

SOCIAL KNOWING

The Social Sense of ‘Scientific Knowledge’

Forthcoming in *Philosophical Perspectives*

Alexander Bird

There is a social or collective sense of ‘knowledge’, as used, for example, in the phrase ‘the growth of scientific knowledge’. In this paper I show that social knowledge does not supervene on facts about what individuals know, nor even what they believe or intend, or any combination of these or other mental (including epistemic) states. Instead I develop the idea that social knowing is an *analogue* to individual knowing, where the analogy focuses on the functional role of social and individual knowing.

1 Introduction—varieties of Social Epistemology

Traditional epistemology has typically been individualistic in two respects. First, it has focussed on attributions of knowledge to individuals. Secondly, it has ignored an individual’s social environment or has treated it as akin to an individual’s physical environment. What is known as ‘social epistemology’ broke initially with traditional epistemology as regards the second of these respects. Epistemologists began to focus on social causes, testimony (e.g. Coady 1992; Lipton 1998) especially, of an individual’s state of knowledge as being particularly important or even disanalogous to the physical cases. ‘Individual–social epistemology’ retains the focus on individual knowers while socializing their epistemic context, and is exemplified by Alvin Goldman’s (1999) *Knowledge in a Social World*.

A second kind of social epistemology diverges from traditional individualistic epistemology by considering groups or social structures themselves as legitimate possessors of epistemic states. It is common to assert the existence of a de-personalized species of knowledge, as in the examples: ‘a proof of Fermat’s last theorem is now

known' and 'North Korea knows how to build an atomic bomb'. What one may call 'social–social epistemology' emphasizes not only the social nature of the epistemic context but also the social nature of the epistemic subject. According to social–social epistemology, reference to a social epistemic subject is not merely an elliptical way of talking about certain individual epistemic subjects.¹

Social epistemology, whether concentrating on individual or social subjects, further divides according to its attitude towards truth and knowledge. Goldman (1987; 1999: 5), for example, takes what he calls a 'veritistic' approach, according to which truth—conceived of as something objective and typically world-involving—plays a central role in epistemology. This contrasts with those who reject truth as important in the study of knowledge, or who regard the term 'truth' merely as an honorific label attached to beliefs the speaker shares, or to beliefs that are widely or institutionally accepted. Aligned with the latter opinions are views according to which 'knowledge' refers, as Goldman (1999: 7) puts it, to 'institutionalized belief'. Non-veritistic social epistemology is epitomized by Steve Shapin's (1994) *A Social History of Truth*, which rejects the distinction between what is true and what is accepted as such. It is also a central feature of the Strong Programme in the Sociology of Scientific Knowledge (Barnes and Bloor 1982) and many other species of opinion within Science Studies, enshrined in the equivalence (or symmetry) principle, according to which the truth of a belief plays no distinctive explanatory role in its aetiology.²

Having mapped briefly the principal strands in social epistemology, I will now sketch the view I promote below and locate it within this map. My aim is simply to give an explication of what knowledge is in the context of social–social epistemology; I explain what it is for a social epistemic subject to have knowledge. My view is unashamedly veritistic, and is an extension of traditional analytic epistemology and its conviction that knowledge entails truth.³

¹Social–social epistemology is not always clearly distinguished from individual–social epistemology. For example, Robert Wilson (2005) argues that Maurice Halbwachs' (1992) discussions of collective memory do not distinguish between the case of a group having a memory and an individual having a memory whose nature is social in some way (e.g. because it requires a social context for its manifestation).

²See also Fuller (2002). It should be noted, however, that Shapin (1994: 4) says that he does not argue with the 'restrictive notion of truth' that distinguishes between what is true and what is locally held to be true, although it is difficult to be sure that he endorses the existence of such a distinction as opposed to acknowledging the social utility of making such a distinction.

³For arguments for veritism in individual–social epistemology see Goldman (1999). His arguments carry over to my version of social–social epistemology.

Social–social epistemology takes the social epistemic subject seriously, in addition to individual epistemic subjects, and so supplements the work of individual–social epistemologists. Goldman (1999: 5), for example, also identifies this approach, but gives it minimal discussion. Compared with individual–social epistemology (and *a fortiori* traditional individualistic epistemology) social–social epistemology is somewhat underdeveloped.⁴ Popper (1972) was one of the first to emphasize a non-individualistic notion of knowledge, but did little to expand upon the idea. But some scholars, principally Margaret Gilbert, J. Angelo Corlett, Raimo Tuomela, and Frederick Schmitt, have taken up the challenge and have articulated detailed accounts of collective belief and collective knowledge. While their work in social–social epistemology represents a radical (but not necessarily revolutionary) departure from traditional individualistic epistemology, I shall argue that treatments of social epistemic states have hitherto not been radical enough, in that they continue to make social epistemic states depend too closely on states of individuals. In particular they still hold that social epistemic states *supervene* on mental (including epistemic) states of individuals; this I think is a mistake. In the view I outline below social epistemology is liberated more fully from linkage to individualistic perspectives by being seen as strictly *analogous* to individual epistemology rather than supervening upon it. Although we cannot dispense with individuals altogether, in the sense that societies cannot exist without them, they and their mental states play a minimal role in my account of social epistemic subjects, according to which we should focus on certain social structures as performing functions analogous to the functions of individual cognitive faculties. Furthermore, we must recognize that those structures can have material as well as human components. The account draws upon functionalist approaches in sociology and the allied ‘biological’ or ‘organismic’ view of society that originates with Durkheim (1893, 1894). It is the functional parallels between social and individual knowing that form the basis for the analogy between them.

I shall be concentrating on social knowledge of scientific and technological facts, as prime examples of socially generated knowledge. These examples are instructive because even though the facts in question may be quite abstruse, we are willing to attribute knowledge of them to broad collectivities, as indicated in these examples:

(I) The pure maths group in our department hopes that within five years it

⁴In his twenty-seven page survey of social epistemology, Frederick Schmitt (1994b) devotes only three paragraphs to social–social epistemology, although his own (1994a) in the same volume is in exactly that area.

will know whether the Riemann hypothesis is true.

(II) The North Koreans know how to build an atomic bomb. (*Or* North Korea knows how to build an atomic bomb.)

(III) We now know that peptic ulcer is often caused by the bacterium *Helicobacter pylori*.

(IV) The world knows that global warming is a major threat.

(V) The growth of scientific knowledge has been exponential since the scientific revolution.

That these collectivities are broad is important, since other authors, who often refer to ‘group knowledge’, have conceptions of a group that are more specific and constrained than these. These collectivities are the subjects of the attributed knowledge. If a scientist’s discovery adds to scientific knowledge, that discovery not only adds to the knowledge of that scientist but also (potentially) adds to a pool of knowledge in which many people share in the sense that they are members of the collectivity. There is a high degree of division of labour in the production and consumption of scientific knowledge, so that the individuals included in the collective possessors of knowledge in these examples may be far removed from having individual knowledge of the propositions in question.

2 Does social knowing supervene on the mental states of individuals?

My first task is to argue that there is a distinctive place for social–social epistemology. That will be true to the extent that there is no straightforward dependence of social epistemic states on individual ones. A simple reduction of the former to the latter would answer my question ‘what is social knowing?’ while also depriving social–social epistemology of a distinctive and significant role. I shall rehearse arguments against the obvious possible reductions below. Although the reductionist programme is already widely believed to be mistaken, I shall nonetheless present arguments against reductionism for two reasons. First, I do not agree that all the arguments presented hitherto against reductionism work—in some cases better arguments are required. Secondly, consideration of my arguments against reductionism will lead us to the rather more radical (and novel) conclusion that social knowing does not even *supervene* on individual mental states. Even those social epistemologists, such as Gilbert, Corlett, Tuomela,

and Schmitt, who reject reduction betray a commitment to supervenience, by linking social states to individual mental—including epistemic—states.⁵ I shall show that two societies may be exactly alike as regards the mental states of its individuals, whereas the social–social knowledge possessed in each may differ. Correspondingly what is known in a society may be tied only very loosely to what individuals know, believe, intend, or are committed to. More importantly, what a society knows will depend in significant ways on factors other than the mental states of individuals (e.g. non-human mechanisms for storing information). Initially this may seem counter-intuitive. But it becomes explicable when, in subsequent sections, we see that social knowing is an *analogue* to individual knowing and so is parallel to it rather than dependent upon it. (I should note that while some authors, such as Corlett, do give social knowledge special consideration, others, such as Gilbert, concentrate primarily on belief. The arguments considered below can be construed so as to refute reductionist and supervenience claims both as regards knowledge and as regards belief.)

2.1 The failure of summative reduction

Anthony Quinton (1976: 19) says the following:

We do, of course, speak freely of the mental properties and acts of a group in the same way we do of individual people. Groups are said to have beliefs, emotions and attitudes and to take decisions and make promises. But these ways of speaking are plainly metaphorical. To ascribe mental predicates to a group is always an indirect way of ascribing such predicates to its members. With such mental states as beliefs and attitudes the ascriptions are of what I have called the summative kind. To say that the industrial working class is determined to resist anti-trade-union laws is to say that all or most industrial workers are so minded.

Let us apply Quinton’s claim to the predicate ‘knows’ thus:

(K1) It is socially known that p iff everyone individually knows that p .

Universal individual knowing, which is also called ‘mutual knowledge’ (Vanderschraaf and Sillari 2005), is both necessary and sufficient for social knowing according to (K1). But clearly mutual knowledge is not necessary for social knowing. The organization of

⁵I include epistemic states, such as knowing, among the mental states of an individual (Williamson 1995). Those who find this disputable should read ‘mental’ in what follows as ‘mental or epistemic’.

science is characterized by the division of intellectual labour on a massive scale. Not everyone reads *Nature* from cover to cover. So many of the discoveries it publishes will not be known by everyone in science, indeed one expects that pretty well all recent discoveries will not be known by everyone. If (K1 \rightarrow) (i.e. the left-to-right implication of (K1)) were correct then none of these discoveries would be contributions to scientific knowledge (in the usual social sense), but that is just what we think (or hope) they are. Very similar considerations also refute the left-to-right implication of:

(K2) It is socially known that p iff most people individually know that p ,

since that would still restrict what counts as scientific knowledge to a very few widely-known facts. Although some major advances get published in *Nature*, most contributions to scientific knowledge appear in more specialized journals that will only rarely be consulted beyond the limited specialism. But contributions to such journals can be contributions to scientific knowledge in the general sense.

