

Epistemic dependence and collective scientific knowledge

Jeroen de Ridder

Received: 7 July 2012 / Accepted: 11 April 2013 / Published online: 15 June 2013
© Springer Science+Business Media Dordrecht 2013

Abstract I argue that scientific knowledge is collective knowledge, in a sense to be specified and defended. I first consider some existing proposals for construing collective knowledge and argue that they are unsatisfactory, at least for scientific knowledge as we encounter it in actual scientific practice. Then I introduce an alternative conception of collective knowledge, on which knowledge is collective if there is a strong form of mutual epistemic dependence among scientists, which makes it so that satisfaction of the justification condition on knowledge ineliminably requires a collective. Next, I show how features of contemporary science support the conclusion that scientific knowledge is collective knowledge in this sense. Finally, I consider implications of my proposal and defend it against objections.

Keywords Social epistemology · Scientific knowledge · Collective knowledge · Epistemic dependence · Testimony · Justification

1 Introduction

Modern science is a thoroughly collective enterprise. Collaborative work is the norm rather than the exception. Research teams and scientific communities are characterized by a more or less official division of epistemic labor. And even without any official division of epistemic labor, scientists rely on epistemic work done by colleagues: formulating and testing hypotheses, data gathering, processing, and analyzing, and experimenting. This reliance extends to contemporary colleagues, but also to past colleagues. Typically, then, no individual scientist has the cognitive resources to oversee

J. de Ridder (✉)
Faculty of Philosophy, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam,
The Netherlands
e-mail: g.j.de.ridder@vu.nl

all the epistemically relevant aspects of the research projects that she—or: her team or community—is engaged in. Epistemic dependence among scientists is pervasive (Hardwig 1985).

In spite of wide recognition of these facts by philosophers of science, the idea that scientific knowledge is collective knowledge, i.e., knowledge possessed by a collective rather than individuals, remains controversial. Kitcher, for instance, insists that taking account of the social aspects of knowledge does not “demand a break with the traditional conception of knowledge as something that is located in (or possessed by) an individual subject” (1994, p. 118). And notwithstanding their enthusiasm for distributed cognition, Giere (2006, 2007) and Thagard (1997, 2010), too, explicitly deny that knowledge can be had by collectives.

In this paper, I argue that the epistemic dependence among scientists makes a lot of scientific knowledge collective knowledge. In Sects. 2 and 3 I look at existing proposals for understanding collective knowledge and argue that they are unsatisfactory. I take a step back in Sect. 4 and propose a different way to construe collective knowledge. In Sect. 5 I show that contemporary scientific knowledge is indeed collective knowledge in my proposed sense. In Sect. 6 I consider implications of my proposal and discuss various objections. Section 7 concludes the paper.

2 Collective knowledge: summative analyses

Although the idea of collective knowledge is controversial, there have been a few proposals which attempt to make sense of it.

First, collective knowledge can be understood in a purely summative fashion. A collective knows that p iff all (or most) of its members know that p . This basic idea would have to be supplemented with some account of what it is for individuals to compose a collective, because we don’t want to say that every random set of individuals who all happen to know that p has collective knowledge.

This proposal, however, fails to give us a robust sense of collective knowledge. On a summative analysis, talk of collective knowledge is shorthand for saying that all individuals in a collective know something. There isn’t anything distinctively collective about summative knowledge. Says Anthony Quinton:

We do, of course, speak freely of the mental properties and acts of a group in the 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. (Quinton 1976, p. 17)

One might try to remedy this by beefing up the basic summative analysis with additional conditions that do involve the collective in some sense. For instance, an additional condition for collective knowledge could be that it also requires that all (or most) individuals in the collective know or at least believe of the others that they know that p . Knowledge that p could then be common or shared knowledge. Such a condition secures that collective knowledge requires something of each individual member that involves the other members.

Nonetheless, collective knowledge so construed remains wholly reducible to the knowledge and beliefs of the individual members of the collective. On a summative understanding, what it is for a collective to have knowledge can be explicated fully as a mere sum of what individuals know and believe. The collective's knowledge is nothing over and above its members' knowledge and beliefs. Talk about collective knowledge can always be eliminated in favor of talk about individuals and their knowledge and beliefs. Summative analyses are thus unsatisfactory: They don't give us an account of robustly collective knowledge (Gilbert 1987, 1994; Bird 2010; Fagan 2011).

The upshot of this is as follows. Some of our everyday talk that involves ascriptions of knowledge to a collective certainly lends itself to a summative analysis. I will argue, however, that not all such talk can be analyzed this way. There is also a more robust conception of collective knowledge.

3 Collective knowledge: joint commitment

A second proposal for construing collective knowledge derives inspiration from the thought that collectives can adopt a view that isn't reducible to its individual members' views, at least not in any straightforward sense.¹ A committee can decide that job candidate X is the best candidate, even when no individual committee member thinks this (e.g., X is everyone's second pick). A professional organization of, say, medical doctors can issue a statement about the risks of certain behaviors, even when the doctors individually haven't made up their mind, or are divided about what the evidence really points to.

Common to such scenarios is that there is a formal or informal normatively binding decision procedure that generates a group view. Group members are expected to let the view that is formed through proper application of the procedure stand as the collective view. In order for the results of such a procedure to count as the collective's view, it is also necessary that the individual members give their informed consent to the procedure. This need not require much and may be largely implicit, but at least group members should understand that being part of the group entails a commitment to the procedure.

The exact nature of the decision procedure may vary, from a simple majority vote to a multi-layered deliberative process. It can be one time and ad hoc or formally institutionalized. A procedure can be dictatorial or oligarchic, taking the views of only one member or a few members as input. It can even be to let a formally appointed individual or group outside the collective determine the collective's view.

