WHEN IS THERE A GROUP THAT KNOWS? DISTRIBUTED COGNITION, SCIENTIFIC KNOWLEDGE, AND THE SOCIAL EPISTEMIC SUBJECT

Abstract

I pose three questions. (i) When does a collection of individuals form an entity that is more than just the mereological sum of its constituent persons? (ii) Given that there is a group of this sort, under what conditions does it know (or believe etc.)? (iii) When we talk of, for example, 'the growth of scientific knowledge', can we regard this scientific knowledge as an epistemic state of some social entity? I draw upon ideas from distribution cognition and Durkheimian sociology to provide responses to the first and second questions, and thereby give a positive answer to the third.
Keywords: collective belief; group knowledge; joint commitment; distributed cognition; functionalism; Émile Durkheim; Margaret Gilbert.

1 Introduction

When is there a 'group', 'collectivity', or 'social system' that knows? That question may be broken down into two more specific questions. The first is the ontological question of the existence of a group: (i) when does a collection of individuals form an entity that is more than just the mereological sum of its constituent persons? The second, epistemological, question concerns the epistemic (and doxastic) states of such a social entity: (ii) given that there is a group of this sort, under what conditions does it know (or believe etc.)? In this paper I address these questions in order to answer a third: (iii) we talk of 'scientific knowledge' in a broad sense, as when we refer to the growth of scientific knowledge'; can we regard this scientific knowledge as an epistemic state of some social entity?

We can find many terms to talk of social entities (as indicated in my opening sentence). I shall disregard any possible differences between their senses there may be. On the assumption (yet to be justified) that there are social entities that can know things, I shall talk of a 'social epistemic subject'. So our third question can be framed as asking whether science involves a social epistemic subject?

My approach will be to compare two approaches to group knowledge, which I call the *commitment* model and the *distributed* model. The commitment model is the one most favoured by authors endorsing the claim that there are social epistemic subjects. I shall argue, however, that the distributed model offers a better picture of the social epistemology of what I will call 'wider science'. One might conclude from this that (wider) science does not after all involve a social epistemic subject. But I shall argue that there is indeed a social epistemic subject for science. And so we

should conclude that satisfying the commitment model is not a necessary condition for the existence of a social epistemic subject. I shall also ask what constitutes the social epistemic subject in the case of wider science.

2 Two models of social cognition

The commitment model

Collective epistemic activity often takes the following form. A committee is formed in order to make recommendations on a particular question. The committee gathers evidence. It discusses the evidence. By some agreed procedure, it reaches a conclusion and submits its recommendations. What constitutes the existence of the committee in addition to its members? Relevant considerations include the fact that the committee has a purpose, a purpose to which the committee's members are individually committed in their role as members of that committee. The pursuit of that purpose is guided by certain norms, which may be explicit or implicit, to which the members are also committed. And when the committee reaches its conclusion, by the appropriate mechanism, the individual members are committed to that conclusion, *qua* committee members. That commitment may, according to some views, involve public endorsement of the conclusion, even if the individual's private opinion differs.

This committee provides one kind of model for thinking about social epistemic subjects. In this model two related features are prominent: first, the norms governing the committee, and, secondly, the commitment of the individuals to those norms. Opinions may differ about what commitments and norms must be involved. Nonetheless, the picture is clear of what an ideal social epistemic subject is like; it is an epistemic version of a sporting team—a tug-of-war team might be the best analogy—with each participant committed to their collective goal and to an agreed, commonly understood mechanism for reaching that goal. The justification for engaging in this form of group cognition is the thought that two (or more) heads are better than one.[1]

The commitment model of group epistemic and doxastic states is supported by a number of philosophers, including Raimo Tuomela (1992, 2004), Margaret Gilbert (1987, 1989, 2004), and Frederick Schmitt (1994). Philip Pettit (2003) also endorses the commitment model as a necessary but not sufficient condition on the existence of a 'social integrate'. With regard to belief, Gilbert (1989: 306) and Schmitt (1994: 262) hold that:

> (G) A group G believes that p just in case the members of G jointly accept that p, where the latter happens just in case each member has openly expressed a willingness to let p stand as the view of G, or openly expressed a commitment jointly to accept that p, conditional on a like open expression of commitment by other members of G.

Gilbert (2004) summarises the kind of commitment involved in (G), whereby several individuals come together and commit themselves as a body to something, as *joint commitment*. So (G) says that a group believes that p when its members are jointly

[1] Whether that is in fact true is another matter. See Kanai and Banissy (2010).

committed to believing as a body that *p*. Assuming that knowing entails believing, then (G) places necessary conditions on group knowledge.[2]

As a consequence of the shared norms and commitments entailed by (G), individuals making up such a epistemic social subject will typically be conscious of and involved the epistemic work of the committee. If indeed public commitment to the conclusion is required, then we would expect that individuals, in order to be comfortable with such a commitment, will want to know how the conclusion is reached and to be able to influence that conclusion.[3] Even without the requirement of commitment to the conclusion, merely being associated with that conclusion will often be enough to provide the motivation for understanding of and involvement with the decision-making process. Some authors argue that such consciousness is among the norms governing the social epistemic subject.[4]

Since this model has the characteristic that the group's members will be aware of one another as members of the group and aware of the group's modus operandi, we might also call this the *mutual awareness* model, following Pettit (2003).

The distributed model

Often we encounter circumstances such as the following. A complex task is information driven. But the information cannot be obtained and processed by any one individual. So several individuals are given roles in gathering different pieces of information while other individuals have the task of coordinating this information and using to complete the task. A famous example from Edwin Hutchins's (1995a) *Cognition in the Wild* concerns the process of bringing a large ship, the *USS Palau*. safely into port. Several crew members are given different landmarks whose bearings they are required to record and to communicate to a plotter who determines the ship's position and course. Another of Hutchins's (1995b) examples concerns the distribution of tasks between the two pilots (the Pilot Flying, PF, and the Pilot Not Flying, PNF) on a commercial aeroplane, in describing the task of determining the correct airspeeds at the various point of the plane's descent when coming into land.