In the brief discussion above, nothing peculiar to the nature of knowledge was at work, and so we may take these example to refute the left-to-right implications of the the analogous claim about social *belief*:

(B1/2) It is socially believed that p iff all/most people individually believe that p .

So it seems that being known by all or many people is not necessary for social knowing, and likewise being believed by all or many is not necessary for social belief. Nevertheless, is widespread or even universal knowledge (or belief) *sufficient* for social knowledge (or belief)? I shall argue shortly that it is not. First I shall consider arguments of other philosophers to the same conclusion.

Corlett (1996: 88) argues that if every member of the Zuni tribe believes ‘the world exists’ for reasons entirely unconnected with their tribal membership it would not be a belief held by the Zuni tribe. He thinks this shows that common knowledge does not suffice for social knowing. I am inclined to disagree. While he is correct that ‘this hardly counts as a distinctively Zuni belief’, that is no reason for withholding the ascription. The same belief is hardly distinctive of the individuals in question either, but we do not withhold ascription of the belief to them on that account.

Schmitt (1994a: 261) gives a similar example, according to which the Audubon Society does not believe that the sky is blue, while all its members do, even if the latter fact is also known to those members. Schmitt holds the view according to which an active and collective commitment to a proposition is required for social belief, and that

commitment is lacking in the case of ‘the sky is blue’. Gilbert (2004: 97–8) discusses a union, a court, and a poetry-reading group; in these, she claims, there are processes or mechanisms (formal and informal) by which the group belief is formed. These are not merely summative. The individual members of the court may have strong personal opinions, but the court itself has no opinion until the case is concluded, discussed, and voted upon. Hence the sum of individual concurring beliefs is not sufficient for the group belief. (Nor is it necessary, since the mechanisms of group belief formation need not reflect the individuals’ beliefs—one does not have to vote as one believes.⁶) While Gilbert’s argument concerns belief and so (B1/2 \leftarrow), it clearly carries over to (K1 \leftarrow) and (K2 \leftarrow).⁷

Gilbert (2004: 96) makes a useful distinction between sets of people that constitute *established groups* and those who do not. Her argument against (B1/2 \leftarrow) applies at most to established groups that have some mechanism for creating a group belief out of the actions and commitments (such as voting) of its members. Gilbert (1989: 306) and Schmitt (1994a: 262) favour the following account of group belief:

(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 .

Correspondingly, what makes S part of a collective to which knowledge or belief is ascribed is the fact that S has contributed to the collective belief-forming process. This collective acceptance view of group belief, shared by Tuomela (2004), means that mutual belief is neither necessary for group belief (because the commitment of individuals does not require them to believe) nor sufficient (because mutual belief does not involve collective commitment).

Assuming that there are established groups that may have special conditions for group belief, then (G), if a correct account of those conditions, would indeed have the consequence that there are groups which can lack beliefs shared by all members. However, is (G) an adequate account of group belief even in established groups? Note

⁶Wray (2001) argues that Gilbert’s is an account not of collective belief but of collective *acceptance*. Wray’s view is very plausible. If there is a social state of acceptance parallel to individual acceptance, what Gilbert describes seems to be it. If Wray is correct, then Gilbert’s arguments do not really address collective belief at all.

⁷The ‘ \leftarrow ’ signifies the right-to-left implication of the biconditional in the corresponding principle.

that (G) has the consequence that an established group may not be ascribed belief even in an obvious logical consequence of the propositions it does believe, since it may not have gone through the process of getting its members collectively to endorse the logical consequence. But that should not preclude group belief, since the logical consequence may play just as much a role in guiding the collective actions of the group as the propositions it has endorsed formally. The same is true of other propositions, especially those that are common knowledge. Imagine that some committee reaches a resolution in a manner conforming to (G). For example, the committee agrees that it believes that p (\equiv the department should hire a logician). The discussion leading to that decision depended essentially on certain propositions that were common knowledge in the group (e.g. $q \equiv$ philosophy students need to be taught logic and $r \equiv$ the department's only logician is about to retire). It would be absurd to say that the committee believed p but deny that it believed q or r , simply because it had failed actively and collectively to endorse the latter propositions. It didn't need to endorse those propositions precisely because they are common knowledge. Moreover, without ascribing group belief in q or r we cannot give a social-social explanation of the decision it came to: it would be natural to say that the committee decided on p because it believed that q and r . But that group-level explanation is ruled out by (G).

So the grounds for rejecting (B1/2 \leftarrow) based on the existence of established groups may need to be bolstered by an improved replacement for (G) that permits established groups to share in what is common (individual) knowledge. Let us put this concern aside, assuming that (G) may be patched up accordingly. A more significant worry is that such arguments, appropriate as they may be to the target presented by Gilbert and Schmitt, *viz.* established groups, do not tell us whether widespread or even universal individual knowing suffices for social knowing of the kind referred to in (I)–(V) above. For in these non-established groups there is no mechanism of group belief formation in which all members may participate. Social knowing of the kind implied by the phrases, 'the world knows', or 'the North Koreans know', or 'scientific knowledge' concerns groups or societies to which neither (G) nor anything like it applies. (I) may seem to present a possible exception here, where the collective in question is a groups of researchers in the same institution. But as we shall see, while there are mechanisms of social belief formation in science, they are not of the established kind, and a group may have social beliefs that only a very few members of the group played any role in forming.

Another argument is required then to see that (K1 \leftarrow), (K2 \leftarrow), and (B1/2 \leftarrow) are false for non-established groups. Imagine that everyone individually knows that p but each person thinks that they are the only one to know it and desires to keep this knowledge secret. So if anyone were to ask whether p , all would claim ignorance. It is not at all obvious that it is true to say ‘in this society it is known that p .’ We may imagine Hans Christian Andersen’s story of the Emperor’s New Clothes as one in which everyone individually knows (because they can see it) that the emperor has a birthmark on his shoulder. However, no-one would admit to this (instead, they all claim to be able to see the wonderful clothes). Although the existence of the birthmark is a fact known by everyone individually, it is not social knowledge. It is certainly not common knowledge, which requires that individuals not only know that p but also know that everyone else knows that p , and that everyone knows that everyone knows that p , and so forth (Lewis 1969). Common knowledge, thus analyzed, and social knowing are different, although common knowledge is sufficient for social knowing. Common knowledge plays a useful social function, because everyone can rely on other people having the same knowledge, which in turn means that it can guide group action without further discussion. But in the Emperor’s New Clothes story we have only mutual knowledge with the addition of a reason for each individual not to reveal his or her knowledge. In such a case the widespread mutual knowledge plays no social role at all—from the social perspective everyone might as well be ignorant. Social knowledge, even if it falls short of Lewis’s common knowledge, is something other than just mutual knowledge.⁸

⁸Gilbert (1989: 257–8) makes a related point with regard to belief—it would be incorrect to say that the Zuni tribe believes that p if every member of the tribe believe that p but is too afraid to share this belief, thinking that he is the only one with this strange view. Gilbert also rejects the sufficiency of widespread belief for collective belief on the ground that it conflicts with the *members’ knowledge principle*, that the group G believes that p only if the members of G know that the group believes that p . This principle reflects the fact that Gilbert’s model of a social group is a highly organized group, one in which the members all contribute to the belief-forming process (see below). The principle is clearly not true for social scientific knowledge, in which contributions to scientific knowledge can be made such that only a few people know that they have been made. Furthermore, it is not obviously true even for Gilbert’s established groups. A committee may form decisions and beliefs on the basis of a quorum that is less than 100 percent. Those absent will be ignorant of the decisions and beliefs of the committee of which they are a part. Gilbert (1989: 311–12) considers explicitly whether such a committee meets the demands of her joint acceptance account of group belief and argues that it does. But she does not notice the inconsistency with the members’ knowledge principle. The members’ knowledge principle has a general problems since it makes group belief a *luminous* condition (Williamson 2000: 95) and may be refuted by considering borderline cases where the conditions for group belief have been met (just) but not all members of the group are in a position to know that they have been met.

What lessons may be drawn from the failure of mutual knowledge (or belief) to suffice for social knowledge (or belief)? One possible explanation of why mutual knowledge does not entail social knowledge is that the former does not involve group commitment whereas the latter does. Above we concluded that this explanation is problematic in its own right (it does not allow for groups to share in common knowledge), and furthermore it not allow us to account for cases (I)–(V) where the collectives may be very large and clearly do not include people who have any role in the belief-forming process.

A better explanation of the failure of mutual knowledge to suffice for social knowledge is that the latter involves socially organized access to that knowledge. In the case of the Emperor’s New Clothes there is no social access to the knowledge that individuals have—a newcomer could not come to know what everyone knows individually. As a consequence that knowledge has no social effects—everyone acts as if they do not have the knowledge in question. So mutual knowledge, in this case, is sociologically speaking, an idle wheel.⁹ In the scientific case there are mechanisms (e.g. libraries) for the dissemination of knowledge, so that in principle anyone can use that knowledge. It is because modern science is open to the rest of society that the groups to which collective scientific knowledge can be ascribed are so large; the ‘we’ in (III) encompasses not only gastroenterologists may be thought of covering most people across the globe. So the explanation for the failures of (K1 \leftarrow), (K2 \leftarrow), and (B1/2 \leftarrow) is not to be found in the requirement of individual commitment or involvement in the belief-forming process. Rather it is the fact that widespread individual knowledge does not guarantee a social mechanism for the dissemination—and hence social use—of that knowledge.

2.2 The failure of supervenience

(K1) and (K2) proposed that social knowing is a matter of being known by all or many; let us now consider a claim at the other end of the spectrum:

(K3) It is socially known that p iff some person individually knows that p .

Since being known by all does not suffice for social knowing, *a fortiori* being known by some is not sufficient either. Discoveries kept secret by individuals or small groups of individuals and not used in a wider social context are not socially known by the larger social groups to which they belong.