This second proposal leads to something that is more robustly collective. The fact that there is an agreed upon procedure makes it the case that the collective view is not reducible to the mere sum of the individual views. Although the collective's view depends on the members' individual views, it really is something over and above the individuals' beliefs and knowledge, brought into existence by the intentional actions of the individuals in the collective. A full description of the situation would have to

¹ Gilbert (1989, 1994), Tuomela (1992, 2004), Schmitt (1994b), Rolin (2008, 2010), and List and Pettit (2011) contain detailed proposals.

include the group's decision to let a view stand as the group's view and must therefore contain more than just what the individuals in the collective believe and know.

While this may yield a plausible construal of what it is for a collective to believe or accept a view, more is needed for collective knowledge. A decision procedure for generating a collective standpoint can aim at many different things besides truth: efficiency, speed, fairness, social acceptability, or other goals. Procedures aimed at such non-epistemic goals may incidentally latch onto truth-conducive considerations and generate views that are true, but their outcomes would not be knowledge. This is because knowledge is widely held to be incompatible with *luck* of a specific kind.² There are forms of luck that are compatible with knowledge (Unger 1968; Pritchard 2005). You may be lucky to exist and have the cognitive faculties you do (perhaps you narrowly escaped a tragic accident), facts you know about may be lucky (that John's ticket won the lottery), you may be lucky to have acquired a specific piece of knowledge (you just happened to be walking by, or you made a serendipitous discovery).³ Knowledge is incompatible, however, with what Pritchard (2005, p. 146ff), calls *veritic epistemic luck*: If it is lucky that your belief is true, in the sense that you could have easily obtained a false belief (while keeping your circumstances and methods of belief formation the same), you don't have knowledge (you don't *know* that it's noon if you look at a clock that stopped exactly 24 h ago).⁴ Since decision procedures that are directed towards non-epistemic goals are not responsive to evidence or other truth-conducive factors, they will easily lead to false beliefs. True group views generated by them thus suffer from knowledge-undermining luck.

Next, a decision procedure must not only *aim* at truth, it must also be sufficiently reliable in achieving this aim. Exactly how reliable is impossible to say, but it should at the very least be more likely than not to generate a true view when properly applied in its intended domain of application, and presumably significantly more likely.⁵ Only procedures that aim at truth and are sufficiently reliable can generate collective knowledge.

How does this pan out for the idea that scientific knowledge is collective knowledge? That depends on whether groups of scientists have and use decision procedures of the right kind. At first sight, it may seem that they do. They publish joint papers or research reports, write grant proposals together, or collaborate on research projects. All of these

² Consult any introduction to epistemology and one will come across a statement of this idea. According to Pritchard (forthcoming), it is one of two 'master intuitions' about knowledge.

³ Roberts (1989) gives wonderful stories about the role of luck in some scientific discoveries. Thanks to an anonymous referee for pointing me to this source.

⁴ In a sense, it is also lucky that we do not live in a demon world or other skeptical scenario. Since most of us aren't skeptics, however, we are committed that such luck is not (always) knowledge-undermining (cf. Pritchard (2005), Ch. 9).

⁵ Note that there are two different ways to think about the reliability of a decision procedure: some procedures are reliable only on the condition that their input is already of high epistemic quality. For instance, a majority vote among reliable experts may be reliable in this sense. Other procedures are reliable in the more robust sense that they not only conserve the epistemic quality of their input, but improve upon it. Even if the input is of bad or mixed epistemic quality, a rational deliberative process in which evidence is filtered and arguments are weighed may reliably produce true outputs.

activities involve a joint commitment to letting certain claims stand as the collective's view, at least for the time being and the purposes at hand (cf. [Wray 2007](#)).

One could object to these examples on the grounds that there is usually no established procedure to generate the collective's view. This is not a good objection. Informal, ad hoc, or one time procedure can also serve to generate a group view. Nonetheless, the objection points to an aspect of scientific practice that makes trouble for those who want to use this construal of collective knowledge to account for collective knowledge in science, such as [Wray \(2007\)](#), [Rolin \(2008\)](#), and [Gilbert \(2000, 2004\)](#). In general, there are no decision procedures through which scientists generate a collective view, let alone normatively binding ones. Formation of a collective view, if it happens at all, is usually more organic or emergent. And when a collective view has formed organically, it is not normatively binding.⁶ When an individual scientist is convinced she has good reasons to reject the view, she is not under any obligation to refrain from doing so, at least not solely in virtue of her being a member of a collective.

This means that the construal of collective knowledge under consideration has limited application in science. The use of procedures to form a collective view which are reliably aimed at truth and normatively binding is at best rare in scientific contexts. Research teams may sometimes have them, but certainly not scientific subfields or disciplines as a whole. In general, scientists do not have decision procedures, let alone formally established ones, to generate a collective view.⁷

Even those situations in science that seem most amenable to an analysis in terms of a normatively binding decision procedure—such as the writing of joint papers or research proposals or consensus conferences—may fail to give us collective knowledge in the sense under consideration. This is because it is unclear that the procedures used to decide what will be in a paper or proposal or what will be the consensus are in fact reliably directed at truth. Other aims besides truth can play a major subsidiary role and can sometimes even trump truth. A research proposal aims at obtaining money from a granting agency. The best way to do this is often not to state the truth as accurately as you can, but to overstate your case a bit, to present results in a more favorable light, or to omit qualifications that would be necessary to be wholly truthful. Even when no sentence in the proposal is literally false when considered in isolation, the communicated content as a whole may still be partly false, because inconvenient truths have been left out and necessary qualifications not made. Writing a paper, too, involves much more than a sober attempt to state the facts of the matter (cf. [Nickles 1992](#)). Similar considerations apply to consensus conferences. The extent to which non-epistemic factors really compromise the reliability and truth-directedness of decision procedures is an open question that I cannot answer here. But these considerations do pose a problem for the proposal under consideration: Even in those cases in science where there is a decision procedure to generate a group view that may appear to be reliably directed at truth (and thus capable of generating collective knowledge), potentially distorting non-epistemic factors are present, which could render the resulting views only luckily true and hence not instances of knowledge.

