



Collaborative knowledge: Where the distributed and commitment models merge

Spyridon Orestis Palermos¹

Received: 2 March 2021 / Accepted: 18 October 2021 / Published online: 28 February 2022
© Crown 2022

Abstract

Within analytic philosophy, the existence of collective knowledge has been motivated by means of two apparently distinct, and in direct competition with one another, theoretical approaches: (i) the commitment model and (ii) the distributed model. This paper argues, however, that to fully account for *collaborative knowledge*—i.e., a special kind of collective knowledge—both models are required. In other words, there is at least one kind of collective knowledge, the account of which requires treating the two models not as competitors but as complementary to each other. If that's correct, not only can we gain a deeper understanding of the specifics surrounding collaborative knowledge but also a clearer picture of the broader debate surrounding collective knowledge.

Keywords Collective knowledge · Joint commitment · Distributed model · Scientific research teams · Transactive memory systems · Social epistemology

1 Introduction

Epistemic collaborations such as Scientific Research Teams (de Ridder, 2014; Giere, 2002a, 2002b, 2006, 2007; Palermos, 2020), Intelligence Agencies, Transactive Memory Systems (Wegner, 1986; Wegner et al., 1985) and even Social Machines like Wikipedia (Palermos, 2017) are already powerful, social engines of knowledge production. Little doubt should there be, then, that, as time goes by, such collaboratively produced knowledge is bound to become increasingly important—especially if we factor in the rise of Web 2.0 and similar technologies. Besides its wide, practical

The article belongs to the topical collection on New Directions in Social Epistemology, edited by Adam Carter and Christoph Kelp.

✉ Spyridon Orestis Palermos
palermoss@cardiff.ac.uk

¹ School of English, Communication and Philosophy, Cardiff University, 1.44, John Percival Building, Colum Drive, Cardiff CF10 3EU, UK

importance, collaborative knowledge is also remarkable from a theoretical perspective: In several authors' minds (e.g., de Ridder, 2014; Giere, 2002a, 2002b, 2006, 2007; Palermos, 2020), this special kind of knowledge is a paradigmatic case of a wider, and somewhat controversial, phenomenon usually referred to as *collective knowledge*—i.e., knowledge that is irreducibly social.

Within analytic philosophy, the existence of collective knowledge has been motivated on the basis of what appear to be two distinct theoretical approaches: (i) the commitment model and (ii) the distributed model (Bird, 2014). While I have previously argued that a specific version of the distributed model is better suited for explaining the sense in which collaborative knowledge is collective knowledge (i.e., irreducibly social),¹ my aim in this paper is slightly different. Here, I will argue that, on closer examination, the commitment and distributed models are not as distinct as the debate so far may seem to suggest. To do so, I will demonstrate that to fully understand the workings of epistemic collaborations, we need to also pay close attention to central notions in the commitment model—for example, to such notions as *common knowledge*, *joint commitment*, *joint reasons*, and the *right to rebuke*. Doing so won't only let us acquire a better grasp of how the two models relate; it will also offer a better insight into the efficient design of future epistemic collaborations.

In the following section, I will present several versions of the commitment model (Sect. 2.1) and the distributed model (Sect. 2.2) and I will then explain why the debate so far may seem to suggest that the two models are in competition with one another (Sect. 2.3). In Sect. 3, I will argue that, even if we accept that the distributed model is better suited for accounting for the sense in which collaborative knowledge is collective knowledge, the commitment model is still necessary for fully understanding the way in which epistemic collaborations give rise to collaborative knowledge. Thus, in Sect. 4, I conclude that, at least in the case of collaborative knowledge, the commitment model and the distributed model should not be viewed as competitors but as complementary to one another.

2 Collective knowledge

According to Anthony Quinton (1975) the ascription of mental states to groups is only metaphorical and simply a means to refer to the sum of mental states possessed by the individual members of the group:

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 a summative kind.