⁹So long as it remains merely mutual knowledge. In the story, the mutual knowledge is transformed into social (and common) knowledge by the child who shouts out that the emperor has no clothes.

Just as (K3 \leftarrow) is obviously false, it may seem that (K3 \rightarrow) is obviously true. If it is socially known that p (e.g. it is part of scientific knowledge that p) then surely someone must individually know that p .¹⁰ This view is expressed by Tuomela (2004) ‘... a group cannot know unless at least some of its members know the item in question. The general ground for this assumption is that group properties supervene on their members’ relevant properties.’¹¹ In addition to asserting (K3 \rightarrow), Tuomela commits himself to the seemingly obvious:

(SUP) Social knowing supervenes on the mental states of individuals.¹²

Although my concern is with veritistic approaches to social knowing, it may be noted that commitment to (SUP) is very wide. It includes, for example, Martin Kusch’s (2002) communitarian epistemology. As the quotation from Tuomela makes clear, this supervenience claim about knowledge is a special case of a more general supervenience claim, that the social supervenes on the individual, endorsed, for example, by Philip Pettit (2003). In what follows we shall see that (SUP) is false.

Consider a process whereby some discovery becomes a contribution to scientific knowledge in the social sense. A scientist carries out some research which produces an important, novel result. The method employed is entirely reliable. Let it be then that the scientist has produced some individual knowledge for herself. She sends her work to a leading journal in her specialist field, the *Journal of X-ology*. This journal, although read by a relatively small number of experts, is taken by every library of a serious university or research institute and its articles are indexed and abstracted by the leading indexing and abstracting services in that area of science. The journal’s editor sends the submitted article to referees, who seeing the quality and significance of the work recommend publication. They too have individual knowledge, as may also the editor. The paper is published. A small number of experts read the paper straight away, gaining more individual knowledge, and in due course the paper is referred to and cited in other papers and there is further increase in individual knowledge. At some point in

¹⁰Gilbert (1989) calls this view (as applied to belief) ‘correlativism’.

¹¹It is unclear whether Tuomela is referring to group knowledge as restricted to established groups and the like, or generally to any kind of group knowledge. (G) above clearly entails (SUP) as applied to established groups; correspondingly I do not take the failure of (SUP) as applied to the case of non-established social knowing to refute (G) regarded as a claim about established groups.

¹²In the light of footnote 5, where I take mental states to include epistemic states such as knowledge, (SUP) should not be understood as saying that social epistemic states supervene on the *intrinsic* states of individuals. Clearly social knowing depends on what is true, but so then does individual knowing.

this process of growing individual knowledge the discovery becomes socially known—it adds to the stock of scientific knowledge. What point is that? Could it be (a) when the scientist first makes the discovery? (b) when the discovery is published? (c) when a sufficiently large number of scientists have individual knowledge of the discovery? Answer (a) is not right: if the scientist were a hermit scientist, refusing to reveal the discovery, or had died very shortly after making the discovery, that discovery could have been made in the same way but without being socially known. The contribution to scientific knowledge in the social sense comes after the original discovery. Nor does (c) appear to be correct. There are some concepts for which the analogue to (c) would be a correct answer, such as the concept *crowd*. A central factor determining whether a collection of individuals forms a crowd is the number of individuals involved. But that is not the case for social knowing, for which weight of numbers does not seem relevant.¹³ A solid but specialized piece of scientific (social) knowledge might be individually known by only a few experts, while the Emperor’s New Clothes story shows that something may be very widely known without being socially known. The explanation of the latter was that the knowledge in question was not socially accessible and could serve no social function. That answer also provides the reply we seek to our current question. What make the difference between being individually known by some (or many) people and being socially known, is the accessibility of the knowledge in question. Thus we should take publication, being placed in the public domain, as (in this case) the point at which the knowledge becomes socially known in addition to being individually known. (It is *typically* the case that public accessibility of a piece of information is required for it to be able to perform a social function. But I later shall consider cases where this need not be the case.)

Now we are in a position to see why (SUP) is false and that being individually known is not even necessary for social knowledge.

(Case of Dr N.) 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

¹³A further piece of evidence in favour of this remark is that the collective term ‘crowd’ can be modified to reflect the fact that it lies on a continuum, e.g. ‘large crowd’, which is akin to the modification of ‘tall’ by the intensifier ‘very’. But there is no modifier for ‘scientific knowledge’ that serves the function of indicating that the case in question falls further along the continuum. An expression such as ‘widely known scientific knowledge’ is not an intensification, but rather indicates a conjunction: the proposition that is scientific knowledge is also known by many individuals.

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.

Was Dr N.'s discovery part of scientific knowledge? I argue that it was so *throughout* the period in question. There is no doubt that it was at the end and also at the beginning. By publishing in a well-known, indexed journal, Dr N. added to the corpus of scientific knowledge in the way that many hundreds of scientists do each month. Now consider the intermediate time t . As regards its status as a contribution to scientific knowledge, it seems irrelevant that Dr N. and others who had read the original paper had died or forgotten about it. What is relevant is that the discovery was in the public domain, available, through the normal channels, to anyone, such as Professor O., who needed it. The position at t provides a counterexample to (K3 \rightarrow)

Contrast this case with one akin to that of the secretive military scientists:

(Case of Dr Q.) Dr Q. carried out his research in secret. He typed up his results and locked them in a drawer. He died; only ten years later did someone, by chance, come across his typescript in an attic. The results eventually reached the hands of Professor R, who arranged for their publication and cited them in her work, as a result of which they were read and discussed widely.

In this case, it is not correct to say that Dr Q.'s work was a contribution to scientific discovery until it was published by Professor R.

Dr N.'s case shows that a discovery can be a part of scientific knowledge even though no-one has individual knowledge of it. When combined with the case of Dr Q. we see that social knowing does not supervene on individual knowing. Consider the intervening period in the two cases. The mental states of all individuals in the two cases are the same. In neither case does any individual know anything or have any mental state relevant to the discoveries. The two cases, in the middle phase, are identical with respect to individual mental states. But in the case of Dr N., his discovery was part of

(social) scientific knowledge during that period, whereas in the case of Dr Q, it was not. Hence one can have a difference in social knowing without a difference in individual knowing or any other mental states of individuals. Hence (SUP) is false.

What is clearly significant in Dr N.'s case, which makes it differ from Dr Q.'s case, is that his results were available to anyone who needed them. His paper was available to be read in any university library and it was indexed in the appropriate indexing services. As soon as Professor O. needed information on the area of Dr N.'s expertise she found his paper immediately. In Dr Q.'s case no-one had any way of accessing his results until they turned up on Professor R.'s desk accidentally. I shall consider the ramifications of this in a little more detail below. But for now it is worth pointing out that those sympathetic to taking the claim, for example, that 'libraries are repositories of knowledge', to be literally true should reject (SUP) straight away, if we add that the knowledge reposing there is social knowledge. (Consider a library containing unique copies of certain books. The library burns down and the knowledge those books contained is lost. If this event has no witness, then it is a loss of social knowledge without any individual mental state undergoing a change.)

So far I have taken (K3 \rightarrow) and (SUP) to be tensed claims. That is, they claim that something is socially known at a time only if some individual knows it at that time, and that social knowledge at a time supervenes on the mental states of individuals at that time. So the case of Dr N. is one where the proposition that is socially known at t is not individually known at t . Nonetheless, the proposition in question was known by someone at some earlier time and indeed at some later time. Can the rejection of (K3 \rightarrow) and (SUP) be strengthened, so as to deny not only the synchronic dependence of social knowing on individual knowing or other mental states of individuals, but also any diachronic dependence? For example, can we find cases where what is socially known is *never* individually known?

I believe we can indeed describe counterexamples to:

(K4 \rightarrow) If it is socially known that p at t then some person individually knows that p at some time or other.

Consider the Manhattan Project to build the first atomic bomb. This, like many other scientific-technological projects, involved division of cognitive labour. While individual team members knew how their bit of the project worked, they did not know the details of other parts. The project leaders knew how the various parts fitted together, but did not have detailed knowledge of each and every one of those parts. So the team as a whole knew how to build an atomic bomb, but no individual member of the project

had this knowledge. (Cognitive division of labour will be a recurrent theme in what follows, and we will come in due course to the example of the Large Hadron Collider project at CERN, which is a much bigger programme than the Manhattan Project and involves division of labour on a much larger scale.) So here we have an example of social know-*how* where no individual has that knowledge. Of course, a natural response is that this is not propositional knowledge-*that*. Nonetheless, if we accept the argument of Stanley and Williamson (2001) that all knowledge-how implies knowledge-that, then we may conclude that this case does involve social knowledge-that without any corresponding individual knowledge-that.

A rather different case also provides an example of social knowledge that q where no individual need know that q at that time or at any earlier time. Dr X., a physicist, and Dr Y., a mathematician, are collaborating on a project to demonstrate the truth of the interesting and important claim that q . They realise that the 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 easy, an application of modus ponens to the results of parts one and two. Accordingly they draft the outline paper with three parts: part one is blank except the conclusion, that p ; part two is also blank except the conclusion, that if p then q ; part three is complete, being the aforementioned application of modus ponens. The draft paper is left with a secretarial assistant with instructions to complete part one with the results of Dr X.'s work and to complete part two with the results of Dr Y.'s work, and then send it off to a leading journal, one that is read widely by the scientists in the field. In due course Dr X. completes his experiments and physical reasoning for the conclusion that p and sends the completed part one of the paper to the secretarial assistant. At the same time Dr Y. finalizes his proof that if p then q , and immediately sends the proof to the assistant as part two of the paper. Their work complete, the two scientists suddenly die, before they have communicated with one another about their success. The assistant mechanically completes this assigned task, without giving any thought to the content and sends the completed paper to the journal. The second part of the tale is a mirror image of the first. The journal's editor is too busy to evaluate the paper herself in any detail, but noticing the logical structure of the paper, she divides it up, sending the first part to a leading physicist and the second part to a leading mathematician. She leaves instructions for the editorial assistant (who has no knowledge of the paper's contents) that if both referees confirm that the arguments of their respective parts are sound, providing knowledge of

the conclusions of those parts, then the assistant should go ahead and publish the whole paper. That indeed occurs and the paper is published. Thus we have the following position: the published paper contains the conclusion that q , and the process leading to publication guarantees that it is indeed true that q and that any individual competent to understand the whole paper would come to know that q on reading it (as indeed might well happen after publication). Consequently, on publication the proposition that q becomes a contribution to scientific knowledge. In this case no individual knows that q at any time until the published article is first read in its entirety. Dr X. knew that p and Dr Y. knew that if p then q , but neither knew that q since neither knew what the other knew before they died. Neither their assistant nor the journal's editorial assistant understood the content of the paper. The editor knows the form of the argument, but, according to this story, does not know what the referees have reported (she leaves the matter entirely to the editorial assistant, or quits the editorship before the reports are in, or dies, etc.). So on publication the proposition that q becomes socially known, even though at no time has any individual known that proposition. Of course that proposition may become individually known once scientists read the article. But it need not. Clearly any *subsequent* individual knowledge is irrelevant to the possibility of social knowledge that q without any individual ever knowing that q .