The key feature of such examples of the production of knowledge by a group is that the task is broken down into components, which are given to different members of the group. Membership of the group in this case is a matter of having a particular function within the overall system. This kind of group knowledge production is known, following Hutchins, as *distributed cognition*. The distributed cognition approach can be characterised in a number of ways; for current purposes, I propose the following. Distributed cognition identifies a *system* for producing knowledge, and studies how the various components, usually performing distinct sub-tasks, contribute to that overall task.

[2] Gilbert (1987: 192; 2004: 99) regards the commitment in (G) as strong, imposing an obligation on group members not to disavow the group belief *p*, such that any disavowal may be rebuked by other group members. I am unsure that this is an inevitable feature of the groups that (G) seeks to describe, and suspect that it varies from case to case. Perhaps group members may still affirm that *p* is the view *of the group* even if not of all of them individually. For example a court may come to an opinion that some of the judges openly disagree with in a minority opinion.

[3] This point will have greater force where the commitment has the strength claimed by Gilbert, as mentioned in footnote 2.

[4] I note though that Gilbert does not require that all members of the group always know what views the group collectively espouses. The group may come to a conclusion in the absence of some members, who are nonetheless committed to the group view. They may then ask, 'What do we think bout X?'.

In typical cases, including those pertinent to social epistemology, the components of interest are usually individual human subjects. There is division of cognitive labour—the different human participants are given distinct cognitive sub-tasks. As Hutchins has emphasized from the outset, almost always in such systems there are non-human vehicles for representing various pieces of information; these must be taken fully into account in articulating how the system achieves its cognitive goals. For example, in the navigational task there are distinct cognitive tasks for the sailors involved, such as taking bearings and plotting the vessel's course. In carrying out those tasks an alidade is used to measure and record the bearings, while the plotter has specially constructed maps and instruments in order to use the bearings to calculate the correct course. In the aeroplane, the PF controls the plane while the PNF operates the aircraft systems and has responsibility for completing the checklists required for each phase of flight. In addition, the pilots are aided by a variety of instruments, not only for recording important information, such as airspeed and aircraft weight, but also for calculating the correct airspeeds for the changing wing configurations (the configurations of flaps and slats that change in order to maintain lift as the aircraft reduces its speed during descent). When the correct airspeeds for the different stages of the descent are determined, these are represented by 'speed bugs', moveable markers around the rim of the airspeed indicator; this is done in advance of the information being required, since the pilots will be too busy to calculate the correct arises in real time; hence the need for a system to represent or 'remember' its correct airspeeds.

Some differences between the two models

Groups engaged in distributed cognition typically do not satisfy the commitment model. It may be correct to say that the *USS Palau* knows (and so believes) that it is heading NNW at 14.2 knots, but that is not something which all the crew members have expressed a joint commitment to accept, not even all those involved in navigation. A midshipman may be on deck taking bearings and contributing to the generation that knowledge, but that knowledge is not something he is aware of let alone committed to.

The commitment model leads to questions and puzzles such as Pettit's discursive dilemma (2003). Commitment doesn't guarantee group rationality even if every individual is rational. This is because the opinions of the group's members might be aggregated in a way that leads to a paradoxical outcome.[5] Because distributed cognitive systems employ the division of cognitive labour, rather than aggregation of beliefs, Pettit's discursive dilemma is not an issue for distributed cognitive systems *per se*.

We may note, however, that a committee, the paradigm of the commitment model, may sometimes show some features of distributed cognition. For example, the committee may form specialist subcommittees or may choose to defer to the

[5]Imagine three subjects, A, B, and C, considering three propositions, p, q, and $p \wedge q$. A holds p to be true but q to be false, and so holds $p \wedge q$ to be false. B holds p to be false but q to be true, and so also holds $p \wedge q$ to be false. C holds both p and q to be true and so holds $p \wedge q$ to be true. Now imagine that A, B, and C have committed themselves in conformity with (G) to deciding which propositions to regard as the beliefs of the group by majority vote. Accordingly they regard p as true (A and C vote in favour), and q as true (B and C vote in favour), and they also regard $p \wedge q$ as false (since A and B both vote for the falsity of $p \wedge q$). Consequently the group beliefs are inconsistent. This is the discursive dilemma and according to Pettit a group only becomes an agent if it adopts a belief aggregation procedure that can avoid the dilemma.

expertise of particular members. (By doing so it may also reduce the chances of the discursive dilemma arising.) Even so, a committee will usually bring the distributed knowledge back together to the whole committee so that all can share in the any fruits of distributed cognition. In that way the committee would be atypical of distributed systems, which do not usually share the knowledge produced across the system. The process by which the group comes to know according to the commitment model is not a distributed process—it is a process that is the same for all members of the group, involving them all individually in the same process of cognition.

3 The two models and science

Does either the commitment model or the distributed model provides a satisfactory model for understanding group knowledge in the sciences? The commitment model is generally regarded by its proponents as a model for group knowledge in general, and so if science can generate group knowledge then the model should encompass science too. Tuomela (2004: 109) implies this when he open his paper 'Group knowledge analyzed' saying 'One can speak of knowledge in an impersonal sense: It is accepted as knowledge that copper expands when heated ...'. Distributed cognition is intended to be an approach with widespread application. Nonetheless, its central examples are cases of a well-defined cognitive task, implemented by a carefully designed system, with sub-tasks assigned to individuals who have some understanding of their role within the larger task. While some group activities within science, such as an experiment carried out by a team in a laboratory, will have these features, it is less clear that science on a larger scale is like this. And so it is a genuine question as to whether the distributed model can be applied to the interactions of scientists on a larger scale.

