the belief – is not what makes it collective. He also makes the following point: a belief's being collective may *account for* its generality. One may be helped to see how this second point could hold by considering the observation discussed in the next section.

### An observation concerning the standing to rebuke

In attempting to formulate an alternative to the simple summative account the following should be born in mind. Once a group belief is established, the parties understand that any members who bluntly express the opposite belief lay themselves open to rebuke by other members. Opposed to a 'blunt' expression is one that makes it clear that the speaker is, as he might well say he is, 'speaking personally'.<sup>18</sup>

A rebuke is a form – albeit a mild form – of punishment.<sup>19</sup> One needs a special standing to rebuke someone just as one needs a special standing to punish him or her in other, less verbal ways. In other terms, if asked 'what puts you in a position to punish me?' one cannot answer 'nothing' and still reasonably claim that one is indeed *punishing* the person in question. One can, of course, act in a 'punishing' manner without any special standing. That is, one can inflict some kind of unpleasantness upon the people in question in response to some action of theirs. But a special standing is required for punishment proper.

That one has the standing to punish someone for certain behavior does not mean that in a given situation punishing him will be *justified* all things considered. It means only that one's punishing that person is *possible*.

Why is this matter of standing important? I take it that many people would hesitate to respond in a 'punishing' way to someone's behavior without an understanding that they had the standing to punish that person. In other words, *given* such an understanding, they are more likely to respond in a punitive fashion. If one is averse

to unpleasant treatment, then, the knowledge that certain others have the standing to punish one for certain behavior is apt to dissuade one from engaging in it.

Here is a hypothetical example of these processes in action. Bob and Judy collectively believe that the conservation of species is an important goal. When this collective belief was formed Judy herself personally believed the opposite. Now that the collective belief is in place, Judy recognizes that Bob has the standing to rebuke her for expressing a view in conflict with it. This leads her to suppress any inclinations bluntly to assert that the conservation of species is unnecessary. Such inclinations initially come to her at times, but their incidence diminishes as a result of such regular suppression. Indeed, the belief that prompts them subsides. This will leave the field open for what she and Bob collectively believe to become her personal belief. If she regularly mouths the collective belief, or hears Bob mouthing it, this may well happen.

Thus the collectivity of a belief may lead to its being generally held at the personal level in the population in question. The process envisaged accords with Durkheim's characterization of collective beliefs (and related phenomena) as having *coercive power* and certainly helps to substantiate his claim that the collective nature of a belief may lead to its generality.

That the other parties have the standing to rebuke the member in question appears to be a function of the collective belief itself. An adequate account of collective beliefs as these are ordinarily conceived of, then, should explain how a collective belief gives the members this standing. In addition to the other problems with it, it is hard to defend the simple summative account in light of this consideration.<sup>20</sup>

More complex summative accounts may be proposed. These take the simple summative account and add further conditions such as common knowledge of the widespread nature of the belief in question. I shall not pursue any such account here. As I shall explain, I believe that a quite different type of account is more promising.<sup>21</sup>

### The plural subject account of collective belief

Three criteria of adequacy for an account of collective belief are suggested by the foregoing discussion. It should explain how the existence of a collective belief that *p* could give the parties the standing to rebuke each other for bluntly expressing a view contrary to *p*. It should not suppose that all or most of the parties must personally believe that *p*. Nor should it suppose that if all or most of them believe something then they collectively believe it.

Among other things, the account of collective belief that I have developed meets all of these criteria. It also respects an important suggestion derived from examples like those considered here. Roughly, it is both necessary and sufficient for members of a population, *P*, collectively to believe something that the members of *P* have openly expressed their readiness to let the belief in question be established as the belief of *P*. In its present formulation, this suggestion begs for some unpacking. Its gist is evidently this: the members must express *something* such that – as all understand – once all have openly expressed it, the truth conditions for the ascription of the relevant belief to population *P* are satisfied.

So, what is this something that must be expressed? This question is best approached by first elaborating the *outcome* of the giving of these expressions. To be sure, one way of describing this outcome is ‘The members of population *P* collectively believe that *p*’. There is another way of describing it, however, that is more helpful. In terms I shall shortly explain: ‘The members of Population *P* are jointly committed to believe as a body that *p*.’ In other words, my account of collective belief is roughly as follows.<sup>22</sup>

*A population, P, believes that p* if and only if the members of *P* are jointly committed to believe as a body that *p*.