The case just described may be somewhat contrived, but the conclusion that (social) scientific knowledge can be generated without any corresponding individual knowledge is increasingly a reality. The collection and statistical analysis of scientific data is frequently automated and where this data is also published automatically, on a reliable website for example, this data can be a contribution to scientific knowledge in which no human had any part. The publication of meteorological data is a good example of this. Automatic Weather Stations (AWS) collect data and transmit it to a computer that can be configured to publish that data on a live weather internet site, all without human intervention.¹⁴ In a paper published in *Nature*, Ross King (King et al. 2004) describes how a functioning laboratory robot can formulate hypotheses, design experiments, and interpret results to determine the function of specific genes in yeast. While

¹⁴AWS and software for automatic analysis and internet publication are commercially available. Where data integrity is especially important, e.g. measuring pressure at airfields, manual checks are made on instruments. However, the instrument makes its own internal checks and AWS software also checks data for exceeding limits or showing step changes (which might indicate an engineering fault). In addition to web publication, the data may be assimilated into numerical forecasting model, which will itself include data checks. So even much of the process of checking the data quality is automated. (Mike Molyneux, Met Office, personal communication)

in that first phase of the work, the hypothesis formulated and verified was one already known, King's robot has recently formulated twenty original hypotheses concerning gene coding for enzymes and has confirmed twelve of them (King et al. 2009). Since the purpose of King's research is to demonstrate the possibility of robotic science, the results of that science have not been published as independent studies. Nonetheless, it requires no great leap of the imagination to see such robotic science becoming sufficiently reliable and routine that it is produced, published, and even consumed with minimal human intervention.¹⁵

Once the possibility of the automation of social knowledge production is appreciated, it is easy to see the falsity of a diachronic version of (SUP):

(SUP_{dia}) Social knowing supervenes on the total history of the mental states of individuals (i.e. two possible worlds whose total histories are exactly alike with respect to the mental states of individuals will also be exactly alike at all times in the distribution of social knowledge).

In w_1 information about the weather at L, a remote uninhabited spot, is published on the internet as the product of an automated system whereby raw data is collected from satellites and weather stations and processed by computer. The published information, the proposition that p , constitutes social knowledge that p . World w_2 is like w_1 except that (a) the weather at L is somewhat different (so it is false in w_2 that p), but because of the remoteness of L, this affects no individual's mental states; (b) there is an unlikely accidental fault in the automatic data collection, processing, and publishing system, leading to the published output being the proposition that p . Since the published output in w_2 is false, it does not constitute social knowledge. Worlds w_1 and w_2 are alike with respect to the mental states of individuals, at the time of the time of the publication of the weather report and at all earlier times (and, quite conceivably, at all later times also). But they differ with respect to the social knowledge in those worlds. In one it is socially known that p but not in the other.

This section has tested the relationship between social states of knowing or belief and individual mental states. While a simple, summative account of the relationship has been rejected on all sides, those who concentrate on established groups do see a reduction of the relevant social states to individual states. On the face of it (G) is committed to such a reduction. It might be argued that what counts as 'open expression' depends on features of the world that go beyond the mental states of the individuals

¹⁵This field now has its own journal, *Automated Experimentation*.

concerned. Nonetheless, it is clear that as far as established groups are concerned, the connection between the group states and the individual states is a close one.

Established groups, however, do not exhaust the social groups there are and which may be the possessors of collective epistemic states, as exemplified in the statements of collective scientific knowledge (I)–(V). When we consider the latter, we see that social knowing in those non-established groups does not even supervene on individual knowing nor indeed on individual mental states of any kind. The examples of the failure of supervenience suggest a much looser connection between social knowing and individual mental states than in the case of established groups. In the cases considered, the accessibility of the social knowledge to individuals was a key consideration. It should not be concluded that accessibility is *always* an essential element in social knowing. Rather accessibility is a particular (and common) way of instantiating a more general condition on social knowing to be articulated in Section 4.

3 Commitment, solidarity, and kinds of social group

Corlett's (1996: 83) account of group belief is unashamedly summative: a collective C believes that p to the extent that each member of C believes that p *qua* member of C . Even those who reject the summative approach, nonetheless often make the social state dependent on some state or states of the individuals in question—they endorse (SUP). Tuomela (2004: 115) contends that 'when a group knows as a group ... the members ... are collectively committed to the content of the knowledge and to each other relative to it.' Gilbert's (G) likewise tells us that group belief exists in virtue of the commitments of individuals (but note the reservation above concerning the 'open expression' of those commitments). More generally Gilbert (1989: 417) asserts the *human intentional states requirement*: viable sociological collectivity concepts will entail that facts about human collectivities, in particular about their actions, are constituted by facts about the ideas and acts of will of human beings. Corlett, Gilbert, Schmitt,¹⁶ Pettit, and Tuomela all require the truth of (SUP).

While the failure of (SUP) requires rejecting the more general claims about the dependence of social states on individuals, that failure does not itself refute *all* the claims under consideration that concern groups knowledge or belief. For we may interpret some of those claims, such as (G), as claims specifically about established groups.¹⁷

¹⁶Schmitt (1994a: 259) endorses Gilbert's human intentional states requirement.

¹⁷On the other hand, Gilbert (1989: 293–4) does explicitly ask regarding her example of collective belief

However, insofar as these accounts concentrate on established groups, those whose social epistemic states satisfy (SUP) or come close to doing so, they ignore the existence of the looser collectivities that are also frequent possessors of states of knowledge, scientific knowledge in particular.

In order for a collection of individuals to be more than a mere collection, but to constitute some kind of group which possesses social properties in a non-trivial way, we must give an account of the cohesion of the group, what it is that binds the members of the group together. Furthermore, the account of cohesion should play a role in explaining why it is that the group possesses the non-trivial epistemic states that it does. For example, Gilbert, Tuomela, and Schmitt hold that established groups get their cohesion from a mechanism for forming joint commitment.¹⁸ It is because an established group exists in virtue of its joint commitment that for it to have a belief requires joint commitment to that belief.

The groups we are interested in, such as scientific communities, are not not established groups—they are what I call *organic* groups. Since their cohesion does not come from joint commitment, we need another account of where their cohesion comes from. This section provides such an account. Organic groups are those that are held together by *organic solidarity*. Émile Durkheim (1893, 1894) characterized two sources of social cohesion: mechanical solidarity and organic solidarity.¹⁹ The former occurs where

in a poetry group, ‘The poetry group, like a committee and a household, is a small group, whose members all know each other and frequently interact. Yet we also ascribe beliefs to large groups which are not like this, tribes, say, firms, and even nations. Can we expect an account of collective belief derived from the poetry group to cover such cases?’ She goes on to suggest that we can, ‘My discussion so far suggests that the core phenomenon is something like expressed willingness to let a certain view stand as the view of the group. There is no reason to suppose that the members of large groups cannot participate in such expressions or know that others have done so.’

¹⁸Corlett’s view is somewhat different. The idea is that the group must have sufficient cohesion and presence that it is a cause of the corresponding belief in its members (Corlett 1996: 83–4). That is typically brought about by ‘some decision-making structure or organizational process.’ But this requirement is too strong for reasons we have already explored. It would also seem to preclude social belief in a case where all the individuals concerned in the collective decision-making process antecedently had the belief in question quite independently, and with sufficient conviction and justification, that their membership of the group played no part in causing or sustaining their belief.

¹⁹Gilbert (1989: 243–54) also draws upon the work of Durkheim in explicating her account of collective belief—not, however, this aspect of Durkheim’s thought, but rather his emphasis on the significance of social facts as non-summatve entities. Durkheim’s view does not distinguish between particular non-summatve accounts (e.g. between Gilbert’s and mine). I note that Gilbert’s discussion of Durkheim emphasizes the capacity of social facts to constrain or coerce individuals, which may explain Gilbert’s claim that collective belief generates a right to rebuke group members who avow dissenting beliefs. However, I think it would

individuals are bound together by their similarity, while the latter involves bonds that arise out of difference, primarily the inter-dependence brought about by the division of labour. The key feature of the division of labour is that individuals and organizations depend on others who have different skills and capacities.

Organic solidarity operates both on the large scale and on the small scale. A collaborative research project is the exemplification of Durkheimian division of labour par excellence. A scientist does not need to carry out research in all areas of interest to her, since she is often able to access the results she needs in the pages of the relevant journals. The picture of the scientist who is able to obtain all the relevant evidence for himself is myth (Hardwig 1991). The division of science into subfields and specialties shows division of labour at work throughout science and its effect in binding different scientists together. For example, a palaeobiologist is investigating the relationship between certain extinct animals. A significant part of the relevant evidence concerns the age of the rocks in which the fossils of the animals were found, and thus depends on the work of geologists. The geologists, in dating the rocks, depend in large measure on techniques that concern the radioactivity of rock samples, and thereby depend on theories and equipment developed by physicists. While scientists may feel some degree of mechanical solidarity in virtue of all being scientists, and thus by sharing certain general values and attitudes as well as professional interests, the *intellectual* structure of science is held together by organic solidarity.