The commitment model and science

Research teams in science might sometimes satisfy the commitment model. The team discusses the results they have been getting and their significance, possible interpretations, problems with the experiment and so forth. They may reach a view which can be expressed as the conclusion of the group, which may be reported as 'We have discovered that *p*; we believe that this shows that *q*'. One would normally expect the team members to avow such conclusions as the view *of their group*, even though some members may have reservations. The larger the research teams the less well they fit the commitment model and the more clearly they exhibit the structure of distributed cognition. Karin Knorr Cetina, in her (1999) *Epistemic Cultures: How the sciences make knowledge* described the High Energy Physics (HEP) experiments at CERN (the predecessor experiments to those carried out with the Large Hadron Collider). Knorr Cetina points out that the size, complexity, and long duration of the HEP experiments means that expertise cannot be centralised—the research leaders, she says, are not at the top but in the middle. It is neither necessary nor possible for participants in a HEP experiment to share all that they know with other participants, let alone to come together to agree on a common view.[6]

[6]Gilbert (2000: 37–49) applies the commitment model scientific groups in particular; she argues that scientists are committed as groups to certain sets of beliefs. I think that this is quite possibly true for collaborative research teams, but I do not think that larger groups in science show this kind of collective

We may think of scientific knowledge as being produced and also possessed by groups larger even than the teams at CERN. For there is a social, non-personal sense of 'knows' that is employed in saying that it is known that copper expands when heated (to use Tuomela's example) or that anthropologists know *Homo sapiens* originated in Africa. In such cases the group in question is not the small, organised research team, but a much larger collection of people, whose boundary is not clear and which is not organised (in some cases the 'we' might seem to include everyone). This is what I shall call *wider* science, in contrast to the more local science of a research group. Participants in wider science clearly do not think of themselves as a group of which every member has expressed a willingness to let some p stand as the view of the group and have not openly expressed a joint commitment to accept p conditional on a like open expression of commitment by all other members.

The fact that the wider enterprise of science does not fit the commitment model does not necessarily refute the model's claim to be a general model of group knowledge. For proponents of that view might argue that there is not any group knowledge beyond the defined research team in wider science. While Tuomela for one does seem to want to include wider science within the remit of group knowledge, it might be that this is a mistake. And so in a later section I shall argue that we ought to regard wider science as involving a social epistemic subject.

The distributed model and science

Prima facie, the wider enterprise of science seems to fit the distributed model rather better. Science does have the key feature of distributed cognition, division of cognitive labour. The HEP experiments exhibit this very clearly; the participating scientists have their own areas of expertise and communicate with others on the basis of what needs to be shared for the local coordination of their parts of the project, what Knorr Cetina calls 'management by content'. While she says that this is something *like* distributed cognition, Ron Giere (2002a) argues, this *is* distributed cognition. Giere points to the similarities between the HEP case and the *USS Palau* (while recognising that there are also important structural and cultural differences between the two cases).[7]

Scientists build on the work of other scientists. That includes using both the results of their experiments and their theoretical conclusions. In modern science a scientist is almost never able to reproduce all the experimental results upon which her work depends for its epistemic justification. Trust in the scientific work of oth-

commitment. It is true that shared beliefs are important to a scientific community, but the importance arises less from commitment than from the fact that shared background beliefs (among other things) are a prerequisite for cooperation and mutual understanding, as explained by Kuhn. Those who reject beliefs core to a paradigm will not have reneged on a commitment; but they may find themselves excluded from the community because the basis for effective cooperation no longer exists. Gilbert proposes that her shared commitment explains why old hands are reluctant to buck the consensus: neophytes and outsiders are helpful to scientific change because they are not bound by their joint commitment to that consensus. There are other explanations. Kuhn (1962: 62–5) argues that training with exemplars builds certain cognitive habits that will make alternative views invisible to those entrenched within a scientific community. Furthermore, established scientists may have more invested in a certain way of doing things; threats to the consensus are threats to their expertise and to the basis of their prestige.

[7] Giere says, rightly in my view, that Knorr Cetina's description of an intended contrast to the HEP experiments, a molecular biology laboratory with centralised management, is also a case of distributed cognition. Its structure has some of the management structure of the *Palau* case. In his (2002b) Giere argues persuasively for the conclusion of this section, that science should be seen an instance of distributed cognition. Brown (2009) usefully supplements the case, in particular relating it to approaches in the sociology of science and learning theory.

ers is a ubiquitous and nowadays inevitable feature of science (Hardwig 1991). This goes also for theoretical results and well as for the outcomes of experiments. For although it is in principle easier for a scientists to follow and so internalise the reasoning of another scientist, in practice this too is often a matter of trust. Often those who use a scientific result are not themselves working in the field that produced that result. Palaeontologists may date a fossil from the age of the rocks in which it lies. That dating will be supplied by radiometry which depends on theories of radioactive decay. The typical palaeontologist will understand the basic ideas of that theory, but will have to accept from the physicists that the detailed reasoning involved is cogent and well-supported.

Science also displays another feature of distributed cognition, the role played by non-human vehicles of cognition. Hutchins focusses on physical means of representing information—alidades and sextants, cockpit dials and speed cards, and so forth. These are important mediators between components of the distributed cognitive system. These are present in science also—printed and now online journal articles, reference resources, and datasets are the obvious example for wider science, for these are the principal means by which scientific information is communicated between scientists; we may include informal means of communication, letters, emails, and blogs, as well as educational resources, television broadcasts, podcasts, and the like. We ought also to include the non-human means of generating and representing information, which may include experimental equipment, satellites, computers running data analysis software, and even robot scientists. Giere (2002b: 293) argues that visual representations in science can be thought of not merely as aids to cognition but as parts of systems engaged in cognition.[8]

There are however differences between wider science and Hutchins's examples of distributed cognition. The structures of the distributed cognitive systems in those cases have been designed to perform specific cognitive functions. The structures of wider science as a whole have not been designed (although some of their components have been). For example, the system of peer-reviewed journals is one that has evolved. That system may in due course be superseded—this is an issue of current debate; if it is superseded that will most likely be determined by what scientists find useful; scientific organisations may play a role in that debate, but there is no central world organisation that could decide what will happen. In Hutchins's example of the cockpit of a commercial airliner, the system is designed to meet a specific goal—knowing at each point during the plane's descent what the correct speed and wing configuration should be. While local science may have well-defined goals around which the systems are designed, wider science is not like this.