The phrase ‘joint commitment’ is a technical one.<sup>23</sup> I take the *concept* it expresses to be implicit in everyday discourse.<sup>24</sup> Indeed, I have argued that it is fundamental to our everyday conceptual scheme concerning social relations.<sup>25</sup> It is not, then, a technical concept.

First I should say something about the relevant

general concept of *commitment*. This is the concept involved in the intuitive judgment that if Pam, say, decides to go shopping today, she is committed to doing so. In the case of a personal decision, the commitment is *personal*. By this I mean that the one whose commitment it is creates it unilaterally and can unilaterally rescind it. I take it that, once committed, Pam *has reason* to go shopping today, unless and until she changes her mind. Her commitment does not make her shopping today more valuable in itself; unless and until it is rescinded, however, she is rationally required to go shopping. For the sake of a label such commitments may be termed *commitments of the will*.<sup>26</sup>

I do not claim to have plumbed the depths of the general concept of commitment at issue here. Doing so is an important task. Meanwhile I find promising the idea that commitments of the will may usefully be represented as involving a type of self-directed order or command.<sup>27</sup>

The concept of a *joint* commitment is the concept of a commitment of two or more people. A joint commitment so conceived is not something composite, a conjunction of the personal commitment of one party with the personal commitment(s) of the other(s). Rather, it is simple. A joint commitment is the creation of all the parties to it, rescindable only with the concurrence of all. Insofar as it involves a type of self-directed order, it involves an order issued jointly by all the parties to all the parties.<sup>28</sup>

Generally a given party will be physically capable of acting against an unrescinded joint commitment, if he or she so desires. But then the commitment will have been violated, as all of the parties will understand.<sup>29</sup>

A joint commitment is created only when each of the parties has, in effect, openly expressed his or her personal readiness to be party to it. That these expressions have been made openly must be common knowledge in the relevant population.

By virtue of their participation in a *joint* commitment, the parties gain a special standing in relation to one another’s actions. Here I summarize without argument important aspects of that standing.

The parties to a joint commitment are *answerable* to one another with respect to their

non-conformity. Further, one who violates a joint commitment has *offended against* all of the parties to the joint commitment, as such. The offense in question can plausibly be characterized in terms of a *violation of right*. In other words, when I am subject to a joint commitment requiring me to do certain things, all of the parties to the commitment have a right to the relevant actions from me.<sup>30</sup> Correlatively, I am *under an obligation* to all of them to perform these actions.<sup>31</sup> These are obligations and rights of a distinctive kind, deriving immediately from the joint commitment. Once it exists, they exist also, irrespective of the surrounding circumstances.<sup>32</sup>

In consequence of the existence of these rights and obligations, failing special background circumstances, those who are party to a joint commitment have the standing to demand that others conform to it, if non-conformity is threatened, and to rebuke one another for defaults that have taken place. For the following can be argued.

Conformity to the commitment is, or was owed to him as a party to the commitment. It is therefore in an important sense his. This gives him the standing to demand that it be 'given' to him or to rebuke one who has withheld it from him.<sup>33</sup>

A given joint commitment can always be described in a sentence of the following form: the parties are jointly committed to *X as a body*. Acceptable substitutions for 'X' are psychological verbs such as 'believe', 'intend', and so on.

What is it to be jointly committed to *X as a body*? As I understand it, this joint commitment has two related aspects. First, the parties are jointly committed to bring it about, as far as is possible, that the parties constitute a single body that *Xs*. Second, it is understood that their doing so is to be a function of the joint commitment in question.

Putting these points together: the joint commitment to *X as a body* is a joint commitment to bring it about that, as far as is possible, the parties constitute a single body that *Xs*, and to do so in light of the joint commitment in question. Evidently, the end in question – bringing it about that, as far as is possible, the parties constitute a single body that *Xs* – is to be brought about by virtue of the activity of the individual parties, each acting in light of the joint commitment.

The guiding idea of a *single body* that *Xs* includes nothing about the intrinsic nature of the single body in question. In particular, it does not imply that it is in some way made up of two or more distinct bodies that are capable of *X-ing* on their own.

I use the phrase 'plural subject' as a technical term to refer to those who are jointly committed to *X as a body*, for some *X*.<sup>34</sup> Those who are jointly committed to *X as a body* constitute, by definition, the *plural subject* of *X-ing*. Accordingly, I shall dub my account of collective belief the plural subject account.