The contrast between organic and established groups is that between cohesion based on inter-dependence stemming from division of labour and cohesion based on individual commitment to the group. Tuomela and Gilbert have a strong conception of the latter: ‘Social collectives (or groups) . . . are taken to be capable of action in a member-binding sense’ (Tuomela 1992: 286); ‘collective belief is able to give individual parties the standing to rebuke a member for expressing a view contrary to the group belief’ (Gilbert 2004: 99).²⁰ By contrast, participation in organic groups need not im-

be mistaken, even from a Durkheimian perspective, to restrict attention to the constraining capacity of social facts. Some social facts—such as states of social knowing—may be empowering of individuals.

²⁰The best examples of this is the doctrine of cabinet collective responsibility holds that it is so for members of the British government, and so a minister can correctly say, ‘the Government believes that the war in Iraq was justified’ even if many individual members do not believe that. Even so they are each of them individually committed to not disavowing the war in public. It is arguable whether there are many other convincing cases where group belief gives a standing to rebuke. Gilbert does mention courts as examples of groups that have beliefs. But even courts constituted by a panel of judges do not possess this degree of joint commitment. A majority verdict may make it true that the court opines that the defendant committed a breach of trust. Yet the dissenting minority may be entitled, as in many British, American, and other courts,

ply that anyone has the standing to rebuke those who take a different view, merely on that ground. In particular, in the case of scientific social belief, there need not be a joint commitment to a proposition, nor any basis for criticizing those who publicly abjure those beliefs. In normal science, by definition, there is common or widespread, perhaps universal, commitment to certain beliefs. But, as Gilbert emphasizes, this is not enough for joint commitment. So there is no reason to suppose that the social groups underlying scientific knowledge are ones that require individual commitment to the group. Even more is it the case that group belief does not supply a standing to anyone whereby an individual may be rebuked for diverging for the group view. On the contrary, diversity of opinion is frequent and, in principle at least, welcomed. (This is not to deny that maverick scientists may not be rebuked on other grounds. There exists a Kuhnian conservatism because working with *shared* beliefs, standards, experimental techniques, etc. is a *logical* precondition of working within the same field of normal science and making progress within a normal science tradition (Bloor 1997; Bird 2003). Scientists will not be able to rely on one another's results, engage in fruitful dialogue, or collaborate on a project unless they share these beliefs, standards, and techniques. Criticisms of dissenting views will be primarily epistemological and are distinct from the rebuke that Gilbert refers to, which arises from breaking a *commitment* the individual makes to the group. It is true also that scientists *can also* form established groups, and a research team might be such (c.f. Gilbert 1994: 237). But such cases are not relevant to the issue at hand.)

The conclusion of this section is that commitment discussed by Gilbert and by Tuomela are not necessary for all kinds of social knowing. Durkheimian organic solidarity is sufficient, and extends beyond scientists to the society of which they are a part. The 'we' in 'we know that the Sun's power source is nuclear fusion' refers to society in general, not just to scientists. I can properly utter that sentence, even though I am not a scientist. As I shall explain, what makes this a proper way to speak is (typically) the fact that scientists make their findings publicly accessible in a systematic way. The division of labour within science is mirrored in the relationship between scientists and the rest of society; lay people cognitively depend on scientists to whom they have delegated certain epistemic tasks. In the next section I shall develop this Durkheimian approach to provide a framework for understanding social knowing. This will also allow us to answer the question, what is it, if it is not commitment, that makes an individual part

to publish its own opinion(s) without that being thought improper or undermining the fact that the court as whole has an opinion they do not share.

of a group or society of which knowledge is predicated.

4 Social knowing as a functional analogue to individual knowing

Social knowing is not supervenient on, let alone reducible to, individual knowing, as we saw in Section 2. Nonetheless there is clearly some connection. It is not that we have a homophonic ambiguity between terms that are unrelated in meaning. What, then, is the relationship? The relationship I propose is primarily one of structural analogy. The rough idea is this: just as individual knowing stands to the individual so social knowing stands to society (or the relevant social group).

4.1 The analogical approach to social knowing

One approach to developing the analogy would be this. Take your favourite analysis of (individual) knowledge and seek social analogues for the various components. So, to take the classical tripartite conception, knowledge as justified true belief, we would need an account of what social belief is, and what it is for a social belief to be justified. While they do not explicitly consider social knowing as analogous to individual knowing, the approach to understanding social knowledge exemplified by Tuomela and Corlett, and endorsed by Gilbert (1989: 313–14), follows the format of providing an analysis of social knowing that mirrors an analysis of individual knowing. Corlett, for example, analyzes knowledge, whether attributable to an individual or to a collective, in reliabilist terms: S knows that p when S has a justified, true belief that p (where a belief that p is justified when the belief that p is produced by a reliable cognitive process, that belief coheres with S's belief system, and S is an epistemically responsible agent). And so, when articulating the social version of this view, 'social epistemic reliabilism', Corlett (1996: 81) then asks, 'What is social belief? What is social truth? Moreover, what constitutes social epistemic justification?'

This approach has several disadvantages. First, it ties one to a particular analysis of knowledge. Without exception all attempts at analyzing knowledge have been subject to counterexamples and other flaws. Should one's preferred analysis of individual knowing suffer the same fate, then so will one's analysis of social knowing, and we will have made no advance in understanding either.²¹ It is preferable to try to understand

²¹For example, Corlett mentions Gettier's problem, but does not indicate how his account avoids it. Indeed

the analogy without such a commitment—particularly as there may be no true analysis of knowledge at all (Williamson 2000). Secondly, this approach can raise problems that are no more tractable than the one we started with. Considering the tripartite view, the notion of social belief may not be more perspicuous than that of social knowledge. Even more difficult, the notion of social justification may be hard to analyze, not least because individual justification is so contested.²² Thirdly, this approach gives us a poor understanding of the relationship between our concept of individual knowing and our concept of social knowing. If it were correct then our willingness to apply ‘knows’ to both individual and social subjects would be grounded in our tacit knowledge of the analyses of ‘knows’ in the two cases *plus* a tacit knowledge of their structural isomorphism.²³ This is very implausible. Our inability to come up with a satisfactory analysis of ‘knows’ suggests that if we have any tacit knowledge of the correct analysis, that knowledge is *very* tacit indeed, amounting to no more than a disposition to judge correctly whether a certain counterfactual case conforms to the analysis or not. It does *not* involve a structured conceptualization of the analysis (e.g. an unconscious belief of the form ‘S knows that *p* if and only if . . .’). But if there is no structured conceptualization of the analysis in either case, it is difficult to see how there could be a tacit recognition that the analysis of individual knowing is isomorphic to the analysis of social knowing.

Given the forgoing, it would be preferable to seek to understand the nature of social knowing in a way that does not presuppose any particular theory of knowledge. While a full understanding of social knowing may in due course require addressing contentious issues in epistemology concerning the nature of justification and belief, and their application to social subjects, I suggest that we can bypass those issues and make progress in understanding the social case by locating the analogy not in the analyses of individual and social knowing but in their *role* or *function*. The approach I shall take is motivated by the conclusion that sophisticated societies are organized organically and by the so-

it seems that his account must fall foul of it: presumably a social belief can meet the social epistemic justification condition (reliably produced, coherent with belief system, epistemically responsible agent) without being true—otherwise the truth condition for social knowledge would be redundant. Since a false social belief can socially justified, a social belief can be socially justified and also accidentally true (that is, its truth is unrelated to its being justified). Such a belief will satisfy the analysis but will not be knowledge for the reasons that Gettier gave.

²²Schmitt (1994a: 276–83) concludes that accounts of justification that apply to individuals cannot be exported to groups, with the exception of reliabilism. But his reasons rely too heavily on an account of group belief that I have argued above is mistaken.

²³For the relationship between analogy and isomorphism, see Holyoak and Thagard (1997).

biological methodology that this gives rise to. Durkheim's work initiated a school of sociology and a metaphor that guided it. The metaphor is the idea that organically structured societies may be likened to biological organisms. Organisms are made up of different parts, such as organs, limbs, and the cardiovascular and nervous systems. These parts and their interactions may be understood functionally, such that the stability and health of the whole is a matter of the proper functioning of the parts and their proper relations to one another (which itself is a matter of the proper functioning of the organism's regulatory systems). Thus the function of the lungs is to oxygenate the blood, the function of the heart is to pump oxygenated blood to the limbs and organs, the function of the diaphragm is to convert oxygen and nutrients into a muscular action that fills the lungs with fresh air, and so on. Correspondingly, the functionalist school in sociology regards a principal route to understanding social structures, institutions, and groups, to be an inquiry into their function, which may be conceived of in terms of their mutual interaction and their contribution to the stability and proper functioning of the larger society, institutions, or groups of which they are a part.²⁴ For example, Durkheim draws an analogy between the function of laws in a society and that of the nervous system in an organism.

I will draw upon the functionalist approach in seeking to understand social knowing as follows. The key will thus be to ask what function social knowing plays. To see how we may answer this question, I will take the organic metaphor seriously, that social states and organizations are analogous to states and parts of organisms. That is, I will articulate the view that we may best understand social knowing as an analogue to knowing in the individual human organism. That in turn requires describing the function of knowing in the individual. In the next subsection I will therefore sketch a view of the role of knowledge in an individual's cognitive economy—how states of knowing relate to the cognitive faculties and to practical and theoretical reasoning. I shall then show how we may naturally develop a social analogue. I will not argue for my view of the role of knowledge, since that would take us beyond the scope of this paper in the fields of the evolution and architecture of cognition. Nonetheless, the view does have some intuitive plausibility and that combined with the fact that it provides the basis for an understanding of social knowing will provide *prima facie* grounds for

²⁴Functionalism is not as dominant in sociology as it once was. It is true that not every difference between societies has a corresponding functional difference. But as most sociologists recognize, functional explanations may well be appropriate in particular cases. I note that Pettit (2003) also grounds a claim that there are collective intentional subjects on the fact they display the functional marks of an intentional subject.

accepting it that are sufficient for this paper. The advantage of this approach is that it is orthogonal to many of the troubling issues in epistemology—which is as it ought to be.