These differences between science and the paradigm examples of distributed cognition may not immediately show that the distributed model does not apply to science. One might reply that whether a cognitive system is designed or has come about by some other process is not relevant to how it functions. But that would be too quick: how a system came into being may well determine what functions it and its components have—whether something is a bug or a feature depends on design

[8] Giere (2002b: 293–4) also argues that scientific theories and models may be too complex to be fully realised as mental models, and so the models are reconstructed as external representations when needed. A slightly different approach, though one that Giere does not endorse, is to think of the mental as extending beyond the skin, so that external representations can be part of the individual's mind, as proposed by Andy Clark (1997). I shall not pursue this debate here, since I am concerned with the social subjects of cognition rather than individuals. But it is worth noting that Giere and Clark agree that there is a distributed cognitive system of which both the biological animal and the external device are parts.

intentions. In order to think of a structure as a cognitive system, whose parts have functions we need to assign a (cognitive) goal to science. Given that science is not a designed system, can we do this?

We think that science progresses—over times science gets better in certain respects, for example its scope widens, its precision and accuracy increase, and our understanding deepens. Why do we think of such changes in science as *progressive* (as opposed to merely occurring)? The best explanation of this, I propose, is that we think of science as having a goal: science progresses when it achieves or gets closer to its goal (or goals—there could be more than one). And given what we think of as contributing to the progress of science (new discoveries, theoretical advances etc.), it is plausible that science has a *cognitive* goal. It is my view that the goal of science is the production of knowledge (see Merton 1973: 270) and that science progresses when we accumulate scientific knowledge (Bird 2007). But the more general argument that the progress of science implies a cognitive goal for science is independent of this particular account of the goal and progress. P. D. Magnus (2007) holds that we cannot find a goal for science, and so the characterisation of science as distributed cognition is problematic. He notes that the goal of knowledge does not allow for the speculative activity of exploring a hypothesis. He also argues that the aim fails to discriminate between worthwhile and pointless knowledge. Both points can be answered. Some instances of 'exploring a hypothesis' do generate knowledge—for example knowledge of the mathematical or logical relationships between the hypothesis and other hypotheses, models, etc. More importantly, exploring a hypothesis is scientifically worthwhile because it can lead to knowledge—for example by allowing us to see what evidence might refute or confirm the hypothesis. If the aim of science is knowledge, then many activities (such as designing an experiment) will count as scientific because they are conducive to that end even if they do not in themselves achieve that end. Magnus is right that the kind of knowledge that science aims at needs to be circumscribed. He himself points to Kitcher's (1993) notion of significant truth. There are problems, as Magnus says, in saying what exactly 'significant' amounts to (I would suggest that significance is relative to a scientific tradition of a Kuhnian kind). But the existence of such difficulties in articulating precisely the goal of science is not a strong reason for denying that there is such a goal, whereas the commonplace idea that there is such a thing as scientific progress implies that science does have a goal.[9]

So although wider science is not designed, it can have a goal or function, just as evolved organisms and their organs can have goals and functions (this biological analogy is one to which I shall return). Given that science does have a cognitive goal (or goals) we can think of it as a system whose components contribute to that goal. Consequently science meets the characterisation of distributed cognition given above.

4 Distributed cognition and the social epistemic subject

We talk as if we are committed to social epistemic states. We use locutions such as 'it is now known that Kepler's conjecture is true', 'we know that smoking causes cancer', 'North Korea knows how to build a nuclear weapon', and so forth. Such

[9] See also Brown (2009) for a response to Magnus.

locutions are not reducible to claims about individuals (see Gilbert 1989; Bird 2010). For example, it it unlikely that any individual North Korean knows how to build a nuclear weapon (I give further examples below). Being irreducible, it is the central thesis of this paper that such statements—statements concerning wider science—are made true by the existence of social entities in epistemic states. In this section I shall argue that the distributed model of science helps us see that (i) the subjects of such statements are indeed social entities, because they are composed in a way that gives them an appropriate unity (that collections of individuals generally do not have), and (ii) such social entities have epistemic states.

It is clear that the distributed model provides a better model for understanding wider science than does the commitment model, which is no surprise. But then the two models were developed with different aims. The commitment model aims at articulating what genuine group doxastic and epistemic states are. The distributed model was developed with the aim of understanding how systems with differentiated components can achieve cognitive goals. Since the conditions that the commitment model lays down are so stringent, it is plausible that if there is any group knowledge at all, then the commitment model describes one way it can come about. The explicit commitment of individuals to the group does as much as could possibly be done to ensure that it is a unity. The distributed model does not have this ontological focus. It is true that Hutchins does *talk* of his systems (rather than just the individuals involved) as having cognitive states—one paper is entitled 'How a Cockpit Remembers Its Speeds'. But he does not pretend to give an argument for this ontological commitment. Hence the defender of the commitment model may respond to the forgoing as I indicated above: *The commitment model provides our best account of how there can be a social epistemic subject. Insofar as science or other examples of the distributed model fail to fit the commitment model, we have no reason to suppose that they are cases involving genuine social epistemic subjects (so this challenge asserts). We should therefore see the distributed model as a model of the distributed production of knowledge, not the distributed possession of knowledge.* The challenge is made sharper by the fact that within the (designed) systems that Hutchins discusses, there is typically a node, a human individual, who does possess the crucial information that the system produces and is able to use it for the purposes of the system—the Pilot Flying in the case of the aeroplane cockpit, and the plotter in the case of the ship. So it might well appear to the social epistemologist that distributed cognition really concerns what I have elsewhere called 'individual–social' epistemology, *viz.* epistemology that is concerned with the epistemic states of individuals as they relate to their social context, rather than with the epistemic states of social epistemic subjects ('social–social' epistemology).