As should be clear from the previous paragraph, I am *not* supposing that those who are jointly committed to *X as a body* are jointly committed to constitute, as far as is possible, a *plural subject* of *X-ing* in my technical sense, or, indeed, in any *other* appropriate sense. By definition, they already constitute a plural subject of *X-ing*, in my technical sense, once they are jointly committed to *X as a body*.<sup>35</sup>

Those who are jointly committed to *X as a body*, then, have as their guiding idea that of a single body that *Xs*. They are to bring it about that, as far as is possible, they constitute such a body. It is up to them to figure out exactly how this may be done.<sup>36</sup>

What the joint commitment requires of each party is that *he should do what he can* to bring it about, *in conjunction with the other parties*, that they constitute, as far as is possible, a single body that *Xs*. This is likely to involve some degree of monitoring of others' behavior on his part.

My account of collective belief involves a joint commitment to *believe* something as a body. I shall now focus on this particular form of joint commitment.

A joint commitment to believe that *p as a body* does *not* require each participant personally to believe that *p*. It would seem to follow from the fact that the requirement at issue is precisely to constitute, as far as is possible, a single body that believes that *p*. It does not concern any *other* bodies that may bear some relation, however close, to the body in question.<sup>37</sup>

More positively, the joint commitment will be fulfilled, to some extent at least, if those concerned say that *p* in appropriate contexts, with an appropriate degree of confidence, and

do not call *p* or obvious corollaries into question. Their behavior generally should be *expressive of the belief that p*, in the appropriate contexts. That does not mean, as said, that they must personally have that belief. In other words, this expressive behavior need not be *the expression of a personal belief that p*.

As just indicated, certain contextual conditions are likely to be understood. Thus, for example, members of a seminar on human rights may in the course of a meeting form a joint commitment to believe as a body that the notion of a group right is a viable one. This would involve a requirement to express that belief at least within the confines of the seminar when it is in session. More broadly, it would require that each party express that belief *when acting as a member* of the seminar. Presumably the parties are not always so acting. If a friend who is not a member of the seminar engages one of them on the topic while they are out on a hike, it would presumably be appropriate for each to speak *in propria persona*, without preamble.<sup>38</sup>

Suppose one is in a context where it is appropriate to act in accordance with a given joint commitment to believe something as a body. It is then open to one to use such qualifiers as 'Personally speaking' to preface the expression of a belief contrary to the collective one. This makes it clear that one is indeed now speaking for oneself and not as a member of the relevant collective.

Thus the plural subject account of collective belief accords with the logic of 'We (collectively) believe that *p*' as this is understood in everyday life. In particular, it allows for the possibility that a party to the collective belief aver without fault that he *personally* does not believe that *p*.

Though such avowals are not ruled out, there is likely to be some cost attached to them. One thereby makes it clear that one's personal view differs from that of the collective in question. Other members may subsequently regard one with suspicion, thinking one more liable to default on the joint commitment, either inadvertently or deliberately. They may begin preemptively to think of one as an 'outsider', as no longer 'one of "us"'. If one does default, of course, they have the standing to rebuke one for doing so. All of this accords with Durkheim's perception of 'social facts' generally as having coercive power.

It is clear that collective beliefs according to the plural subject account are likely to suppress the development of contrary ideas at both the individual and the collective level. This has many practical implications. For instance, it suggests that groups of researchers in various fields of science are likely to benefit from the less constrained insights of outsiders and neophytes – assuming that these fields involve a collective characterized by a network of collective beliefs. There are, certainly, well-known cases in which the insights of such people have transformed a field.<sup>39</sup>

### Further questions

I have offered an account of the phenomenon I take to correspond to everyday ascriptions of collective belief. I have argued that this phenomenon, where it exists, is extremely consequential with respect to human thought and behavior. It is, in other words, a highly significant phenomenon. In what follows I shall refer to this phenomenon as *collective belief*, assuming the accuracy of the account.

Is there any reason that the epistemologist should refuse to take collective beliefs seriously? In concluding, I briefly consider three concerns that might be raised on this score.

The first concern goes back to everyday collective belief ascriptions. It was on the basis of a variety of data relating to such ascriptions that an account of collective belief was proposed. The concern runs as follows. Why should the epistemologist take everyday collective belief ascriptions seriously? After all, there are many ascriptions of belief that should not be so treated. Consider the following possibilities. In a story for children a train is said to believe something. A computer user finds himself saying in response to an error message that his computer thinks he is stupid. Someone says of his cat: 'She thinks she is the President of the United States'.