⁶ [Bird \(2010, p. 29\)](#) and [Thagard \(2010, p. 280\)](#) make a similar objection.

⁷ For this reason, [Wray \(2007\)](#) argues that only research teams can have collective knowledge.

In sum, then, although this second proposal for understanding collective knowledge is robustly collective (unlike the summative analysis from the previous section), it faces grave difficulties as an account of how scientific knowledge is collective.

4 Collective knowledge reconsidered

The conclusion so far is that existing proposals for understanding collective knowledge do not give us an adequate grasp of the sense in which scientific knowledge is collective knowledge. So let's take a step back and compare what it takes to have individual knowledge with what might be required for collective knowledge. According to the default epistemological analysis of knowledge as non-Gettiered justified true belief, there are four conditions on knowledge: (i) a belief condition, (ii) a truth condition, (iii) a justification condition, and (iv) an anti-Gettier condition. Any of these four conditions could give rise to a sense of collective knowledge if its satisfaction requires an irreducibly collective state of affairs. In what follows, however, I will focus on condition (iii).

Neither (ii) nor (iv) seem likely candidates to require irreducibly collective states of affairs for their satisfaction. On a realist conception of truth, truth is a relation between a proposition and the world, so people do not even enter the picture for condition (ii).⁸ Since condition (iv) seems irrelevant for science and is unlikely in general to require collectives for its satisfaction, I will also ignore it.

The proposals discussed in the previous two sections can effectively be understood as attempts to collectivize condition (i). Since I think these attempts have been mostly unsuccessful and my aim is to develop an alternative construal of collective knowledge which does not require collectivized belief, I will make just a brief comment here. There is an obvious sense in which it is impossible for groups to satisfy the belief condition. Groups do not have mental states in any straightforward sense. So if condition (i) is interpreted strictly, as requiring its subject to have the mental state of belief, then group knowledge immediately becomes impossible.⁹ This need not be decisive, however. Wray (2001) argues correctly that groups are capable of *accepting* views. To accept a proposition, according to Cohen, is to 'go along with it', i.e., to choose a policy of positing or postulating it, using it as a premise in one's reasoning, arguments, etc. Cohen (1989, p. 368). This isn't the same as having or acquiring a belief. Cohen argues that you can accept propositions or views you don't believe, because acceptance doesn't require that the proposition strikes you as true (which is plausibly a necessary condition for belief). Second, you can accept a proposition for other reasons, both epistemic and non-epistemic, than its striking you as true. Third, and related, accepting a proposition is under direct voluntary control, whereas belief formation is typically held to be involuntary (Alston 1998). A collective's accepting a proposition is a good analogue for the individual belief condition on knowledge, as long as the reason for

⁸ Perhaps proponents of an epistemic conception of truth would argue that satisfaction of the truth condition on knowledge requires a collective. One worry, however, is that all knowledge will come out as collective then, so that there is nothing distinctly collective about scientific knowledge. I leave it to the friends of epistemic conceptions of truth to explore this further.

⁹ This is why Giere (2006, 2007) and Thagard (2010) reject the possibility of collective knowledge.

the collective's accepting the proposition is properly truth-related. The proposition accepted should strike the collective as true, just as an individual's belief consists of a proposition that strikes that individual as true.

Let's now look at condition (iii), the justification condition. Epistemology is rife with controversy about the nature of justification and its relation to knowledge. This is not the place to enter into these controversies. Instead, I propose to identify several relatively uncontroversial characteristics of scientific justification—i.e., the kind of justification that scientific knowledge must have—to see if they necessitate an irreducibly collective rendering of condition (iii). What characteristics should a belief or accepted view have to count as an instance of scientific knowledge, in addition to its being true?

First, it should be properly based on a process of scientific inquiry. What counts as scientific inquiry varies across disciplines and over time. Formal reasoning, sophisticated observation of naturally occurring processes, experimental manipulation and intervention in nature; all of these can constitute scientific inquiry. Furthermore, astrology and alchemy were once considered as equally reputable branches of science as what we now call physics and chemistry. We can be accommodating about all this: what counts as scientific inquiry is relative to times and places. For a belief to have scientific justification it should be *properly based* on scientific inquiry. I will not attempt to spell out the notion of proper basing in detail, but the idea is that the inquiry should *give rise to* the belief in a non-deviant manner and be an *epistemically proper response* to it (e.g., it should involve an adequate assessment of the probative force of the evidence, correct reasoning, etc.). Here's an example. Suppose someone is presented with the evidence for quantum mechanics, grasps it, yet somehow fails to see and make the proper inference from it to a belief that quantum mechanics accounts for the phenomena correctly. At the same time, however, she does form this belief, but she does so on the basis of wishful thinking: She finds the thought that reality is fundamentally indeterministic immensely appealing and feels that the standard Copenhagen interpretation of quantum mechanics fits this thought perfectly. This person doesn't have scientific knowledge (although she may well be in a position to acquire it easily, because she does have the relevant evidence and could thus easily make the relevant inference) because her belief isn't induced non-deviantly by scientific inquiry.