What considerations are relevant to judging that there is an epistemic social subject in the case of distributed cognition, as found in science? In this section I argue that the hypothesized epistemic social subject has sufficient *unity* to be considered to be a genuine entity; this is because division of labour can provide the principle of composition for some social entities. I then argue on functionalist grounds that this entity can have intentional, including epistemic states. I then turn to the objection that the knowledge in question is subjectless; that is we should accept the argument of the second subsection but not the first—this is a case of knowing without a knower.

The unity of the subject

For any entity that has parts there must be something that unites those parts. In so saying, I am assuming that composition is restricted. (If it is not, then it is trivial that the components of a putative social epistemic subject form a whole.) For most physical objects that unity is mechanical—the parts will be physically joined to one another. For social entities, the principle of composition will not be mechanical. We ought to bear in mind that there might be different ways in which individuals can form social entities and that the resulting different kinds of entity may differ in their existence conditions. A set of individuals may form a social entity of ontological kind G but not of some other kind G′. So we may wish to consider that the principle of composition for a social *epistemic* subject ought to be one that plays some part in explaining how that resulting entity can be the possessor of epistemic states.

One concern that we might have in holding that there are any social entities at all, is that the individuality and free-will of human beings is a barrier to their being unified into a supra-individual entity—have not totalitarian leaders promoting social unity sought to suppress individuality and freedom? Some biologists emphasize the unity of a hive of bees or a colony of ants to the extent of arguing that these are the true organisms, not the individual bees and ants. While these insect groups do show distributed cognition, one might think that it is the 'mechanical', genetically hard-wired nature of their co-operation that allows us to think of them as forming a single entity. But the cognitive cooperation of individuals, especially in wider science, is not like that, being subject to the will.

The fact that a social entity's putative existence may be subject to the will of individuals is not sufficient to deny that that such existence is genuine. After all, the continued existence of my body is subject to my will (as is the existence of this mug, that cake etc.). If there is sufficient strength of will directed towards the unity of the social entity, that ought to allow for this objection to be overcome. So it is perhaps no surprise that the commitment model emphasizes commitment so centrally, since that is a demonstration that the individual members all will that their social entity should exist. And this might explain also the significance for Gilbert and Tuomela of the standing to rebuke, for the existence of that standing will reinforce the commitment of the individuals to the group's existence as an epistemic subject.

Individual commitment to the group is one way that the wills of individuals can support the existence of the group. But it is not the only way. The sociologist Émile Durkheim was concerned with the question, what holds societies together. The means by which social cohesion is achieved he called *solidarity*. While the question of the sources of solidarity was for Durkheim a sociological question—how do societies hold together and function effectively rather than break down or fragment—responses to that question may also be seen as answers to the composition question for social groups. Durkheim (1893) distinguishes between *mechanical* and *organic* solidarity as means by which society can be held together. In societies characterised by mechanical solidarity there is unification through the shared beliefs (experiences, obligations etc.) of individuals. Cohesion is achieved by the individuals in the group valuing their similarities of belief and value, their kinship, and their cooperative practices. In societies characterised by organic solidarity, however, cohesion is achieved not by similarity among individuals but instead by difference. Unification is generated by mutual dependence arising from, above all, the division of labour. Where an economic system operates by division of labour that system's existence does not require commitment of the individual to the system; commitment

by the individual to perform his function within the system will suffice. The fruit-canner does not need to have any conception of the economic system of which he is a part in order for him to contribute to its existence and operation, so long as he is committed to the grocer to supply his needs and to the grower to take his produce—commitments that self-interest will generate.

So individuals can compose to form a social unity when they cohere because of the mutual interdependence that arises from the division of labour. While Durkheim thought principally in terms of economic division of labour, we need to consider whether there is an epistemic analogue—solidarity arising from cognitive division of labour—if we are to think of scientists forming epistemic social groups.[10] And indeed, that is exactly what distributed cognition gives us.[11] Hutchins's examples of a landing a plane and navigating a ship show division of labour among two or more individuals (and also non-human devices). Those individuals are trained for their particular sub-tasks, and various social, professional norms (and self-interest) ensure that they are reliably performed. Just as there is division of cognitive labour in Hutchins's examples, there is division of cognitive labour in the local science of an experiment or research project. It is true also of wider science, although in this case the division of labour has evolved without a plan being imposed upon the participants. That itself does not detract from unity, for the economic division of labour, generating Durkheim's organic solidarity is likewise an un-designed evolution. No modern science depends for its conclusions just on the intellect and the evidence of the senses of a single scientist nor even of a local team of scientists. The work of other scientists will be drawn upon in the design of equipment used, in auxiliary hypotheses, background knowledge, the statistical methods and computer software employed to analyse data. The scientist's conclusions may be worthwhile in themselves but they will frequently be used as evidence or background knowledge in some other scientist's reasoning. So wider science exhibits the division of labour that gives rise to organic solidarity. This is not restricted to fields or sub-fields, for the interdependence of science crosses such boundaries. Wider science is, for the most part, a single entity.

The epistemic subject

I have just argued that there is a principle of composition (other than commitment) that can bind individuals together to form a social entity. This can apply to groups involved in distributed cognition and to science in particular. Given that there is a social entity composed of scientists, we need further to show that such an entity can be the subject of intentional states, in particular epistemic states.

The argument that entities engaged in distributed cognition have epistemic states is a functionalist one. We can identify particular behaviours and states of the

[10] I do acknowledge that there is an element of mechanical solidarity in many social entities even when organic solidarity predominates, including communities of scientists. The latter may well see themselves as a community in virtue of shared values, goals, and interests.