I agree that many belief ascriptions can be excluded from consideration by epistemologists as not *intended* to be taken seriously or literally. Collective belief ascriptions, in contrast, are commonly made with serious, literal intent. That does not necessarily let them off the hook. As I said at the beginning, they could involve some



kind of mistake. That they do involve some kind of mistake, however, needs to be shown.

I suggest that showing it is going to be a harder task than is often assumed. Consider this. A tempting way of arguing that collective belief ascriptions involve a mistake would be to bring up some feature alleged to be essential to belief and to argue that, since so-called collective beliefs lack that feature they are not beliefs – nor, perhaps, can they be counted as any other kind of cognitive state. One who wishes to maintain that collective belief is indeed belief can argue that, on the contrary, if collective beliefs lack the allegedly essential feature of belief, that throws doubt on the claim that this feature is indeed essential to belief.

Another concern that several philosophers have raised both informally and in print is whether collective belief – on the plural subject account – is better thought of as *acceptance* rather than *belief*.<sup>41</sup> I have dubbed this position “rejectionism” and replied to some initial statements of it elsewhere.<sup>42</sup> Discussion on the topic has continued.<sup>43</sup> I shall not attempt to summarize this material here. Indeed, I have not yet had a chance myself to review all of it. Suffice it to say that my own tendency is to go with everyday discourse, as I understand it. If such-and-such a phenomenon is referred to, seriously and literally, as belief, it is hard to argue that it is not, after all, belief *as this is ordinarily understood*. Four related points follow.

First, given that individual human beings differ in important ways from populations comprised of such beings, it is likely enough that collective beliefs differ in important ways from the beliefs of individual human beings. If there are significant contrasts to be made between beliefs of these two kinds, that is certainly worth pointing out.

Second, it is possible that the claim that collective ‘belief’ is acceptance rather than belief could be argued for one or more pairs of stipulative definitions for the relevant terms. I have taken the issue to concern the everyday concept of belief whatever precisely that amounts to. It should be noted, in connection with this point, that different philosophers have proposed a wide variety of contrasts between ‘belief’ and ‘acceptance’.

Third, it is interesting to note that many of the

features traditionally claimed to characterize belief in general can be argued to characterize collective beliefs. In order to see this it is important carefully to distinguish what is true of the members collectively and what is true of them as individuals.<sup>44</sup>

Fourth, I do not deny that there is such a thing as *collective acceptance* according to the various accounts of acceptance that have been proposed. Though I have focused on collective belief in this paper, I assume that, generally speaking, there is a collective cognitive state that corresponds to a given cognitive state of individual human beings: we believe, we know, we accept, we doubt, we conjecture, and so on.

The final concern I address here returns us to the question of the relevance of collective epistemology to general epistemology. I proposed that epistemologists should consider what it is for the members of a population collectively to believe something, before settling on a general account of belief. Now it turns out that if I am right about collective belief, those who collectively believe something must have the general concept of belief. Is there a problem here?

On the contrary, the plural subject account of collective belief suggests a certain test for an acceptable account of belief in general. Suppose that according to a proposed account of belief in general S believes that p (for any S) if and only if S Fs (for some particular F). This will be acceptable only if a population whose members are jointly committed to F as a body thereby Fs.

Here is an example of the application of this test. Suppose one holds that S believes that p if and only if S is generally disposed to act as if p is true. Is it the case that a population whose members are jointly committed to be disposed-as-a-body-to-act-as-if-p-is-true is thereby disposed to act as if p is true? If so, the account passes the test in question. That does not mean, of course, that it is the correct account of belief, only that it is consistent with the plural subject account of collective belief. That an adequate account of belief in general is so consistent is quite possible in principle.

## Summary and conclusion

I have argued for the importance of collective epistemology both in its own right and in relation to general epistemology. As to the latter, if we are to develop adequate general theories of knowledge, belief, and so on, we need to take the collective versions of these phenomena into account.

Apart from the matter of such general theories, if we are to make sense of the human world, we need to understand the nature and functioning of collective cognitive states as well as the nature and functioning of the cognitive states of human individuals. We also need to understand the

relationship between these two types of cognitive state, for example, the ways in which they may influence one another.

In this essay I focused on the topic of collective belief. I argued against an account that is likely to spring to mind when one first considers the question, and in favor of the account I labeled the plural subject account. This alone satisfies the criteria of adequacy for such an account that I specified.