Second, the belief should be an intended outcome of the inquiry, in the sense that it must have been the point of the inquiry to produce evidence for the belief in question. The reason for this condition is that scientists come to know lots of things during a process of inquiry that do not qualify as scientific knowledge, even though these things are, in a sense, properly based on the inquiry. Examples are: knowledge of the number of lab assistants present at a given time, costs of the equipment involved, the color of their coworker's shirt, etc. Such random pieces of knowledge have nothing to do with the goal of the inquiry and therefore shouldn't count as scientific knowledge. As it stands, however, this requirement is too strict, for it cannot do justice to serendipitous findings in science, i.e., findings that were not an intended outcome of a process of inquiry.¹⁰ Instead, then, we should require that the process of inquiry on which a belief is properly based be retrospectively re-interpretable as an intended outcome, i.e., as

¹⁰ Cf. note 3 above.

producing evidence for the belief. For the sake of brevity, however, I will henceforth omit this qualification.

Third, the process of scientific inquiry should also be properly performed. Without such a condition, a true belief based on sloppy or otherwise flawed scientific research could be scientific knowledge and that is absurd. What counts as proper performance will again vary over time and across disciplines. As novel or more advanced methods and techniques become available, standards for proper performance will tend to go up. Naturally, different disciplines will have different quality standards. Again, we can be accommodating about this and let the relevant scientific community's judgement about whether or not inquiry is properly performed be our guide in the satisfaction of this condition.

Fourth, it should be the case that the process of scientific inquiry on which the belief is based is objectively reliable in its intended domain of application, in the sense that the methods and techniques used in the inquiry make it, at the very least, more probable than not that true rather than false beliefs will be formed.¹¹ (Presumably, their reliability must be significantly higher than 50%, perhaps above 90%. Putting an exact number on this, however, would be arbitrary.) What matters is not whether scientists sincerely believe (possibly with excellent reasons) that a process of inquiry is reliable in its domain, but whether it is in fact reliable.¹² This condition is needed because we are thinking about scientific knowledge, and as we saw above, certain kinds of luck undermine knowledge. If, unbeknownst to the scientists involved, a process of inquiry is in fact highly unreliable at discovering truths, the outcomes of such a process cannot qualify as scientific knowledge. Even if no one knows about this unreliability at the time, if the methods and techniques used aren't objectively reliable, their (luckily) true outputs do not amount to scientific knowledge. This is why processes of inquiry in alchemy, astrology, and other more obscure branches of pre-modern science never yielded genuine scientific knowledge, even if people at the time didn't realize this.

Fifth and finally, scientific knowledge is supposed to be an especially high-grade type of knowledge. Although there may be different ways of fleshing out what this amounts to, for the present purposes I want to understand it as involving a kind of higher-order knowledge or understanding.¹³ It is characteristic for scientific knowledge that you do not just arrive at the truth reliably, but that you also grasp how and why this is so. Scientists aren't content merely to have and properly use reliable methods, they also want to know that they do so and understand why their methods are reliable. Scientific knowledge requires having insight in how a process of inquiry

¹¹ Scientific antirealists, as well as scientific realists of various stripes, will balk at my unqualified talk about truth. While I won't be able to satisfy antirealists who lean towards relativism, this reliability condition can be amended to accommodate weaker forms of realism, such as referential realism or structural realism, and even constructive empiricism. To do this, reliability must be understood as making it more likely that successfully referring or empirically adequate theories are found.

¹² What I am proposing here is that scientific knowledge requires what Alvin Goldman (1988) calls 'strong justification'.

¹³ This is not to deny that there may be further ways in which scientific knowledge is high-grade. Perhaps scientific knowledge requires higher reliability than knowledge generally. Perhaps it ought to be embedded in a broader network of knowledge which exemplifies internal coherence, explanatory potential, or other theoretical virtues, thus giving rise to what some authors have called 'understanding' (Kvanvig 2003; Greco 2010).

lends support to its conclusions, grasping the relevant evidential relations correctly.¹⁴ Sosa (2007, 2009) discusses a similar conception of high-grade knowledge: *reflective knowledge*. This, he says, requires of a knower that he has an ‘epistemic perspective’ on his belief, ‘a perspective from which he endorses the source of that belief, from which he can see that source as reliably truth conducive’ (Sosa 2009, p. 135). When you have such a perspective, you know that you know. You are in a position to answer affirmatively when you ask yourself whether you know and you can ward off various sorts of skeptical challenges. Sosa traces this conception of high-grade knowledge back to Descartes and even further to the antique notions *episteme* and *scientia*. My proposal is that scientific knowledge is a strong kind of reflective knowledge, whereby the subject understands that the process of inquiry on which her belief is based was properly performed and objectively reliable, and that it indeed supports her belief.¹⁵

Scientific justification, then, involves the following:

(SJ) A subject S’s belief that p has scientific justification only if it is properly based on a properly performed and objectively reliable process of scientific inquiry, the purpose of which was to gather evidence for the truth of p, and S understands this to be so.

It might be that more conditions are needed to arrive at a fully satisfactory account of scientific justification.¹⁶ (SJ) provides necessary conditions, but perhaps not a sufficient condition. Nonetheless, it will suffice for the present purposes. The questions before us now are whether scientific justification, at least in some cases, can only and irreducibly be had by collectives rather than individuals; if so, in what sense; and whether this gives rise to a proper sense of collective knowledge.

5 Collective scientific justification and collective knowledge

To argue that there is an important sense in which scientific justification can indeed only be had by a collective, I will briefly survey some characteristic features of (contemporary) science and then show how these features support the claim that much scientific knowledge is collective knowledge.

¹⁴ A constraint like this is already implicit in the second requirement above that the belief ought to be the *intended* outcome of a process of inquiry. To satisfy this requirement, scientists must understand how the inquiry is supposed to produce evidence for the belief and thus (assuming that the inquiry is successful and indeed produces knowledge) know that they know.