[11] I argue that distributed cognition is one way or implementing organic solidarity (in epistemic groups). Is there an analogous relationship between entities satisfying commitment model and mechanical solidarity? Collective commitment to shared beliefs does look very much like a form of mechanical solidarity. There may be mechanical solidarity that does not involve commitment, so collective commitment is just one realisation of mechanical solidarity. That said, it should be recognised that mere similarity of belief does not its own constitute mechanical solidarity—it is also required that this fact be mutually recognised and engender a sense of belonging together. Although this falls short of collective commitment, it points in that direction in that both require groups members to have a sense of the group as a group and to know that certain beliefs or intentions are shared by its members.

system and its parts that are best understood as actions and the intentional states that explain those actions. In Hutchins's cases, those systems are designed to fulfil a function—controlling an aircraft's descent carefully, navigating a ship. So these systems have goals and engage in actions to achieve them. The successful execution of those actions to achieve those goals depends on the systems generating and being able to access certain crucial pieces of information, and as described by distributed cognition, the systems are designed precisely to produce that information and deliver it to those parts of the system that need it. The system's goals and actions are thus supported by an information processing function of the system. If a system has goals and performs actions to achieve them, and has an aspect or subsystem whose function is to provide information that guides those actions, then that function is a cognitive one. The states that the system is in as a result of performing that cognitive function are correspondingly cognitive/epistemic ones. For example, the *Palau* needs to enter Hampton Roads at Naval Station Norfolk (goal). So the ship changes course to 265°. Why? Because the *Palau* has taken its bearings and knows that Hampton Roads lies in a direction 5° south of due west.

Can the functionalist argument be extended from Hutchins's central examples to wider science? Wider science is not a designed system (though it does have designed components). Nonetheless, as has been discussed, the lack of a design does not preclude possession of a function. Functional explanations are common in biology without that being any concession to the intelligent design view; rather, evolutionary accounts of function are required. Above I briefly argued that science has a cognitive goal. This claim is reinforced by the Durkheimian perspective, which sees various social institutions, especially those in societies characterised by organic solidarity, as fulfilling various functions, in a manner analogous to the functions of the various organs in a living being. Societies can have collective goals and engage in collective actions (such as those coordinated by governments); there are institutions that provide the information that allows those goals to be pursued and those actions to succeed; science is one of these institutions and an increasingly important one as the goals and actions become increasingly technologically sophisticated.

To the claim that the group can be an epistemic subject, one might respond that this is unnecessarily inflationary given that we can ascribe knowledge to some or all the individuals in the group. For example, in Hutchins's examples we can identify a particular central individual who receives and acts on the information supplied by the other individuals (and objects in the system). The response to the very last point is that while in some systems it may be effective to have a single, central individual in such a role, that is certainly not required, and many distributed cognitive systems lack such an individual—air traffic control is an important example. In large-scale local science, while the Principal Investigator or Programme Director may have an overview of the functioning of the project, she will not necessarily be a cognitive central node in the system; conversely, as Knorr Cetina describes the HEP experiments, those with responsibility for scientific outcomes of sub-projects and their communication to other parts of the experiment find themselves in the middle of the management structure. Such projects may produce papers with dozens of authors, where the content of the paper cannot be said to be fully known (in an individual way) by any of them (papers with hundreds of authors are increasingly common).[12]

[12]Usually those listed are not all 'authors' in the normal sense—they are participants in the project; even so many papers are collaboratively written with different authors contributing different sections according to their expertise.

More generally, group knowledge does not reduce to the knowledge of any individuals in the group. Indeed there can be scientific knowledge without any individual knowing. Consider this example:

> Dr N. is working in mainstream science, but in a field that currently attracts only a little interest. He makes a discovery, writes it up and sends his paper to the *Journal of X-ology*, which publishes the paper after the normal peer-review process. A few years later, at time t, Dr N. has died. All the referees of the paper for the journal and its editor have also died or forgotten all about the paper. The same is true of the small handful of people who read the paper when it appeared. A few years later yet, Professor O. is engaged in research that needs to draw on results in Dr N.'s field. She carries out a search in the indexes and comes across Dr N.'s discovery in the *Journal of X-ology*. She cites Dr N.'s work in her own widely-read research and because of its importance to the new field, Dr N.'s paper is now read and cited by many more scientists.

I claim (Bird 2010) that Dr N.'s contribution to knowledge did not cease with his death. Rather his discovery is a contribution to what is known in wider science in virtue of its publication and remains known thanks to the accessibility of that publication. From the perspective of wider science what is important is not whether some individual knows or not but rather whether they can use that knowledge in their own work if so required, and that is what publication ensures. Contrast that with a parallel case where a reclusive scientist does not publish his work—that work cannot be a contribution to wider science because it is inaccessible.

This example illustrates one aspect of distributed cognition—the importance of non-human forms of representation. For the system to know something, what is in peoples' heads is not important; what is important is the availability of the information known for the various social purposes the system may have. And that information may be made available through non-human means. (Libraries and the web do genuinely contain knowledge.) Another common feature of distributed cognition is the division of cognitive labour, illustrated by the following example (Bird 2010):

> Dr X., a physicist, and Dr Y., a mathematician, collaborate on a project to demonstrate the truth of the conjecture that q. Their project can be broken down into three parts. Part one is a problem in physics, the problem of showing that p, which will be the work of Dr X. alone. Part two is a problem in pure mathematics, that of proving that if p then q, for which Dr Y. takes sole responsibility. Part three is an application of modus ponens to the results of parts one and two. They arrange for an assistant to publish a paper if and only if the assistant receives from X the demonstration that p is true and from Y the proof of $p \rightarrow q$ (the brief final part with the application of modus ponens has been pre-written). We can imagine that X and Y have no other communication with each other or with the assistant and so do not know at the time of publication that q has been proven.