The subject of collective epistemology is still in its infancy. It is to be hoped that more and more theorists will pay it the attention it deserves, so it grows to maturity.\*

---

## Notes

- <sup>1</sup> I take it that if what is referred to on a given occasion as a 'family', for instance, is understood to have but a single member, the ascription of a cognitive state to it would not count as the ascription of a *collective* cognitive state. Nor would a single-membered 'family' count as a paradigmatic *group*.
- <sup>2</sup> On the relatively narrow sense of 'group', or 'social group', in question see *On Social Facts*, chapter 4. I also explore there the conditions for the appropriate use of the first person plural pronoun 'we', which had been neglected by theorists of other such expressions.
- <sup>3</sup> See for instance the essays in Frederick Schmitt, ed., 1994, *Socializing Epistemology*, Rowman and Littlefield: Lanham, MD; see also, in Italian, *Esperienza e Cognoscenza*, 1995, ed. Gianguido Piazza, Milan: Citta Studi, which translates a number of texts that are generally taken to be contributions to social epistemology. Each of these collections includes at least one essay on a collective cognitive state, but the majority of the essays are concerned with other matters.
- <sup>4</sup> Deborah Tollefsen uses this phrase in a roughly similar way in a paper presented at Leipzig University, June 2004, "Collective Epistemic Agency and the Need for Collective Epistemology" (2004ms).
- <sup>5</sup> See, for instance, Charles Taylor, 1980, Critical notice, Jonathan Bennett's *Linguistic Behavior*, *Dialogue*, vol. 19; Charles Taylor, 1985, *Human Agency and Language*, Cambridge University Press: Cambridge, chapter 10, 'Theories of Meaning', section III.1. 'Common knowledge', undoubtedly an important topic for epistemology, has been variously defined. For present purposes what is important about these definitions is that they generally refer to individual human beings, what each knows, what each knows about each one's knowledge, and so on. Two classic sources are David Lewis, 1969, *Convention: A Philosophical Study*, Harvard University Press: Harvard; and Stephen Schiffer, 1972, *Meaning*, Oxford University Press: Oxford. See also Margaret Gilbert, 1989, *On Social Facts*, Princeton University Press: Princeton.
- <sup>6</sup> Margaret Gilbert, 'Modeling Collective Belief', 1987, *Synthese*, reprinted in Margaret Gilbert, *Living Together: Rationality, Sociality and Obligation*, 1996, Rowman and Littlefield: Lanham, MD, and Gilbert, *Social Facts*. Due to the vagaries of publishing, the 1987 article was written after the 1989 book was sent to the press. I have developed my position in a number of subsequent publications including, most recently, 'Belief and Acceptance as Features of Groups', 2002, *Protosociology*, vol. 16, 35-69.
- <sup>7</sup> Raimo Tuomela, 1992, 'Group beliefs', *Synthese* vol. 91, 285-318, responds to my proposals in 'Modeling' and *Social Facts* and presents an alternative but in many ways similar account. See

also his book *The Importance of Us*, 1995, Stanford University Press: Stanford. Other philosophers who have addressed one or more aspects of the topic in the 1990s or later include, in alphabetical order, Alban Bouvier, Austen Clark, Angelo Corlett, Christopher MacMahon, Antonie Meijers, Philip Pettit, Gianguido Piazza, Gerhard Preyer, Abraham Sesshu Roth, Frederick Schmitt, Deborah Tollefsen and Bradley Wray. Preyer edits the online journal *Protosociology* in which relevant publications by a number of these authors have recently appeared in special issues: web address: [www.protosociology.de](http://www.protosociology.de).