¹⁵ My analysis of scientific justification is thus explicitly *internalist*, since this final condition requires cognitive access to the grounds of belief and to how they support the belief. Such conditions are the hallmark of what is often called access internalism (BonJour 2010). Note, however, that this is perfectly compatible with holding that knowledge or justification in general must be analyzed externalistically. As I hinted at above, it might be that scientific knowledge is high-grade exactly because it meets internalist constraints. Ordinary knowledge need not require any such thing.

¹⁶ Post-Gettier epistemology literature contains multitudes of cleverly devised and increasingly outlandish counterexamples to proposed analyses of justification and knowledge. I will not attempt to safeguard (SJ) against such counterexamples here. Moreover, since it is far from clear that outlandish counterexamples would ever materialize in actual scientific practice they don’t pose a very pressing threat to (SJ). See, e.g., Pritchard (forthcoming), Sosa (2007), and Greco (2010) for discussion and recent attempts to offer a satisfactory general analysis of knowledge.

Contemporary science is a thoroughly collective effort, in a variety of ways (Hardwig 1985, 1991; Goldman 1999). First, many research projects in various areas of contemporary science are carried out in research teams consisting of scientists and assistants with different kinds of expertise. Such teams have a more or less official division of cognitive labor, with various members taking care of specific parts and aspects of the overall research projects and relying heavily on each other's work. While team members will typically share their results and conclusions with each other through testimony, they typically do not share all the original justification that underpins these results and conclusions. This is not just more efficient but simply inevitable, because it is impossible for a single individual to possess the relevant expertise, knowledge, and skills needed to understand and evaluate all epistemically relevant aspects of the entire inquiry. Teams can get very big. A well-known example is CERN's Large Hadron Collider in Geneva, Switzerland, where hundreds of scientists have cooperated to home in on the Higgs boson. In biomedical research, randomized clinical trials in which new drugs are tested for efficacy can involve up to hundreds of investigators and assistants, sometimes spread out over several locations in different countries. The efforts of the Intergovernmental Panel for Climate Change (IPCC) require cooperation of enormous amounts of scientists, often from different disciplines. Longitudinal studies of the characteristics and behavior of groups of people in social science are also carried out by collectives of investigators, students, and assistants.

There are two different dimensions in which epistemic dependence is necessary in such collaborative projects. Sometimes the evidence needed to substantiate a conclusion is too much for any individual to collect and process. It is not that doing so requires skills or expertise that no single individual has; it is purely a matter of time. This happens in longitudinal studies in social science. Conducting and processing thousands of surveys is not beyond the cognitive skills of an individual scientist, but it would take too long. We might say that epistemic dependence on others is *practically necessary* in these cases. In other cases, however, there is an even stronger form of epistemic dependence. Where multiple disciplines and forms of expertise are involved, individual scientists and assistants cannot check each other's contributions, because doing so requires expertise or cognitive skills that they neither possess nor can acquire easily. Such mutual epistemic dependence is *cognitively necessary*. The CERN and IPCC examples are cases in point. Interdisciplinary work—prominent in many areas of science—exemplifies such cognitively necessary epistemic dependence in particularly stark form.¹⁷

Second, even apart from formally organized collaboration in official research teams, scientists are mutually dependent in other ways, both on their predecessors and on contemporary colleagues. Research requires high degrees of specialization, which means that a lot of background knowledge has to be taken for granted. Scientists

¹⁷ The sense of necessity that I am appealing to here has slightly vague contours. It is conditioned on scientists' general opportunities and abilities to learn about other scientific subfields, on the development of background and common knowledge among scientists, and on scientists' individual circumstances, which may or may not permit them to invest additional time and effort into data gathering, processing, learning, etc. Nonetheless, it should be uncontroversial that when someone with a Ph.D. in, say, particle physics cooperates with a molecular biologist on a joint research project, they will typically have to rely on each other's expertise *necessarily*.

are embedded in subfields or research programs, which rely on established bodies of results or background assumptions—both substantive and methodological—that have been justified by their predecessors. They use equipment that they often cannot build or test themselves. By learning about the work of contemporary colleagues scientists develop new projects that build on results established by those colleagues. Some studies involve analyzing information from large databases that are created and maintained by broader collectives or even entire scientific subfields. Any contemporary scientist is thus strongly epistemically dependent on many others.

Some of this epistemic dependence could be remedied easily: Scientists could obtain the relevant original justification themselves by investing more time, consulting the relevant textbooks or original research articles, re-analyzing their colleagues' data, replicating experiments, etc. In other cases, however, the epistemic dependence is necessary in either the practical or the cognitive sense identified above.

I take these characteristics to show that, *in those cases where mutual epistemic dependence is cognitively necessary, only collectives of scientists can satisfy (SJ)*, in a sense that I will specify in due course. First, both the design and proper performance of a process of inquiry often requires the efforts of a collective. Many questions asked in contemporary science make teamwork practically and cognitively necessary, because answering them requires much more work, time, theoretical and practical knowledge, and cognition than any individual scientist could muster, even if she had plenty of time on her hands. Both thinking about what a reliable process of inquiry should look like—which is required if the process aims at the formation of a belief about what the answer to the research question is as its *intended outcome*—and *properly performing* the inquiry are things that only teams can do successfully. In doing so, moreover, they might well be relying necessarily on others who have established previous results, designed experimental procedures or apparatuses, etc.