The following then would clearly qualify as a summativist approach:

¹ Earlier attempts to motivate this claim can be found in (Palermos, 2015; Palermos 2017; Palermos and Pritchard 2013; Palermos and Pritchard 2016). Palermos (2020) contains my most recent and developed defence of this idea.

Mutual Knowledge

A group G knows that p if and only if (a) p is the case and (b) all or most members of G know that p .

For example, in Hans Christien Andersen's story, *The Emperor's New Clothes*, this is the kind of knowledge that all subjects have with regards to the emperor's nudity. And yet, despite being possessed by all or most members of the group, their unwillingness to openly express it prevents it from becoming common knowledge.² A stronger kind of summativist collective knowledge would thus be the following:

Common Knowledge

A group G knows that p if and only if (a) p is the case; (b) everyone in G knows that p ; (c) everyone in G knows that (b), and so on.³

Clearly, this kind of knowledge is more collective than the previous kind. And yet both of these kinds of collective knowledge are summative, because in both cases, the group G 's knowledge that p is reducible to the knowledge possessed by the individual members of the group. Therefore, such epistemic states would fail to count as truly collective—i.e., irreducibly social.

Thus, in response to Quinton's summativism, and to argue that collective knowledge is irreducibly social, several authors have offered non-summativist accounts of collective knowledge. Following Alexander Bird (2014) such accounts can be roughly distinguished as falling under two main models of collective knowledge: (i) the commitment model and (ii) the distributed model.

2.1 The commitment model

The commitment model has been primarily championed by Margaret Gilbert (1987, 2004, 2013). According to her view:

Joint Belief

- (i) A group G believes that p if and only if the members of G jointly accept that p .
- (ii) Members of a group G *jointly accept* that p if and only if it is common knowledge in G that the individual members of G have openly expressed a conditional commitment jointly to accept that p together with the other members of G (1987, p. 195).⁴

² See also (Bird 2010) for discussion of mutual knowledge in relation to the same story. For more on the notions of mutual and common knowledge, see (Vanderschraaf and Sillari 2005).

³ See, for example, (Gilbert 1987), (Lewis 1969, 1972) and (Schiffer, 1972).

⁴ I offer here Gilbert's formulation of her account as it initially appeared in (Gilbert 1987) for the simple stylistic reason that its main elements are packed together in a clear manner that makes for a handy quote. Her account as it appears in (2004) and (2013) remains essentially the same. One minor difference concerns the fact that in (2004) and (2013) Gilbert notes that members of the group must jointly commit to *believe* rather than *accept* that p . This is an inconsequential difference in her presentation, however, since what she essentially requires on the part of the individual members of the group is acceptance rather than belief. As she quickly notes (2004, 101–102 and 2013, 177): "A joint commitment to believe that p as a body does not require each participant personally to believe that p . [...] More positively, the joint commitment will be

(Where *common knowledge* of the joint commitment to accept that p is to be understood along the lines suggested in the previous section.)

According to Gilbert, one of the advantages of her account (1987, 2004, 2013) is its ability to capture what Emile Durkheim referred to as the ‘coercive power’ of groups. According to Gilbert:

once a group believes that p, then, *ceteris paribus*, group members are personally obliged not to deny that p or to say things that presuppose the denial of p in their ensuing interactions with each other. If someone does say something which implies that p is not the case, this person should give some sort of explanation, or qualify the statement, saying something like: "In my own view, p is not the case". A violation of these obligations is understood to be grounds for rebuke.

Of course, as Gilbert further notes, the underlying joint commitment needs not be the result of an explicit agreement between group members: “Evidently one can communicate one’s conditional commitment to participate in joint acceptance in quite subtle, nonverbal ways” (Gilbert, 1987, p. 199). Even in such cases, however, the resulting *right to rebuke* is a significant mechanism of group *cohesion*.