4.2 The functional role of knowing

In the preceding subsection I outlined the functionalist framework that will allow us to see how social knowing can be a functional analogue to individual knowing. In this section I add detail to the framework, first by describing the function of individual knowledge and the characteristics that allow it to fulfil that function. I then argue that the social knowledge provided by science possesses analogous characteristics and performs a parallel function.

Our starting point is the hypothesis is that (individual) knowing plays a particular role in the structure of one's 'mental economy' and that social knowing plays a structurally parallel role in a social analogue. An individual has various cognitive faculties. These all have the same function, to produce a certain kind of output that is the raw material for the processes of reasoning. Those processes comes in two types, practical reasoning and theoretical reasoning. Practical reasoning issues in action or intention to act. Theoretical reasoning issues in an output that can be stored in memory (one of the cognitive faculties) or used as an input in a further process of reasoning that may be practical or, once again, theoretical. Hence the faculty of reasoning (theoretical reasoning at least) is obviously another cognitive faculty. The role of knowledge is this: the function of the cognitive faculties is to produce knowledge as their outputs and correspondingly the appropriate input into reasoning is also knowledge. Other philosophers may accept the general framework but claim instead that the function of the cognitive faculties is to produce true belief as inputs into the reasoning process. Knowledge is what we call that true belief when the system operates in a reliable fashion. The main thesis of this paper will not be affected by which of these approaches one takes.

In both cases we can seek to justify the claim that the function of the cognitive faculties is to produce knowledge (or true belief) as the input for reasoning by appealing to standard accounts of (biological) function, whereby the function of some organ or faculty is ϕ if that organ or faculty confers biological fitness by doing ϕ .²⁵ Clearly, a system that provides knowledge or true belief as the input into practical reasoning will confer a higher degree of fitness than a system providing inputs that have no positive

²⁵For discussion of function see Wright (1973), Millikan (1989), Neander (1991), and Walsh (1996).

correlation with the truth. (The detail of the cognitive architecture, e.g. whether it is highly modular or includes general learning systems, is controversial. Such disputes are not relevant to the current discussion.)

In passing I note that this account of the role of knowledge dovetails nicely with other views concerning knowledge. For example, John Hyman (1999) argues that S knows that p precisely when the fact that p can be used by S as a reason for something. For example, knowing that it is raining is a matter of being able to use the fact that it is raining as a reason for my taking an umbrella when I go outside, or as a reason for thinking that this is the third wettest June since 1900. According to my view, the function of the cognitive faculties is just that, to provide a link between the subject and the relevant facts so that they may be used as the inputs (reasons) in practical and theoretical reasoning. Many philosophers have held that the aim of belief is truth or, in the case of Timothy Williamson (2000), knowledge. The metaphor of ‘aim’ is rarely spelt out. But we may interpret ‘the aim of belief’ as meaning ‘the function of belief producing (i.e. cognitive) faculties’. If so the slogan ‘the aim of belief is truth (or knowledge)’ becomes ‘the function of the cognitive faculties is to output true belief (or knowledge)’, as claimed above. This functional approach will be attractive especially to naturalized epistemologists. Ram Neta (2007) promotes an explicitly functional account of knowledge, while a similar view is articulated by Hilary Kornblith (2002).²⁶

The account of individual knowing provides enough for us to look for a social analogue. To make good on the analogy we need to look for social structures that are the analogues of individual cognitive faculties. Such structures, social cognitive structures, should have the following properties:

- (i) they have characteristic outputs that are propositional in nature (*propositionality*);
- (ii) they have characteristic mechanisms whose function is to ensure or promote the chances that the outputs in (i) are true (*truth-filtering*);
- (iii) the outputs in (i) are the inputs for (a) social actions or for (b) social cognitive structures (including the very same structure) (*function of outputs*).

Most societies do contain structures of this sort and that in advanced societies the structures of science occupy a central position among such structures. Let us take as

²⁶Bishop and Trout (2004) also promote a functional approach to epistemology, but seek to bypass questions concerning the nature of knowledge.

our example the group that comprises semiconductor physicists. I'll leave for a while the question of what *exactly* the 'structure' here is, but for the moment I will take the structure to be the organized community of these physicists.

(i) *Propositionality* The characteristic output of the community of semiconductor physicists is indeed propositional. As researchers the characteristic outcome of their work, whether experimental or theoretical or some combination thereof, is the journal article. The leading journals for these physicists are *Physical Review B*, *Journal of Physics: Condensed Matter*, *Electronics Letters*. Other characteristic outputs are also propositional, in the form of a conference presentation or poster, or a student textbook.

Not *all* the outputs of scientific activity are propositional. Scientists also develop new techniques, apparatus, and personal skills (although, if techniques and apparatus are also to be contributions to the field as a whole, they will usually be describable propositionally in order to be shared). Rather, the claim is that science does have structures and mechanisms whose characteristic outputs are propositional—which remains the case whatever else science does

(ii) *Truth-filtering* There are processes within science designed to ensure that the outputs are true. The most obvious of these is the peer-review of journal submissions. While the journal referee has several duties not directly related to the truth of a submission (e.g. confirming that the work is of interest and value), a central task of a referee is to eliminate error. He or she has to decide whether the reasoning is sound, whether the theoretical proposals are consistent with known results, whether the methodology is appropriate, and so forth. Papers will be rejected that do not fulfil these requirements. These processes filter out those papers likely to be false and thus make it more likely that those that are published are true.

It is useful to make a rough distinction between experimental results and theoretical proposals. As regards the former the truth filters will be pretty robust. The referees will check that the methodology is one that is reliable and that any background theory employed is well-grounded rather than speculative or unsupported. This means that in good journals the results given will be ones that have a pretty high probability of being true and hence are such that other scientists may rely upon them. Where the results are particularly important other scientists will attempt to reproduce them. It is notable that credit will be given for so doing. The editorial guidelines of *Physical Review B* state, 'Papers must contain new results in physics. Confirmation of previously published

results of unusual importance can be considered as new.²⁷

As regards theoretical proposals, the filters may be intentionally less robust. It may be necessary for several well-worked out proposals to be aired and to compete in order for others to be able to test and criticize them and for, eventually, evidence to mount up in favour of one over its competitors. Thus the ultimate goal of truth will be served by permitting the publication of theoretical proposals that may well turn out to be false. Even so filters are still applied: theoretical proposals will be published only if they ‘do not contradict established experimental results’ as the editorial policy of *Physical Review B* states. Proposals will have to meet minimum standards of theoretical plausibility and preferably have some empirical support. Thus proposals that have a low likelihood of truth are excluded.

Not all or even most that gets published in scientific journals counts as a contribution to scientific knowledge; not all that gets published is of equal value or reliability. There are further filters that typically must be passed through before something may be accorded the status of scientific knowledge. Results in a prestigious journal are more likely to be taken at face value than those in a lower ranking journal, partly because it is believed that the better journals operate a more stringent filtering process. Experimental results gained using standard techniques that are not particularly dramatic and which are published in a reputable journal may be accepted as a contribution to knowledge without further scrutiny. But if experimental results are contentious, for example because they are in tension with earlier experimental results or because they appear to refute a well-confirmed and widely accepted hypothesis, then the results will not achieve the status of knowledge until they are suitably replicated. This is why, as we have seen, leading journals are willing to publish confirmations of previous important findings. At the other end of the spectrum, a novel hypothesis will rarely become a contribution to knowledge on publication, but will typically require widespread independent confirmation and probably further theoretical discussion and debate before it achieves the status of knowledge.²⁸

²⁷<http://prb.aps.org/info/polproc.html>; accessed June 2007.

²⁸Do these remarks stand in tension with the case of Dr N. discussed above, which assumes that mere publication can amount to a contribution to knowledge? Such cases are rare, since the filtering and assessment processes typically do not stop there; but they do exist. Experimental results, achieved with a non-contentious methodology, that are themselves not too strongly contentious, and which are carried out by reputable researchers and published in leading journals may well not need replication before becoming part of the body of knowledge. Another case may be formal results in theoretical research. Mathematical results showing that a certain theory has a given mathematical consequence may stand without further work

Truth-filters of the kind already discussed recur throughout science, from the criteria used to decide whether a student gets her PhD to the assessment of proposals for research funding. The student's training will include the acquisition of standard techniques that have a proven track-record of yielding results that are likely to be true. More generally young scientists are taught a respect for the truth while many scientists are driven by curiosity, which is the desire to know and to understand.²⁹ And one cannot know what is false and one cannot understand something if one's beliefs about it are erroneous. Curiosity involves inherently a preference for truth over falsity. Secondly, given the prevalence of truth filters in science, it is instrumentally good for a young scientist to acquire a respect for the truth (cf. Papineau 1988).

(iii) *Function of outputs* The top-ranking outputs of science are employed as inputs into (a) 'social actions' and (b) systems of the kind that we have been discussing. The research results of semiconductor physics will often have direct applications in technology. For example, a discovery may lead directly to the design of a new kind of computer chip, the construction of a prototype, testing the latter, and eventually mass manufacture of the chip and marketing it. Such actions are themselves social in nature, since every stage involves the collaboration of teams of individuals. In many cases the interface between science and the production of a new piece of physical technology will involve large-scale division of labour. An extreme example of this is the design and construction of the new Large Hadron Collider at CERN. This involves enormous scientific and technical resources to overcome entirely novel challenges in three areas: the physics of the particles beams, the magnets controlling the beams, and the cryogenics needed to supercool the magnets. The project is managed by six departments (two of which are themselves purely administrative) and five scientific committees. The departments and committees oversee ten sub-projects and twenty-two working groups. Additionally the project depends on collaboration with non-European agencies as well as non-CERN projects in member states. It is the working groups that have to deal with the nitty-gritty of getting their sub-part of the project to work whereas the supervising committees have the responsibility of coordinating this work. Clearly no one individ-

since there may be no question concerning their correctness—which is not to say that formal work cannot be contentious, only that it is sometimes not contentious. Thus some pieces of scientific work may be suitable as the contents of the Dr N.'s research.