Hence, for scientific knowledge claims that make teamwork practically and cognitively necessary, the total evidence that bears on them and that is gathered in a process of inquiry, is not had in a non-testimonial way by any one individual, but only by research teams or broader collectives of scientists, often with different training, expertise, and skills, and possibly from different disciplinary backgrounds. For cases where teamwork is a practical necessity, it could be argued that once the evidence is all collected, properly stored, and made accessible, individuals can have it non-testiminally. This is impossible, however, for research projects that involve cognitively necessary teamwork. Even if all the evidence were made accessible, individual scientists couldn't 'have' it all, because they lack the expertise and skills to grasp and assess all of it adequately.

This is crucially important in relation to last clause of (SJ); the reflective knowledge condition. Scientific justification requires that the subject *understand* that her belief is properly based on a properly performed and objectively reliable process of scientific inquiry, the purpose of which was to gather evidence for the truth of her belief. As I have explicated it above, satisfying this condition involves grasping the evidence for your belief, appreciating how it bears on your belief, and thus seeing that your belief is properly based on reliable evidence. For research where mutual epistemic dependence is cognitively necessary, this means that there is an important sense in which only the collective satisfies (SJ). Each individual scientist grasps only parts of the evidence,

appreciates only some of the evidential relations, and thus sees only partly how the overall conclusion is properly based on the inquiry. Every individual could in principle let herself be informed about all the evidence and evidential connections, but she would be taking her colleagues' testimonial reports on trust. In virtue of the different expertise and skills involved, she herself is in no position to verify what they say directly.

To say that only a collective can satisfy (SJ) in cases where mutual epistemic dependence is cognitively necessary is not to deny that individuals can (also) have scientific knowledge. The point is rather that an individual's knowledge in such cases is derivative of the collective's knowledge. To see this more clearly, suppose an interdisciplinary group of scientists wants to know whether *p*, where answering this question requires a process of inquiry that involves different scientific specializations. Assume that the process is properly performed, objectively reliable, and that the group comes to know that *p* by properly basing its belief on this process.¹⁸ Now consider how an individual scientist in the group, who was responsible for a part of the inquiry, can have knowledge that *p*. She will have collected some evidence and will be able to appreciate how that evidence bears on *p* (or on another proposition, which in turn bears on *p*). She also realizes that the evidence she understands falls short of fully supporting *p* and that she has to rely on her colleagues to provide and assess other bits of evidence. And that this is matter of necessity; she cannot obtain the relevant expertise and skills relatively easily enough to remedy her epistemic dependence. The best she can do is to collect evidence for the reliability of her colleagues' testimony. In doing so, she obtains testimony-based justification and knowledge that *p*. This, however, falls short of satisfying (SJ), because in relying on testimony-based justification, she doesn't really *understand* in a direct way that the process of inquiry on which her belief is (ultimately) based is properly performed and objectively reliable, and that the evidence it produces indeed supports *p*. In other words, she doesn't satisfy the last 'reflective knowledge' clause of (SJ).

It could be objected that she does satisfy (SJ) in a way. She understands how the testimonial evidence she does have bears on the non-testimonial evidence for *p*. So she understands that her own non-testimonial evidence in combination with her testimonial evidence supports the claim that a reliable process of inquiry has been properly performed, which has resulted in good evidence for *p*. By properly basing her belief on this combined evidence, she has perfectly good justification for *p* and also understands this to be so.

In response, I want to say that this is a good characterization of how an individual scientist in the collective can know that *p*. However, in doing so, she doesn't satisfy (SJ). Strictly speaking, she only has direct cognitive access to good evidence (i.e., her colleagues' testimony) that there is good non-testimonial evidence for *p*, but she doesn't have access to all of that non-testimonial evidence herself, because it is partly beyond her cognitive reach.¹⁹ Hence, she doesn't fully understand all the evidence for

¹⁸ By writing that 'the group comes to know' I don't mean to beg the question. Those who remain unconvinced up to this point may paraphrase this locution in their preferred way.

¹⁹ Cf. also Goldberg (2010) and Faulkner (2011) for accounts of a similar distinction between 'original' and testimony-based justification.

p and how it supports p, which would be required to satisfy (SJ). Only the combined individuals, i.e., the collective, can thus satisfy (SJ) directly.²⁰

Perhaps a comparison with another activity is helpful to draw the argument to a close. Think of a soccer team and suppose that an offensive player scores a goal. We might say that she scored the goal and deserves the credit for doing so. But although we do often talk that way, this is a loose way of putting it. The player could never have scored the goal without the combined efforts of her team members. As an offensive player, she relies necessarily on skills of other players that she herself doesn't have (and couldn't easily acquire). A better way of putting it, then, is to say that the achievement of scoring a goal is first and primarily a collective achievement. It is an individual achievement only in a derivative sense. I have argued that the same goes for instances of scientific knowledge that involve cognitively necessary epistemic dependence. Such knowledge is first and primarily collective knowledge, because only the collective can satisfy (SJ). Individual scientists can acquire such knowledge, but their individual knowledge is wholly derivative of that of collective.

To sum up, I have argued that some scientific knowledge is collective knowledge, because it is such that only a collective of scientists can satisfy the justification condition on scientific knowledge, i.e., (SJ). While individuals can also come to possess such knowledge, their knowledge remains derivative of the collective's.

6 Clarifications and objections

In this penultimate section I want to clarify my proposal further. I'll first draw attention to some of its implications and then consider several possible objections to it.

First, the proposed construal of collective knowledge leaves open whether any of the other conditions on knowledge also require a collective for their satisfaction. In particular, I am not committed to claiming that collective knowledge also requires collective belief or acceptance. I consider this an advantage, since the proposals for collectivizing the belief condition discussed in Sects. 2 and 3 face serious problems. However, if some suitable collectivized version of the belief condition were available for scientific knowledge, then nothing stands in the way of combining it with my proposal.