Obviously, Gilbert’s account can only motivate collective *beliefs*. Raimo Tuomela (2004) has noted, however, that if we require that group members not only collectively accept the target proposition but also some reason(s) in its support, then we can use the commitment model to also provide an account of collective knowledge. For present purposes, we can capture Tuomela’s account in the following way:⁵

Joint Knowledge

A group G knows that p if and only if:

- (i) It is the case that p
- (ii) The members of G jointly accept that p
- (iii) The members of G jointly accept some reason(s) for the claim that p
- (iv) It is common knowledge in G that the individual members of G have openly expressed a conditional commitment jointly to accept (a) that p and (b) their reason(s) for p, together with the other members of G

(Where *common knowledge* of the joint commitment to accept that p and the reasons for p is to be understood along the lines suggested in the previous section.)

From the point of view of the joint commitment model, the main difference between Gilbert’s and Tuomela’s account concerns the addition of (b) in condition (iv)—i.e., that group members jointly accept not only p but also some reason(s) in its support. Here is what Tuomela notes about this addition in length:

Knowledge (justified true belief) is supposed to guide action and, what is central here, to do it better than mere belief and even true belief. Why can this be said?

Footnote 4 continued

fulfilled, to some extent at least, if those concerned say that p in appropriate contexts, with an appropriate degree of confidence, and do not call p or obvious corollaries into question. Their behavior generally should be expressive of the belief that p, in the appropriate contexts. That does not mean, as said, that they must personally have that belief.”

⁵ The following is my gloss of Tuomela’s account. For more details, see his (2004); especially section II.

The justification will in general tell why the belief is more than an accidental truth, and in the case of e.g. repeated action and changing circumstances, true belief with justifying reason will fare better than mere true belief. Furthermore, in the group case knowledge based on a shared group reason will functionally fare better than that the justified true beliefs the group members might have that would be based on their individual reasons. What we thus are comparing here from the group's point of view is intentionally acting on (possibly true) belief or even merely individually justified true belief versus acting on group knowledge (viz. on the basis of a good group reason). Clearly, the group can act as a group more successfully and reliably (from an objective point of view) and take more risks (from its internal, "groupjective" point of view) in the latter than in the former case. Thus there will be both more successful acting and acting of another kind when the requirement of justification is satisfied and taken by the group members to be satisfied (2004, p. 114).

In other words, Tuomela seems to suggest, a jointly accepted reason in support of the group's jointly accepted belief will add a further layer of group *cohesion* that will increase the group's ability and resolve to successfully carry out its current and future epistemic pursuits.

2.2 The distributed model

The *distributed model* is an alternative approach to collective knowledge. Alexander Bird introduced the term to juxtapose his own view to the commitment model. The term, however, can be fruitfully employed as an umbrella term, encapsulating several related takes on collective knowledge. Below, I offer a review of a few representative ones, starting with Bird's own.

Bird is interested in social knowing as it applies to the wider scientific community; for example, in sentences such as: "a proof of Fermat's last theorem is now known" (Bird, 2010, pp. 1–2). Moreover, given that the scientific community as a whole is not, according to Bird, an established group that is capable of forming joint commitments, he turns away from the joint commitment model of collective knowledge. Instead, to account for the scientific community's social knowing, Bird introduces the alternative of the distributed model.

Bird's version of the distributed model is what he calls a functionalist approach to collective knowledge that draws on Durkheim's notion of *organic solidarity*. Bird starts with identifying the functional role of individual epistemic processes such as perceptual faculties. The functional role of such cognitive faculties, Bird claims, is to produce knowledge that will act as input to reasoning. With that in place, he then proceeds to argue that this functional role can be realised not only by brain structures, but also by *organic* social structures, essentially characterised by *division of labour*. One such structure, on Bird's view, is the wider scientific community; so, in this way, we can make sense of the claim that certain pieces of knowledge can be attributed to the wider scientific community. Bird also notes that, on his view, and in juxtaposition with the commitment model, the wider scientific community is not held together by means of joint commitments but on the basis of the mutual interdependence of its members.