- <sup>8</sup> Relatively few of the authors mentioned in the last footnote specialize in mainstream epistemology.
- <sup>9</sup> See Emile Durkheim, 1982, *The Rules of Sociological Method*, translated from the French by W. D. Halls, Free Press: New York (first published 1895). I suggest an interpretation of key passages in this work in *On Social Facts*, chapter 5, and subsequently in 'Durkheim and Social Facts', 1994, in *Debating Durkheim*, Herminio Martins and William Pickering, eds., Routledge: London, reprinted in French translation in Margaret Gilbert, *Marcher Ensemble: Essais sur les Fondements de la Vie Collective*, 2003, Presses Universitaires de France: Paris. As I interpret him Durkheim presents us with an intriguing, nuanced proposal that requires further elaboration. My own proposals on the topic indicate one way in which such elaboration out might proceed.
- <sup>10</sup> In *On Social Facts* I referred to my approach as 'conceptual analysis', a phrase open to various construals. The preceding description of my approach may be less tendentious. The main point is that everyday judgments about when, say, a group believes something, and when it fails to believe something are carefully considered to arrive at a description of what (according to everyday understandings) collective belief is. No substantive 'theory of concepts' is presupposed.
- <sup>11</sup> Angelo Corlett pays attention to the question of collective knowledge in his book *Analyzing Social Knowledge*, 1996, Rowman and Littlefield: Lanham, MA.
- <sup>12</sup> See, for instance, Anthony Quinton, 1975, 'Social Objects', *Proceedings of the Aristotelian Society*, vol. 75. Quinton assumes the simple summative account *en passant*. Related assumptions have been made for specific forms of belief or judgement, such as moral judgement. On this see Margaret Gilbert, 'Shared Values, Social Unity, and Liberty', forthcoming in *Public Affairs Quarterly*, e.p.d. January 2005.
- <sup>13</sup> More extensive discussions are to be found in Gilbert, 'Modeling', and Gilbert, *Social Facts*.
- <sup>14</sup> I develop an example involving a poetry discussion group in more detail in Gilbert, 'Modeling'. The example derives from an actual, unexpectedly long-standing group founded by the author together with Mairie McInnes and Priscilla Barnum.
- <sup>15</sup> I take interpersonal discussions generally to be a form of acting together or collective action. I offer accounts of such action in *On Social Facts*, chapter 4, and elsewhere. See, for instance, 'What is it for *Us* to Intend?' 1997, in *Contemporary Action Theory*, vol. 2, *The Philosophy and Logic of Social Action*, G. Holmstrom-Hintikka and R. Tuomela, eds, reprinted in Margaret Gilbert, 2000, *Sociality and Responsibility: New Essays in Plural Subject Theory*, Rowman and Littlefield: Lanham, MD. These accounts are formally similar to the account of collective belief sketched here.
- <sup>16</sup> For an example of how there might be a complete lack of corresponding personal beliefs see Gilbert, *Social Facts*, p.290.
- <sup>17</sup> This might not be so in special circumstances. For instance, the group has previously opined that a given poet is so talented he is incapable of writing a poem that is less than brilliant.
- <sup>18</sup> See Gilbert, 'Modeling'; also 'Remarks'.
- <sup>19</sup> Cf. H. L. A. Hart, 1961, *The Concept of Law*, Clarendon Press: Oxford. Hart sees 'informal reproofs' and legal punishment as in the same category.
- <sup>20</sup> For related discussion see Gilbert, 'Modeling'; Gilbert, *Social Facts*, Chapter 5.
- <sup>21</sup> For detailed discussion and critique of two such complex summative accounts see Gilbert, *Social Facts*, Chapter 5.
- <sup>22</sup> I have formulated this account differently on different occasions. Essentially the same idea is

expressible in different ways. The Introduction to my book *Living Together: Rationality, Sociality, and Obligation*, 1996, Rowman and Littlefield: Lanham, MD, pp. 7–10, explains the way the main formulations I have used relate to one another. See also Gilbert, *Social Facts*, Chapter 7, on different formulations of the general schema of which this is an exemplification.

<sup>23</sup> Other theorists have used the same phrase; sometimes they have not indicated what they mean by it; at other times the sense explicitly given has been different to the one I give. In any case I am concerned with joint commitment in my own technical sense here.

<sup>24</sup> For a more extended discussion see the Introduction to *Living*, pp. 7–15.

<sup>25</sup> See Margaret Gilbert, *Social Facts, Living, and Sociality and Responsibility: New Essays in Plural Subject Theory*, 2000, Rowman and Littlefield: Lanham, MD.

<sup>26</sup> For more on such commitments see Margaret Gilbert, forthcoming, 'Towards a Theory of Commitments of the Will: On the Nature and Normativity of Intentions and Decisions', Wlodek Rabinowicz, ed., *Patterns of Value 2*, Lund University: Lund. For a distinction between 'reasons' and 'rational requirements' see John Broome, 'Are Intentions Reasons? And How Should we Cope with Incommensurable Values?', 2001, in *Practical Rationality and Preference: Essays for David Gauthier*, eds. Christopher Morris and Arthur Ripstein, Cambridge University Press: Cambridge, pp. 98-120.