Second, the collectivity of a piece of scientific knowledge is not an essential property of it. In fact, something that is collective knowledge at one point may become individual knowledge later on. We come to understand better how the world works and more clever or powerful methods and techniques become available as science progresses. This makes it possible for individuals to acquire scientific justification for knowledge claims that previously required a collective.²¹ For instance, desktop computers make it possible for individuals to perform calculations or analyses much more efficiently so

²⁰ When pushed, I might grant that we could stipulate that the partially testimonial justification an individual scientist in the group acquires for p is perhaps an *indirect* way of satisfying (SJ). Nonetheless, I want to insist that only collectives can satisfy it in a *direct* way.

²¹ Note that this is a direct consequence of the fact that the cognitive necessity of epistemic dependence is premised on various changeable factors. Cf. note 17 above.

that they can now easily carry out computations that previously required a collective effort.²²

Third, as I emphasized above, if a group of scientists collectively has scientific knowledge that *p*, then individual scientists in the group can also come to know that *p* by combining their own evidence for *p* with testimonial evidence for *p* from their colleagues. Their individual knowledge then derives from the collective, which is the primary subject of knowledge in these cases. Scientists or other people outside the group can of course also acquire knowledge that *p* through regular mechanisms of testimony. (I take it that this latter point is nothing special; any account of collective knowledge should allow that individuals can acquire knowledge from a collective through testimony.)

Fourth, there are quite substantial amounts of collective knowledge on the proposed view. With the advent of ‘big science’, interdisciplinary science, and the increasing amounts specialization required in contemporary science, cognitively necessary epistemic dependence—and hence collective knowledge—becomes more and more common.

Next, I want to consider some objections that might be raised against my proposal.

Objection 1: This proposal fails as a construal of collective knowledge, because genuine collective knowledge requires that satisfaction of the belief condition on knowledge involves a collective (cf. Fagan 2011, p. 249).

Reply: First, it is unclear why this should be so. It isn’t an axiom of epistemology that collective knowledge stands or falls with a collective rendering of the belief condition. Given that there are several conditions on knowledge, any of these conditions could give rise to a sense in which knowledge is collective. Since (a) it seems intuitively compelling that (some) scientific knowledge is collective knowledge in some sense and (b) attempts to collectivize the belief condition hold little promise of capturing this sense, we have good reason to look for alternative accounts of the sense in which scientific knowledge is collective. My proposal is an attempt to develop such an account. It is thus explicitly meant to be an alternative to proposals that seek to make the belief condition collective. To object that it doesn’t do so is thus to beg the question.

Second, even if I were to accept that collective knowledge indeed requires collective belief, it doesn’t follow that it must also be irreducibly collective belief. If the justification of a claim already requires an irreducibly collective state, the belief condition can perhaps be understood along summative lines, e.g., as acceptance of the claim by the majority of the members of the collective. Since satisfaction of the justification condition is already irreducibly collective, the objection that a summative rendering of the belief condition is reducible to the beliefs and knowledge of individuals is less pressing.

²² This example potentially does introduce another form of cognitively necessary epistemic dependence, namely on the developers of hardware or software.

Objection 2: You've only shown that the *production* of scientific knowledge often irreducibly involves collectives, not that knowledge can properly be *attributed to* or *had* by collectives.

Reply: This objection assumes that collectivity in the production of knowledge is somehow eliminable once the knowledge is produced and attributed to a subject. This seems warranted in cases where it is only contingent and avoidable that a collective is involved in the production of knowledge. Presumably, even so in cases where epistemic dependence is a practical necessity in the sense specified above. However, when epistemic dependence is cognitively necessary, only a collective can satisfy (SJ). Since having scientific knowledge requires satisfying (SJ), the collective primarily *has* scientific knowledge. Individuals in the group have knowledge only derivatively, in virtue of the group's having it. Other scientists and non-scientists outside the group can acquire such knowledge through testimony.

Objection 3: This proposal overgeneralizes in objectionable ways. First, it makes a lot of testimonial knowledge collective knowledge, since it is often the case that people's epistemic dependence on others is cognitively necessary. When I rely on my financial advisor's claim that, say, taking out a mortgage is too risky, my dependence on her is necessary. So it seems the proposal entails that she and I have collective knowledge. Second, scientists often rely on background knowledge that has been established by their predecessors, and for which they couldn't obtain scientific justification themselves relatively easily. It thus follows that much scientific knowledge is had by collectives that are scattered throughout space and time.

Reply: The first part of this objection overlooks a crucial element of my proposal. It's not the case that cognitively necessary epistemic dependence always gives rise to group knowledge. Only *mutual* cognitively necessary epistemic dependence among a group of scientists (or non-scientists) who can be seen as being jointly engaged in a research project generates group knowledge, because only in such cases is it the case that the group (and not any individual on her own) can satisfy (SJ) directly. The financial advisor case is crucially different. There, it seems that the financial advisor has all the relevant expertise and simply uses your input to find out what is best for you financially—and she does all this individually. This is thus a straightforward case where you acquire testimonial knowledge.

Closer inspection also reveals the second part of the objection to be harmless. Although scientists seldom replicate the experiments behind established results that are part of their background knowledge, they do typically understand the evidence for these results very well. Or, if they don't, either they are in a position to easily acquire (or re-acquire) such understanding (for instance, by consulting original papers or textbooks), or there are others in their scientific community who are in such a position. In other words, even though scientists have not collected the original justification for their background knowledge themselves, they (or their contemporary colleagues) do typically have easy access to it and understand how it bears on their background assumptions. But this is just to say that there is no cognitively necessary epistemic dependence on historically crucial scientists such as Newton or Einstein, and hence

no need to concede that the subjects of collective scientific knowledge are quirky spatiotemporally scattered collectives.