However, because Bird wants to argue that the collective epistemic subject of social knowing can be as wide as the entirety of the scientific community, he further notes that the members of the socio-epistemic agents he has in mind do not have to actually collaborate in order to produce the relevant piece of knowledge. Instead, it is sufficient that they only be indirectly connected with each other. For example, according to Bird (2010, pp. 32–33), individual *S* may be a member of the community *G* that knows that *p*, even if *S* plays no role in the production of *p*, but simply benefits from *p* by using, for example, some technology that relies on *p*:

An individual [can be] included in *G* [i.e., the epistemic subject that knows that *p*] 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 (ibid. p. 32).

Moving on, Jeroen de Ridder (2014) has also put forward a version of the distributed model in order to account for collective knowledge within science. Unlike Bird, however, de Ridder's focus is not on the wider scientific community but on scientific research teams. His motivation for seeking an alternative to the commitment model is, among other things, the observation that when scientists collectively form beliefs, such beliefs are not, usually, normatively binding: "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" (de Ridder, 2014, p. 41).⁶ On the contrary, on the commitment model, each member of the group is obliged to hold on to the jointly accepted belief or else other members can rightfully rebuke them.

To account for the sense in which some scientific knowledge is collective knowledge, de Ridder starts by putting forward an account of scientific knowledge and justification, according to which in order for the epistemic agent *S* to know that *p* it is necessary that (i) *p* 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 (ii) "*S* understands this to be so." Then, like Bird, de Ridder notes the pronounced division of labour present within scientific research teams:

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. (ibid. p. 46)

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. (ibid p. 46)

⁶ According to de Ridder (2014, p. 41), another reason for turning away from the commitment model when accounting for collective scientific knowledge is that, in cases where scientists do employ normatively binding procedures for forming beliefs (such as the writing of joint papers or research proposals or consensus conferences), such procedures do not seem reliably directed to truth.

Thus, drawing on such considerations, de Ridder notes that the second condition of his preferred account of scientific knowledge and justification is often satisfied not by any individual alone, but by the scientific team as a whole. This is the effect of what de Ridder calls mutual epistemic dependence: Many times, understanding that a piece of information is derived from a properly performed and objectively reliable process requires individual contributions from the entirety of the scientific research team. Given that such understanding is a necessary condition on scientific justification, in such cases, de Ridder holds, we may talk of collective scientific knowledge: Such knowledge belongs to the entire scientific team comprising of all the mutually epistemically dependent members.

To argue specifically for collaborative knowledge (as it might emerge, for example, on the basis of Scientific Research Teams and Transactive Memory Systems (TMSs)), I have also proposed (Palermos, 2020) a version of the distributed model that I call *Distributed Virtue Reliabilism*. And like Bird and de Ridder, I too have shied away from the commitment model. This has been mainly because of its focus on what seemed to be the *aggregative* effect of individual members' intentional attitudes; a focus, which, until recently, I believed to run counter to my interest in collective knowledge that results specifically from team members' *collaboration*. Thus, to account for collaborative knowledge, I proposed instead, we need only focus on the distinctive feature of this kind of collective knowledge—i.e., the group members' *dense interactions*.

To offer such an account, I sought to combine virtue reliabilism from mainstream epistemology with the hypothesis of distributed cognition from philosophy of mind and cognitive science. To start with, I assumed that epistemic collaborations, such as scientific research teams and TMSs, are complex dynamical systems ultimately to be modelled on the basis of Dynamical Systems Theory. Thus, helping myself to insights and formalisms offered by this widely used area of mathematics, I noted that when group members engage in continuous reciprocal interactions, new properties emerge that belong to an integrated, distributed system—consisting of all interacting individuals—as a whole. In the specific case of epistemic collaborations, I argued, group members' ongoing interactions allow the group to *self-organise* and *self-regulate*, which, respectively allow the group to perform in an *epistemically reliable* and *responsible* manner. Moreover, I noted, according to virtue reliabilism, epistemic reliability and responsibility are the main two ingredients of epistemic justification. Thus, when in the process of forming a belief, group members engage in continuous reciprocal interactions to self-organise and self-regulate, the group's epistemic reliability and responsibility—i.e., the group's justification—are emergent properties that belong to the group as a whole. Consequently, I concluded, any resulting knowledge will be the collective knowledge of the whole group, consisting of all the densely interacting (i.e., collaborating) individuals.