These examples show two things. First, that what the social epistemic subject knows does not depend on what the individuals know. Pettit (2003: 191) and Tuomela (2004: 112) subscribe to a supervenience claim, that the knowledge (and other intentional states) of a group supervene on the relevant mental states of the

individuals.[13] The supervenience claim is false for states of social knowing brought about by distributed cognition. Note that the falsity of the supervenience claim does not show that there are mysterious group-level entities that act independently of other things. Rather, it indicates that the true supervenience basis for facts about social epistemic subjects extends beyond the mental sates of individuals: the supervenience basis includes non-human entities as well as human ones and also facts about the relationships between individuals (which need not themselves supervene on individual states). Secondly, and as a consequence of the preceding point, we cannot reduce states of social knowing to states of individuals; so positing such states cannot be regarded as *merely* inflationary, inflating ontology without bringing explanatory benefit.

Subjectless knowing?

In the two preceding subsections I have argued first that we have reason to believe that there are social entities and then that we can sometimes attribute epistemic states to them. This approach assumes that if there is to be knowing at all, there must be a subject who is the knower, whether individual or social. Some philosophers of science are reluctant to ascribe a subject to this social knowing; they will agree with the claim that there is social knowing, but deny that we have to find an entity to which we can ascribe this state. Popper (1979: 109) famously held that 'knowledge in the objective sense is knowledge without a knower: it is knowledge without a knowing subject'. Giere comments on Knorr Cetina's suggestion that in a large-scale experiment, it is the experiment that knows (Knorr Cetina 1999). Giere (2002a) prefers the idea that the knowledge produced by the experiment does not belong to any epistemic subject. The forgoing arguments suggest that we do not need to avoid identifying a knowing subject. On this view Knorr Cetina is right, when we take 'the experiment' to be the total system of the the team working on the experiment and its analysis, plus the equipment and infrastructure that enables them to reach their results. We might prefer the terms 'project' or 'program'. It does not sound at all odd to assign epistemic properties to such entities, saying, for example, 'the Manhattan project discovered how to harness the power of nuclear fission'.

A further reason for rejecting the subjectless knower view is that it is difficult to reconcile with the best argument that there is social knowing at all. The best argument is the functionalist one articulated above. That argument requires, in outline, that we regard social states as having functions and where the functions can be regarded as cognitive, the states are epistemic. But states of *what*? One cannot identify a state as having a function without there being a reasonably determinate thing that has that function. One cannot identify functions independently of the things they are functions of. Hence our best argument for social knowing requires that there is a social entity that does the knowing.

[13] Gilbert (2000) also says, 'According to the conception of collective belief I put forward, collective beliefs are a matter of how it is with the individual members of the group in question'; as show goes on to emphasise, the 'how it is' is not just a matter of the corresponding beliefs of the individuals. Other features of individuals are relevant according to her view, for example, as has been noted above, commitments by the individuals.

5 When is there a social epistemic subject, and of what is it composed?

We have established that there are grounds for holding that there are social subjects with epistemic states. The arguments given allow us to answer the following questions:

1. When is it that a social epistemic subject exists?
2. Of what is the social epistemic subject composed?

A social subject exists when there is sufficient social glue (cohesion, solidarity) to join the components together. I have argued that the glue need not be joint commitment. Division of labour is another source of solidarity. Put less abstractly, the fact that certain people have mutually interacting jobs or roles may be sufficient to bind them together in the relevant way, as for example it is on a naval ship. That division of labour, those different job-roles, will give those individual functions that will contribute to the various functions possessed by the social whole thus bound together. If the latter include cognitive functions, then the social whole is an epistemic social subject. (While it may be vague when there is enough of the right kind of cohesion and distinctness that a social entity exists, that is no objection to the claims being made. It is often vague when a physical object has enough cohesion to be regarded as an entity, e.g. a cloud, melting iceberg etc.)

In what does this social subject consist? At least the individuals performing the functions contributing to the function of the whole. That much is clear. The lesson of distributed cognition and the anti-supervenience arguments is that we must include relevant non-human entities also. Describing the function of the navigation crew on the *Palau* makes no sense without including the crew's navigational equipment and indeed the ship herself. While that answer gives us the correct supervenience basis for the epistemic state, that does not mean that these are the only components of the social epistemic subject. Jane knows that Edinburgh is the capital of Scotland but John does not; Jane's hand is not part of the supervenience basis for this knowledge. Still, that hand is part of the knowing subject (Jane) whereas John's hand is not. That is because Jane is not just an epistemic subject; we identify her has having cognitive functions part and parcel of identifying other functional states. If she knows which direction is North, she may demonstrate this by pointing with her hand. Turning now to the *Palau*, the whole of interest extends beyond just the navigation team. For the cognitive function of knowing the bearing of Hampton Roads is related to the locomotive function of steering the vessel into Hampton Roads. And there will be crew members performing that function who have no role in the cognitive function. Given the connections between the various functions, those functions must be attributed to an entity comprising the whole crew.

What are the consequences of this approach for wider science? Let us start with not-so-wide science, the science in the North Korean nuclear weapons programme. That scientific programme is functioning as part of the North Korean state, with a view to enabling that state to fulfil its goals. However, the social entity in question goes beyond just the officials of the state; it will include the citizens of North Korea also, since they are functionally connected to the weapons programme in a number of ways: their labour (even if coerced) pays for the programme; they have other obligations to the state and military (e.g. military service); they live within the boundaries of the country that the weapons programme aims to protect. So those

citizens have a functional relationship to the knowledge generated by the weapons programme that the citizens of other nations do not. This explains why it is appropriate to say 'North Korea knows how to build a nuclear weapon, but Burma does not' (or 'the North Koreans know how to build a nuclear weapon, but the Burmese do not').