²⁹Williamson (2000: 31) asserts without argument that curiosity is the desire to know. Whitcomb (2010) argues at length and convincingly for this claim. My addition 'and understand' is not an additional component of curiosity but a special case of the more general desire to know—the desire to know how or why.

ual can have all the expertise required by the project. There is considerable intellectual and managerial division of labour. The working physical object will be the product of the large team at CERN and its external collaborators in the U.S., Canada, India, and elsewhere, and will depend on the scientific results generated both by the team and by its predecessors and other researchers in particle physics, cryogenics, magnetism, and so on. Scientific outputs, generated both within and outside the project, are in this case the inputs into a social activity of creating this remarkable machine.

Technology of this sort as well as of a more commercial kind is a consequence of ‘social action’ whose inputs are scientific results. This is the analogue to the use of individual knowledge in practical reasoning that leads to individual action. As individuals we may engage in reasoning whose outcome is not an action or intention to act but is further knowledge, which we may enjoy for its own sake, store in our memories, or use immediately in some further cognitive process. The same is true of the social knowledge produced by science. Pure research is valuable for its own sake, satisfying a deep desire to understand our world. But much scientific research does not answer directly the questions we have but provides evidence that may help us decide between rival hypotheses. So the role of that knowledge is as the input into a process of social reasoning which in due course may provide the answers we are looking for. Here ‘social reasoning’ means the process of debating rival hypotheses, submitting them to scrutiny, further development and investigation, a process which may take place in a variety of social fora—in the pages of journals, in conferences and workshops, as well as in discussions between individual scientists. Again we see a large scale division of labour, this time in the production of scientific knowledge rather than a physical item such as a computer chip or the Large Hadron Collider. This division of labour may be deliberately organised, as in the Human Genome Project, but the division of labour in science is usually closer to that of market economies, developing organically rather than in a managed fashion. Projects are pursued because of the intrinsic interest of their results and because those results may be useful in assisting other scientists, whose identities may not be known, by providing data useful to their projects (and professional credit is given for so doing).

Not all the structures that yield social knowing must have analogues to the cognitive capacities of human individuals; we do not think that all knowing individuals including those of other species should have the same cognitive capacities as a normal human adult. Nonetheless, it is worth remarking that there are broad analogues. The process whereby experimental data become social knowledge has a structural analogy to indi-

vidual perceptual knowledge (which is *not* to say that individual experimental knowledge is merely perceptual knowledge); discussion of theoretical hypotheses parallels individual theoretical reasoning; libraries and databases are analogous to memory.

As regards the last of these, in fully understanding a society we must understand the material entities that are employed in human interactions—art, trade in goods, mechanisms of communication and travel, and so forth. Modern society could not exist in anything like its current form if it did not possess the material means of rapid medium- and long-distance transport and communication. Given that material entities can make societies what they are, it should not be surprising that the cognitive structures of society should depend on and in important respects be constituted by material things, such as libraries and their contents. Scientific knowledge is typically stored in an accessible fashion, often in print form in libraries, but increasingly in one electronic form or another, available over the internet. Indeed, in order to play the role of output of and input into processes of social reasoning on the scale of modern science, social knowledge *must* be stored in some physical (i.e. non-mental) form. Even if early science did depend on personal communications between scientists, modern science could not function that way. The existence of material structures such as books and journals or the internet is not an accidental feature of modern science, but is essential to it. The social cognitive structures that mirror individual cognitive structures inevitably have material components. For this reason social knowing does not supervene on states of individuals alone. This does not imply that social knowing could in any sense do without individuals. It is not *merely* the material inscriptions in libraries that constitute knowledge. Just as what makes something a human faculty of memory is its role within an individual's larger cognitive psychology, so what makes a library a repository of social knowledge is its role within a larger set of institutions.

4.3 Functional integration

To conclude this section I shall show how Durkheim's functionalism and the organismic analogy can help us answer the question of the composition of the relevant social group of which social knowledge is being predicated. Gilbert's account of social belief in terms of commitment implies a parallel account of group membership also in terms of that commitment. In rejecting an approach that is at best limited to established groups, I rejected such a view of group membership. A sketch of an alternative approach is necessary, for it might be tempting to replace Gilbert's explicit commitment with a tacit commitment of some kind, paralleling the appeal to tacit consent that one

finds elsewhere in social theory. My alternative approach draws on Durkheim's organic analogy. What makes an organ or cell a component of one organism rather than another? The answer, in brief, is *functional integration*. Components are functionally integrated when there is dependence between them for their own proper functioning. Thus the vascular system supplies nutrients and oxygen to all organs, and they are thereby functionally dependent on the vascular system, which in turn is dependent on the heart to move the blood and the lungs to supply it with oxygen. So a component may be integrated with a whole by being a producer of something needed by other components, or by being a consumer of their products, or both. This kind of functional integration may be developed into account of belonging to a social group. Division of labour, for example, leads to functional connections between individuals that bind them together as members of society. One answer to the question, 'who belongs to G in 'G knows that ...'?' would regard an individual as included in G if that individual is functionally integrated into G in a general social sense. Note that this can be true without that individual being in any way a part of the process of production of that knowledge. Nor in any direct way need the individual be a consumer of that knowledge, for other kinds of integration, e.g. economic, might be sufficient for membership of the group. Thus not only may an individual not know that *p*, yet be part of some community, G, of which it is true to say 'G knows that *p*', but also that individual might even deny that *p*, for example if *p* is the claim that humans and apes have a common ancestor, and the individual is a creationist in the U.S.. Integration into American society more generally will suffice for creationists' membership of that society, and the fact that the relevant scientific institutions are also integrated into that society, means that the former are included in the ambit of the social knowledge created by the latter.³⁰

Thus individuals who are ignorant of, or even reject, certain items of scientific knowledge, may nonetheless be included in a wider collective of which that knowledge can be predicated in the social sense. That inclusion is based upon their integration in primarily non-cognitive (e.g. economic and social) ways. But they will also be integrated in cognitive ways. The various institutions and social systems are multiply connected. And so an individual may be an indirect consumer of knowledge, for example by being the user of products designed using that knowledge: a creationist might be a consumer of biological theorising they reject by using pharmaceuticals developed

³⁰The position is analogous to the fact that when it is true to say that 'the British were at war with Nazi Germany in 1940', the term, 'the British' encompasses British pacifists and even Nazi sympathisers such as Oswald and Diana Mosley.

using that knowledge. Individuals will also be users of an education system which draws upon or is connected to such knowledge. That knowledge will typically, in a loose sense, be accessible to these individuals, through the education system, through public libraries, and through the internet. For much scientific knowledge, this accessibility is very much ‘in principle’ in the sense that those individuals may not have the background knowledge required to understand the information in question, and it may in fact be difficult for a lay person to get access to journals typically found only in university libraries. So the integration of an individual into the cognitive part of the total social system will not be a matter of actual or even potential direct consumption of its knowledge. That species of integration will usually be indirect. As long as those who may need to use the information have access to it, and an individual is cognitively linked to those persons (perhaps indirectly), then that individual is linked into the cognitive structures and institutions of the society of which she is a part. While it is sufficient for individual S to be part of group G when considering attributions of knowledge to G, that S is integrated into G in non-cognitive (e.g. economic and political) ways, it is also often and increasingly the case that S is integrated into G cognitively, even when S is ignorant of, or rejects, the knowledge in question.

Section 2 emphasized *accessibility* as a key consideration in attributing social knowing. But the fundamental point is functional integration—the knowledge plays a social role. While knowledge typically plays that role via socially organized access, this need not mean ‘directly accessible to everyone’. Most scientific knowledge is not directly accessible to most people. Lay people would find it hard to get hold of most scientific journals, and their lack of scientific training would mean that they could not understand what they read. Functional integration does not require universal access, even in principle. What is important is that the information should be accessible to those who need it. By publishing in *Physical Review B*, a researcher in semiconductor physics makes her research accessible to other semiconductor physicists, and to those in related areas of science or industry who might find her work useful. Thus the relationship between a layperson and an item of scientific knowledge may be very indirect: the structure of science and technology make it possible, for example, that this research will feed into other research that will result in some piece of technology (such as computer chip or component of a hospital scanner) of which the layperson may be a beneficiary. Such indirect connections are what one would expect from functional integration, just as the visual cortex and the liver are functionally integrated into one person, although only indirectly connected. Hence it is not the accessibility of the knowledge that is essential

to its being social knowledge; rather it is the capacity of the knowledge to play a social role (e.g. in decision making by the group) in virtue of the structure and organization of the group; accessibility is the principal means by which that is achieved.

And so we should not expect direct access even ‘in principle’ in every case of social knowing. It is true to say that in 1945 that the Americans, British, and Canadians knew how to build an atomic bomb but the Germans, Soviets, and Japanese did not. Yet that knowledge was secret and far from accessible, even to those who would have understood it. In this case the knowledge is functionally integrated, but not via the normal structures of science and commercial technology but via political and military structures.

Accessibility, when relevant, is vague. Some journals have a higher profile than others; they may be abstracted more widely and more freely available on the internet. Information placed on an obscure website is in a sense accessible, but not in the way that information published in *Nature* is. Functional integration may come in degrees. Consequently social knowing is vague (it is vague in other respects too). Such vagueness is to be expected—individual knowing is similarly vague. Memories fade, become corrupted and increasingly difficult to retrieve; their status as knowledge is vague.

5 Conclusion

The problem addressed by this paper is that of explaining what social knowing is. The question is not an idle one. Knowledge matters. Individual knowing is important in the justification of individual action; and arguably it is a constitutive norm of assertion (Williamson 1996). Reasons of this sort transfer to the social analogue. If it is an abuse of assertion for an individual to assert something he does not know to be true, it is an abuse for institutions to make assertions that are not socially known to be true by that institution. Consider the assertions of certain governments regarding the existence of weapons of mass destruction in Iraq in 2003: it is not simply that individuals made assertions without the requisite individual knowledge, but that such assertions were also the assertions of institutions, seeking to justify social actions (such as going to war), without social knowledge of what was asserted. A 2009 debate about the classification of cannabis in Britain raised the question whether the then government was properly committed to evidence-based policy making. The latter, on one view of evidence at least, is a matter of forming policy about social actions on the basis of social knowledge. It is indeed social knowledge that we are concerned with, for the knowledge

must be produced in those socially sanctioned ways that generate social knowledge rather than just individual knowledge. It would not be satisfactory if the Home Secretary's decision were based on experiments carried out privately in his garden shed, however well conducted. That kind of purely private, individual knowing cannot play the role of justifying social decisions and actions.