Now, admittedly, there are interesting differences in the details of the aforementioned accounts—especially with regards to what they require from individuals to count as members of the overall epistemic group agent, which in turn has a significant effect on its potential size and shape.⁷ Nevertheless, if we turn our attention away

⁷ How large is the group agent? Is it the whole scientific community (Bird) or a well-defined research team (de Ridder and Palermos)? And are isolated contributions enough for membership to the group agent (de

from the details and back to the wider debate, we can view the above takes on the distributed model as a family of views whose common point is their focus on some form of ‘distribution’ of the epistemic effort, deriving from the target community’s, or group’s, *division of labour*.

2.3 The relation between the two models

The debate so far suggests that proponents of the two approaches for modelling collective knowledge either ignore or disregard the alternative. One may fairly wonder, then, whether the two models are direct competitors, each to be employed at the exclusion of the other.

Consider, for example, Bird’s approach to the debate. As noted earlier, Bird associates his version of the distributed model with Durkheim’s notion of *organic solidarity* which accentuates the importance of the group’s division of labour. According to Durkheim, however, organic solidarity is to be contrasted with *mechanical solidarity*, which characterises groups whose unity is achieved “through the shared beliefs (experiences, obligations, etc.) of individuals” (Bird, 2014, 10). Specifically, within groups with mechanical solidarity—which Bird appears to associate closely with the commitment model—“cohesion is achieved by the individuals in the group valuing their similarities of belief and value, their kinship, and their cooperative practices” (Bird, 2014, 10). By contrast, he notes, in groups with organic solidarity, “cohesion is achieved not by similarity among individuals but instead by difference. Unification is generated by mutual dependence arising from, above all, the division of labour” (ibid.) In other words, drawing parallels with and following Durkheim’s distinction between the two mechanisms of social cohesion—i.e., mechanical vs organic solidarity—Bird’s dialectics seem to suggest that epistemic collectives can be exclusively characterised *either* by the commitment *or* the distributed model.

Of course, Bird’s dialectics are not alone in generating the impression of a strict dichotomy between the two models. The repeated attempts of supporters of the distributed model (including those of both de Ridder and myself) to motivate their approach over the commitment model is no doubt sufficient for leading followers of the debate into thinking (at least initially) that the two models are, for all intents and purposes, in direct competition with one another.⁸ By contrast, in what follows, I argue that, at least in the case of collaborative knowledge, a combination of the two models is not only possible but necessary: A full grasp of the workings of epistemic collaborations and their process for generating collaborative knowledge requires focusing on elements from both the commitment and the distributed model.

Footnote 7 continued

Ridder) or must there be collaboration in the form of dense bidirectional interactions between members (Palermos)?

⁸ Perhaps, one possible exception to conceiving the dialectical space on the basis of a strict distinction between the commitment and distributed model is Brad Wray who notes that (2007, p. 342) committees that form joint commitments manifest organic solidarity.

3 Reconciling the commitment and distributed models

I am here interested in collaborative knowledge as this is produced by well-defined groups such as scientific research teams (as opposed to broader communities). I am therefore inclined to think that the distributed model, and specifically views such as de Ridder's or *Distributed Virtue Reliabilism* are better suited for explaining why this kind of collaborative knowledge is collective knowledge—i.e., irreducibly social.⁹ Motivating this metaphysical claim, however, is not the focus of this paper. My aim here is to show instead that, even if we do accept that the distributed model is better suited for capturing the collective nature of collaborative knowledge, a full explanation of how epistemic collaborations operate would require employing elements of the commitment model too.