<sup>27</sup> On this see Margaret Gilbert, 1999, 'Obligation and Joint Commitment', *Utilitas*, reprinted in Gilbert, *Sociality*; this cites Anthony Kenny, 1963, *Action, Emotion, and Will*, Routledge and Kegan Paul: London.

<sup>28</sup> For a more extended discussion of joint commitment see Margaret Gilbert, 2003, 'The Structure of the Social Atom: Joint Commitment as the Foundation of Human Social Behavior', in Frederick Schmitt, ed., *Social Metaphysics*, Rowman and Littlefield: Lanham, MD; see also the Introduction, *Living*, pp. 7–15, and 'Obligation'.

<sup>29</sup> For discussion of the consequences of such violation see the Introduction, *Living*, pp. 14–15. I there incline to the view that generally speaking violation by one or more parties renders a joint commitment rescindable by the remaining parties, as opposed to nullifying it.

<sup>30</sup> I focus on the relationship of joint commitment and rights in my book *Rights Reconsidered*, to be published by Oxford University Press.

<sup>31</sup> See Gilbert, 'Obligation', and elsewhere.

<sup>32</sup> For discussion of the distinctiveness of these obligations see my article 'Agreements, Coercion, and Obligation', *Ethics*, 1993, vol. 103, reprinted with some revisions in *Living*.

<sup>33</sup> I assume here that a rebuke is the after-the-fact correlate of a demand.

<sup>34</sup> See Gilbert, *Social Facts, Living*, and other writings.

<sup>35</sup> I emphasize the point in response to various comments.

<sup>36</sup> For more on this see the text below.

<sup>37</sup> See Margaret Gilbert, 'Shared Values, Social Unity, and Liberty', forthcoming in *Public Affairs Quarterly*, e.p.d. January 2005.

<sup>38</sup> On the question of context see also Margaret Gilbert, 'Remarks on Collective Belief', 1994, in Frederick Schmitt, ed., *Socializing Epistemology*, Rowman and Littlefield: Lanham, MD (revised as 'More on Collective Belief' in *Living*).

<sup>39</sup> For further discussion of the role of collective beliefs in science see Gilbert, 2000, *Sociality*, chapter 3, 'Collective Beliefs and Scientific Change', first published in Italian in 1998.

<sup>40</sup> I make this point in response to a query from the developmental psychologist Letitia Nagles in discussion at a University of Connecticut cognitive science colloquium. A related point was brought up by the anthropologist Maurice Bloch in discussion at a roundtable on collective belief held at the Sorbonne, Paris.

<sup>41</sup> See, for instance, K. Bradley Wray, 2002, 'Collective Belief and Acceptance', *Synthese*, 129: 319-333; Anthonie Meijers, 1999, 'Believing and Accepting as a Group', in *Belief, Cognition and the Will*, ed. A. Meijers. Tilburg: Tilburg University Press. 59-71; also A. Meijers, 2002, 'Collective Agents and Cognitive Attitudes', *Protosociology*, 16: 70-86.

<sup>42</sup> See Gilbert, 'Belief and Acceptance'.

<sup>43</sup> See K. Bradley Wray, 2003, 'What Really Divides Gilbert and the Rejectionists', *Protosociology* 17, and related papers, in the same volume, by Christopher MacMahon, Anthonie Meijers, Deborah Tollefsen. See also Abraham Sesshu Roth, 2003ms, 'Remarks on Collective Belief and Acceptance'.

<sup>44</sup> See Gilbert, 'Belief and Acceptance'.

\* A version of this paper, 'Collective Belief as a Subject for Cognitive Science', was given to the cognitive science group at the University of Connecticut, Storrs, CT, February 14 2003, and to the Roundtable on Collective Belief at the Sorbonne, Paris, April 4 2003. I thank those present for discussion, and Ludger Jansen for some pertinent conversation on a related topic, August 2004.

Margaret Gilbert is Professor of Philosophy at the University of Connecticut. Her research has centred on the nature of social phenomena such as social groups (social ontology), social conventions, group beliefs and emotions, shared (or collective) intention and action. Her publications include *Rights Reconsidered* (forthcoming); *Marcher Ensemble: Essais sur les Fondements des Phenomenes Collectif* (2003); *Sociality and Responsibility: New Essays in Plural Subject Theory* (2000); *Living Together: Rationality, Sociality, and Obligation* (1996); and *On Social Facts* (1989).