In considering what social knowing is, I initially considered various reductionist theses. While the most obvious of these have already been rejected by other authors, those authors are still committed to taking social knowing to be dependent, if not on individual knowing, then on other mental states of individuals. But we saw that social knowing does not supervene even on the sum of individual mental states—worlds can differ with respect to social knowing that differ only in physical respects, not in any (individual) mental respect.

We were thus presented not only with the continued challenge of finding an explication of social knowing, but also of showing why it should fail to supervene on individual mental states. The account given here meets this challenge. Social knowing is not reducible to individual knowing, nor is it dependent on mental states in some more complicated way, precisely because it is *parallel* to individual knowing. Analogues need not be, and typically are not, reducible to one another, nor are they supervenient upon one another.

Other philosophers, in discussing group knowledge, adhere to supervenience or positions close to it because they have a specific, strong conception of what can contribute to the social; in particular they make the joint commitment of individuals the central feature social belief. While this may be appropriate for some kinds of group, established groups, this ignores the widespread division of labour and responsibility that is the source of cohesion in many social structures. Political leaders will make decisions whose consequences are social actions engaged in by whole nations, whether or not those nations' populations accept or even know of those actions. Cognitive division of labour works similarly. It is a mistake to think of cognitive division of labour as a matter of dividing cognitive tasks and then explicitly sharing the results among all participants. Instead, we devolve responsibility for cognitive activities to the relevant experts. Their activities, if of the right kind, suffice for social knowing, whether or not that knowledge is shared explicitly.

In demonstrating the failure of supervenience I gave the examples of Drs N. and Q.. The principal factor in making the social difference between the two cases was the availability of information—the presence of an article in a well-known journal in a

library versus its being in a drawer in someone's attic. That is a physical difference.³¹ But it is also a social difference. Societies are more than totalities of individuals. Infrastructure and the means of communication are key to the understanding of the functioning of societies. When it comes to the cognitive functions of institutions within societies, libraries and, increasingly, the internet are central in this regard. However, it is not simply their physical existence that gives, for example, a library and its deposits their social cognitive function. Two conditions must also be satisfied. First, there must be mechanisms for ensuring that those deposits are likely to be true. Secondly, there must be reliable means for accessing those deposits, at least by those who may need to use them, so that they can affect what individuals and groups think and do. The second ensures that the deposits are analogues to *belief*-like mental states; the first disposes them to be states of *knowing*.

In short, a functionalist account of mind and a functionalist account of social institutions together make it plausible that there could be structures in persons and in societies that perform analogous functional roles, and it seems that science is one of the social institutions that performs a social cognitive role analogous to the knowledge-generating cognitive powers of individuals. Finally, I suggest that this points in the direction of a general functionalist conception of knowledge that encompasses both individual and social knowing. A complex structure is in a state of knowing when that state is the output of a properly functioning cognitive substructure. A cognitive substructure is identified as having the function of producing true outputs that are potential inputs either for what may be identified as actions on the part of the whole structure or for further process of substructures of the same (i.e. cognitive) kind. That is it may not simply be that we use 'knows' in the social case because it bears an analogy to the individual case; rather, both cases may be seen as instances of a more general concept of knowledge, and it is for this reason that they have analogous properties.³²

³¹I note the coherence of this argument with the extended mind hypothesis concerning individuals (c.f. also Wilson (2005), who addresses this question). According to Clark and Chalmers (1998) an individual's mind can extend beyond his skin in that, for example, a person can be properly said to know someone's address if that address is in an address book to which they have immediate and habitual access. Hence the extent of individuals' mental states is broader than normally thought, even once epistemic and semantic externalism are accepted. Thus the supervenience basis under consideration in (SUP) is broader than normally thought. Might (SUP) then be true after all? No. Accepting the extended mind hypothesis, one might consider the contents of a person's address book as part of their knowledge, but not the contents of a university library (see Clark 2006). The latter is what would be needed to undermine the argument against (SUP).

³²Acknowledgments: I am grateful to Mike Molyneux (Met Office) and Darrell Rowbottom for information and commentary, to audiences at Belfast, Birkbeck, Bristol, Cambridge, Edinburgh, and the London

References

- Barnes, B. and D. Bloor 1982. Relativism, rationalism and the sociology of knowledge. In M. Hollis and S. Lukes (Eds.), *Rationality and Relativism*. Blackwell.
- Bird, A. 2003. Three conservative Kuhns. *Social Epistemology* **17**: 125–31.
- Bishop, M. and J. D. Trout 2004. *Epistemology and the Psychology of Human Judgment*. Oxford: Oxford University Press.
- Bloor, D. 1997. The conservative constructivist. *History of the Human Sciences* **10**: 123–5.
- Clark, A. 2006. Memento's revenge: The extended mind extended. In *The Extended Mind*, pp. 18–44. Aldershot: Ashgate.
- Clark, A. and D. Chalmers 1998. The extended mind. *Analysis* **58**: 7–19.
- Coady, C. A. J. 1992. *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- Corlett, J. A. 1996. *Analyzing Social Knowledge*. Lanham, MD: Rowman and Littlefield.
- Durkheim, É. 1893. *De la division du travail social (The Division of Labour in Society)*.
- Durkheim, É. 1894. *Les règles de la méthode sociologique (The Rules of Sociological Method)*.
- Fuller, S. 2002. *Social Epistemology* (2nd ed.). Bloomington: Indiana University Press.
- Gilbert, M. 1989. *On Social Facts*. London: Routledge.
- Gilbert, M. 1994. Remarks on collective belief. In F. F. Schmitt (Ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*, pp. 235–56. Lanham, MD: Rowman and Littlefield.
- Gilbert, M. 2004. Collective epistemology. *Episteme* **1**: 95–107.
- Goldman, A. I. 1987. Foundations of social epistemics. *Synthese* **73**: 109–44.
- Goldman, A. I. 1999. *Knowledge in a Social World*. Oxford: Clarendon Press.

School of Economics for helpful discussion, and to anonymous referees for advice.

- Halbwachs, M. 1992. *On Collective Memory*. Chicago, IL: University of Chicago Press.
- Hardwig, J. 1991. The role of trust in knowledge. *Journal of Philosophy* **88**: 693–708.
- Holyoak, K. J. and P. Thagard 1997. The analogical mind. *American Psychologist* **52**: 35–44.
- Hyman, J. 1999. How knowledge works. *Philosophical Quarterly* **49**: 433–51.
- King, R. D., J. Rowland, S. G. Oliver, M. Young, W. Aubrey, E. Byrne, M. Liakata, M. Markham, P. Pir, L. N. Soldatova, A. Sparkes, K. E. Whelan, and A. Clare 2009. The automation of science. *Science* **324**: 85–9.
- King, R. D., K. E. Whelan, F. M. Jones, P. G. K. Reiser, C. H. Bryant, S. H. Muggleton, D. B. Kell, and S. G. Oliver 2004. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* **427**: 247–52.
- Kornblith, H. 2002. *Knowledge and its Place in Nature*. New York: Oxford University Press.
- Kusch, M. 2002. *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: Oxford University Press.
- Lewis, D. K. 1969. *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Lipton, P. 1998. The epistemology of testimony. *Studies in History and Philosophy of Science* **29**: 1–31.
- Millikan, R. 1989. In defense of proper functions. *Philosophy of Science* **56**: 288–302.
- Neander, K. 1991. Function as selected effects: A conceptual analyst's defense. *Philosophy of Science* **58**: 168–84.
- Neta, R. 2007. How to naturalize epistemology. In *New Waves in Epistemology*, pp. 324–53. Ashgate.
- Papineau, D. 1988. Does the sociology of science discredit science? In *Relativism and Realism in Science*, pp. 37–57. Dordrecht: Kluwer.
- Pettit, P. 2003. Groups with minds of their own. In F. Schmitt (Ed.), *Socializing Metaphysics*, pp. 167–93. Lanham, MD: Rowman and Littlefield.

- Popper, K. 1972. *Objective Knowledge*. Oxford: Clarendon Press.
- Quinton, A. 1976. Social objects. *Proceedings of the Aristotelian Society* **76**: 1–27.
- Schmitt, F. F. 1994a. The justification of group beliefs. See Schmitt (1994c), pp. 257–87.
- Schmitt, F. F. 1994b. Socialising epistemology: an introduction through two sample issues. See Schmitt (1994c), pp. 1–27.
- Schmitt, F. F. (Ed.) 1994c. *Socializing Epistemology: The Social Dimensions of Knowledge*. Rowman and Littlefield.
- Shapin, S. 1994. *A Social History of Truth*. Chicago, IL: University of Chicago Press.
- Stanley, J. and T. Williamson 2001. Knowing how. *Journal of Philosophy* **98**: 411–44.
- Tuomela, R. 1992. Group beliefs. *Synthese* **91**: 285–318.
- Tuomela, R. 2004. Group knowledge analyzed. *Episteme* **1**: 109–127.
- Vanderschraaf, P. and G. Sillari 2005. Common knowledge. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2005 Edition ed.). <<http://plato.stanford.edu/archives/win2005/entries/common-knowledge/>>.
- Walsh, D. 1996. Fitness and function. *British Journal for the Philosophy of Science* **47**: 553–74.
- Whitcomb, D. 2010. Curiosity was framed. *Philosophy and Phenomenological Research*. (forthcoming).
- Williamson, T. 1995. Is knowing a state of mind? *Mind* **104**: 533–65.
- Williamson, T. 1996. Knowing and asserting. *The Philosophical Review* **105**: 489–523.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wilson, R. A. 2005. Collective memory, group minds, and the extended mind thesis. *Cognitive Processing* **6**: 227–36.
- Wray, K. B. 2001. Collective belief and acceptance. *Synthese* **129**: 319–33.
- Wright, L. 1973. Functions. *Philosophical Review* **82**: 139–68.