For example, in order to illustrate what collaborative knowledge amounts to, I have previously used such exemplars of epistemic collaborations as scientific research teams and Transactive Memory Systems (Wegner et al., 1985)—i.e., groups of two or more individuals who collaboratively encode, store and retrieve information. Consider the following example by Wegner and his colleagues:

Suppose we are spending an evening with Rudy and Lulu, a couple married for several years. Lulu is in another room for the moment, and we happen to ask Rudy where they got that wonderful stuffed Canadian goose on the mantle. He says “we were in British Columbia...,” and then bellows, “Lulu! What was the name of that place where we got the goose?” Lulu returns to the room to say that it was near Kelowna or Penticton—somewhere along lake Okanagan. Rudy says, “Yes, in that area with all the fruit stands.” Lulu finally makes the identification: Peachland (Wegner et al., 1985, p. 257)

According to Wegner et al., this “observable interaction between individuals entails not only the transfer of knowledge, but the construction of a knowledge acquiring, knowledge-holding and knowledge-using system that is greater than the sum of its individual member systems” (1985, p. 256).

⁹ Between de Ridder's and my view (i.e., Distributed Virtue Reliabilism), the latter might be offering a stronger—albeit more complicated—explanation of why collaborative knowledge is irreducibly social. To see why this might be, consider, first, how de Ridder (2014, pp. 39–40) summarises the way the commitment model attempts to motivate the existence of collective knowledge: “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 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.” The problem with this approach to collective knowledge is that, although collective knowledge may not be fully reducible to the *epistemic states* (i.e., the beliefs and knowledge) of individuals, because it also consists of the individual's intentions to let the target view stand as the group's view, it is still fully reducible to the sum of the relevant individual cognitive states. Unfortunately, a similar worry can be raised against de Ridder's view, which is that a scientific team's knowledge counts as collective, because in order to fully understand its epistemic process the group requires input from all the participating members. The problem is that such understanding and knowledge appears to be again reducible to the sum of individual understanding possessed by the members of the group. By contrast, my view (2017, 2020), which relies on the notions of ‘coupled systems’ and their ‘emergent properties’ from Dynamical Systems Theory, is, arguably, in a better position to resist the reduction of the collective's epistemic properties to the sum of the cognitive properties of the participating individuals.

Now, following Wegner et al. (1985), I have drawn attention (Palermos, 2017) to the observation that there are certain practical preconditions that need to be in place for the members of TMSs to give rise to an integrated system that is capable of collaboratively knowing and recollecting. The first such precondition is best captured by the notion of a *differentiated structure* (Wegner et al., 1985). According to Wegner and his colleagues, the existence of a differentiated structure allows each member to be aware and take advantage of everyone else's knowledge and expertise. In this way members know when it is appropriate to rely on others and when it is time for them to take action themselves. Obviously, and in line with the distributed model, it is on the basis of such a differentiated structure that *division of labour* may be achieved.

As I have also noted (2017), however, according to Wegner and his colleagues, there is yet another prerequisite that needs to be in place for the group's differentiated structure (and the resulting division of labour) to emerge. This more basic, additional ingredient can be usefully captured by the notion of *common knowledge*.¹⁰ That is, according to Wegner et al., (1985), group members need a common culture, language and knowledge of general topics so that they can adequately understand each other, communicate, become acquainted with one another and eventually build their differentiated structure.

Moreover, I have claimed (Palermos, 2017), *differentiated structure* and the underlying *common knowledge* are not practically useful only in the case of TMSs. Rather, they seem to also play important roles in the development of scientific research teams too—the other example of epistemic collaborations I have focused on. To successfully communicate with each other, research collaborators must share some common knowledge “or, in other words, what Kuhn (1962) calls a ‘scientific paradigm’ (briefly, a set of agreed-upon metaphysical assumptions, methodological rules, scientific theories, techniques for using equipment, etc.)” (Palermos, 2017, p. 14). But to “efficiently allocate the workload of performing a complex experiment or solving a research problem,” I noted further, they need to also possess some expertise that will essentially provide their group with a differentiated structure. Only when these practical prerequisites are in place, can the members of the research team begin to collaborate by efficiently interacting with each other.