This is reinforced by considering mathematicians before Fermat's theorem was recently proven (again?). For a time, they did *not* know the theorem to be true; exactly because nobody in the community had access to a proof (even though they did have testimonially-based justification). Only when the theorem was proven did the community come to possess scientific justification and hence scientific knowledge of it. This supports the thought, espoused above, that scientific knowledge is had by a community of contemporary scientists or individual scientists who have relatively easy access to the required justification, and not by a quirky collective of contemporary and historical scientists.

7 Conclusion

It is intuitively highly plausible that some scientific knowledge is collective knowledge. I have argued that existing proposals for understanding collective knowledge, however, do not give an adequate account of this. Instead, I propose that some scientific knowledge is collective knowledge in the following sense: There are instances of scientific knowledge which are such that only collectives of scientists can have the scientific justification that is required for it. My claim is that this is sufficient to make these instances of knowledge collective knowledge.

Acknowledgments I am grateful to the following people: the organizers of the Collective Dimensions of Science conference in Nancy in December 2011 for giving me an opportunity to present an early version of this paper; audience members at that conference, members of the Theoretical Philosophy research group at VU University Amsterdam, and three anonymous referees for this journal for incisive comments on several earlier versions of this paper. Research for this paper was made possible by an Veni grant (275-20-023) from the Netherlands Organization for Scientific Research (NWO).

References

- Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives*, 2, 257–299.
- Bird, A. (2010). Social knowing: The social sense of 'scientific knowledge'. *Philosophical Perspectives*, 24, 23–56.
- BonJour, L. (2010). Externalism/internalism. In J. Dancy, E. Sosa, & M. Steup (Eds.), *A companion to epistemology* (pp. 364–368). Oxford: Wiley-Blackwell.
- Cohen, J. (1989). Belief and acceptance. *Mind*, 98(391), 367–389.
- Corlett, J. A. (1996). *Analyzing social knowledge*. Lanham, MD: Rowman and Littlefield.
- Fagan, M. B. (2011). Is there collective scientific knowledge? Arguments from explanation. *The Philosophical Quarterly*, 61(243), 247–269.
- Faulkner, P. (2011). *Knowledge on trust*. Oxford: Oxford University Press.
- Giere, R. (2006). *Scientific perspectivism*. Chicago: University of Chicago Press.
- Giere, R. (2007). Distributed cognition without distributed knowing. *Social Epistemology*, 21(3), 313–320.
- Gilbert, M. (1987). Modelling collective belief. *Synthese*, 73(1), 185–204.
- Gilbert, M. (1989). *On social facts*. London: Routledge.
- Gilbert, M. (1994). Remarks on collective belief. In Schmitt 1994a (pp. 235–256).
- Gilbert, M. (2000). Collective belief and scientific change. In *Sociality and responsibility: New essays in plural subject theory* (pp. 37–49). Lanham, MD: Rowman & Littlefield.
- Gilbert, M. (2004). Collective Epistemology. *Episteme*, 1(2), 95–107.

- Goldberg, S. (2010). *Relying on others: An essay in epistemology*. New York: Oxford University Press.
- Goldman, A. (1988). Strong and weak justification. *Philosophical Perspectives*, 2, 51–69.
- Goldman, A. (1999). *Knowledge in a social world*. New York: Oxford University Press.
- Greco, J. (2010). *Achieving knowledge*. Cambridge: Cambridge University Press.
- Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy*, 82(7), 335–349.
- Hardwig, J. (1991). The role of trust in knowledge. *Journal of Philosophy*, 88(12), 693–708.
- Kitcher, P. (1994). Contrasting conceptions of social epistemology. In Schmitt 1994a (pp. 111–134).
- Kvanvig, J. (2003). *The value of knowledge and the Pursuit of understanding*. Cambridge: Cambridge University Press.
- List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford: Oxford University Press.
- Nickles, T. (1992). Good science as bad history: From order of knowing to order of being. In E. McMullin (Ed.), *The social dimensions of science* (pp. 85–129). Notre Dame, ID: University of Notre Dame Press.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3), 247–279.
- Quinton, A. (1976). Social objects. *Proceedings of the Aristotelian Society*, 75, 1–27.
- Roberts, B. (1989). *Serendipity: Accidental discoveries in science*. New York: Wiley.
- Rolin, K. (2008). Science as collective knowledge. *Cognitive Systems Research*, 9(1–2), 115–124.
- Rolin, K. (2010). Group justification in science. *Episteme*, 7(3), 215–231.
- Schmitt, Frederick (Ed.). (1994a). *Socializing epistemology: The social dimensions of knowledge*. Lanham, MD: Rowman & Littlefield.
- Schmitt, F. (1994b). The justification of group beliefs. In Schmitt 1994a (pp. 257–287).
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. I). Oxford: Clarendon Press.
- Sosa, E. (2009). *Reflective knowledge: Apt belief and reflective knowledge* (Vol. II). Oxford: Clarendon Press.
- Thagard, P. (1997). Collaborative knowledge. *Noûs*, 31(2), 242–261.
- Thagard, P. (2010). Explaining economic crises: Are there collective representations? *Episteme*, 7(3), 66–83.
- Tuomela, R. (1992). Group beliefs. *Synthese*, 91(3), 285–318.
- Tuomela, R. (2004). Group knowledge analyzed. *Episteme*, 1(2), 109–127.
- Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy*, 65, 157–170.
- Wray, K. B. (2001). Collective belief and acceptance. *Synthese*, 129, 319–333.
- Wray, K. B. (2007). Who has scientific knowledge? *Social Epistemology*, 21(3), 337–347.