Now we can think about wider science, as practiced in universities and publicly-funded research institutes across the world. Scientific collaboration, the division of cognitive labour, and the dissemination of results are international and typically accessible to anyone who has sufficient background knowledge to understand them. Yet others will be indirect consumers of the knowledge, for example, through the technology it creates. Consequently, membership of the social group of wider science may be considered to be very extensive, potentially encompassing much of the world's population. Such a conclusion is neither trivial nor absurd. For if we are all part of a global village, participating, willingly or otherwise, in a highly internationalised economy, it is not absurd to think of others, even if geographically remote from us, as part of some of same social entities.[14] Science in particular has been international from an early date, with scientists in different European countries as well as America communicating with one another regularly from the seventeenth century forward.

6 Conclusion

This paper has contributed arguments in support of three claims: (i) that there are social entities; (ii) that such entities can have epistemic states; (iii) that wider science is one such entity.

The claim (i) that there are social groups is controversial. It is perhaps natural for those wishing to defend this view to take a conservative approach. The group's existence is closely tied to the intentions of its members; their commitment to the group provides it with the unity that makes it a genuine entity. The commitment model adds detail in order to gives its groups epistemic/cognitive states as required by (ii). Doxastic states are provided by certain commitments to propositions. Epistemic states, such as knowledge will come about as a result of the mechanism the group uses to reach its shared conclusion.

However. in order to answer (i), we can see that there are other ways of binding people together into a social unity, forms of Durkheimian organic solidarity, where differences between members, such as those based on division of labour, can connect people via relations of dependence. An account of group epistemic states arises naturally from a model of group unity based on organic solidarity; for example, if organic solidarity arises from division of cognitive labour, then it will be natural to see the group as having cognitive states. Systems exhibiting distributed cognition exemplify such groups. It will be necessary to regard the group as having a cognitive goal, but doing so does not require that the goal has been intentionally chosen or that the system has been explicitly designed (as in many examples of distributed cognition); evolved entities can have goals and functions, and some groups' cognitive goals may be seen in this light.

[14] This point has an ethical dimension. For example, even if we in the West are under an obligation to help relieve poverty wherever it exists, that obligation seems all the stronger when the poor in question are those making the cheap goods we buy.

Finally, claim (iii) requires that we show that wider science can be understood as a social entity with group epistemic states. Above I argued, as have Giere (2002b) and Brown (2009), that wider science is a distributed cognitive system. When added to the arguments for (i) and (ii) this makes it a plausible conjecture that wider science is indeed an entity with epistemic states—states of 'scientific knowledge' in the impersonal sense. To add confirming evidence to the plausibility of the conjecture will require detailed work regarding the structures of science showing how they do have the coordination towards a cognitive goal that one would expect of a distributed cognitive system, even though the system is an evolved rather than designed one. And it will also need to be argued that the ties that bind not only scientists but also the downstream consumers of science are sufficient for genuine unity. These, it strikes me, are promising lines of enquiry.

References

Bird, A. 2007. What is scientific progress? *Noûs* **41**: 64–89.

Bird, A. 2010. Social knowing. *Philosophical Perspectives* **24**: 23–56.

Brown, M. 2009. Science as socially distributed cognition: Bridging philosophy and sociology of science. In K. Francois, B. Löwe, T. Müller, and B. V. Kerkhove (Eds.), *Foundations of the Formal Sciences VII*.

Clark, A. 1997. *Being There: Putting brain, body, and world together again*. Cambridge MA: MIT Press.

Durkheim, E. 1893. *De la division du travail social*. Paris: Alcan. Translated as *The Division of Labor in Society*, by W. D. Halls, New York: The Free Press, 1984.

Giere, R. N. 2002a. Distributed cognition in epistemic cultures. *Philosophy of Science* **69**: 637–44.

Giere, R. N. 2002b. Scientific cognition as distributed cognition. In P. Carruthers, S. Stich, and M. Siegal (Eds.), *The Cognitive Basis of Science*, pp. 285–99. Cambridge: Cambridge University Press.

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

Gilbert, M. 1989. *On Social Facts*. London: Routledge.

Gilbert, M. 2000. *Sociality and Responsibility: New essays in plural subject theory*. Lanham, MD: Rowman and Littlefield.

Gilbert, M. 2004. Collective epistemology. *Episteme* **1**: 95–107.

Hardwig, J. 1991. The role of trust in knowledge. *Journal of Philosophy* **88**: 693–708.

Hutchins, E. 1995a. *Cognition in the Wild*. Cambridge MA: MIT Press.

Hutchins, E. 1995b. How a cockpit remembers its speeds. *Cognitive Science* **19**: 265–88.

Kanai, R. and M. Banissy 2010. Are two heads better than one? It depends. *Scientific American*. August 2012. URL: <http://www.scientificamerican.com/article.cfm?id=are-two-heads-better-than> accessed 25 Sept 2012.

Kitcher, P. 1993. *The Advancement of Science*. New York: Oxford University Press.

Knorr Cetina, K. 1999. *Epistemic Cultures: How the sciences make knowledge*. Cambridge MA: Harvard University Press.

Kuhn, T. S. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.

Magnus, P. D. 2007. Distributed cognition and the task of science. *Social Studies of Science* **37**: 297–310.

Merton, R. K. 1973. The normative structure of science. In N. W. Storer (Ed.), *The Sociology of Science*, pp. 267–73. Chicago, IL: University of Chicago Press.

Pettit, P. 2003. Groups with minds of their own. In F. F. Schmitt (Ed.), *Socializing metaphysics*, pp. 167–93. Lanham, MD: Rowman and Littlefield.

Popper, K. 1979. *Objective Knowledge* (Revised ed.). Oxford: Clarendon Press.

Schmitt, F. 1994. The justification of group beliefs. In F. F. Schmitt (Ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*, pp. 257–87. Lanham, MD: Rowman and Littlefield.

Tuomela, R. 1992. Group beliefs. *Synthese* **91**: 285–318.

Tuomela, R. 2004. Group knowledge analyzed. *Episteme* **1**: 109–27.