So, the upshot so far seems to be the following: Epistemic collaborations are primarily characterised by the existence of division of labour. This division of labour requires that the group possess a differentiated structure, and this, in turn, requires a body of common knowledge between the members of the group. Now, if we recall, common knowledge is an important concept within the joint commitment model. And although common knowledge is not the same as joint belief or joint knowledge, recognising its importance for epistemic collaborations might be a hint that the commitment model might be relevant too.

Before moving on to the more central aspects of the commitment model, let us look into common knowledge a little further. Recall that common knowledge that *p* is stronger than mutual knowledge that *p* in the sense that not only all or most members of the group know that *p* but, also, all or most members of the group know that all or most

¹⁰ Wegner et al. (1985) do not explicitly use the label ‘common knowledge’ to refer to this aspect of Transactive Memory Systems. Their discussion, however, makes it clear that this is what they have in mind.

of the other members know that p . At the same time, it is weaker than joint belief and knowledge in the sense that if a group member denies that p , this does not entitle other members to rebuke them. Thus, common knowledge provides common grounds for members of a group to sufficiently communicate with each other, but denying aspects of it does not warrant non-epistemic criticism from others in the group.¹¹ Perhaps, others can openly express their surprise, but, since no joint commitments have been violated, they have no right to chastise their ‘abdicated’ colleagues.

With the above in mind, let us briefly return to my past remarks on common knowledge. As noted above, in (Palermos, 2017), I suggested that, in the case of scientific research teams, all parts of a paradigm may constitute common knowledge. On second thought, however, this is, strictly speaking, false—metaphysical assumptions, for example, are obviously excluded since, by definition, they cannot be known. Nevertheless, the accuracy of experimental techniques or the validity of methodological rules and scientific theories may indeed count as part of scientists’ common knowledge, depending on the role they play within the project of their group. If these are elements at the periphery of the team’s research project, constituting, if you like, part of the scientists’ ‘common language,’ but with no direct impact on their actual collaboration, then they may constitute common knowledge indeed. Such common knowledge can play a facilitatory role in letting scientists communicate with each other about broader issues of scientific interest, but since it does not bear on what they are actually working on, such that no (explicit or implicit) joint commitment to accept them is necessary in the first place, denying aspects of it cannot and should not elicit reproach (most fitting in such cases, perhaps, is mere intellectual debate).

Nevertheless, within scientific research projects, there are propositions whose rejection by members of the corresponding research team would constitute violation of a joint commitment to accept them, in a way that would clearly warrant rebuke. Imagine, for example, a team of astrophysicists discovering that one of their colleagues denies one of the postulates of Einstein’s Special Relativity (e.g., that the laws of physics are the same across all inertial frames of reference).

Thus, while the role of common knowledge should not be underplayed, it seems more is required with respect to certain propositions in order to sustain collaborations. In a way, this should not come as a surprise at all. Humans are complicated entities, with particular yet constantly changing idiosyncrasies, personal preferences, priorities and agendas. So, with respect to matters of opinion, it is not at all rare to change our minds. But the same is true with respect to the things we take ourselves to know and reasonably believe. For example, it is not uncommon to suddenly find ourselves presented with defeaters that call into question what we take ourselves to know. In such cases, it is all too easy to reject our knowledge or at least suspend judgment, even if, objectively speaking, many times this is the wrong response—after all defeaters can be, unbeknownst to us, misleading. But, whatever the reasons of such doxastic and epistemic changes of mind might be, usually, such changes also have an effect on our actions and intentions, such as the kind of inquiries and projects we are prepared to participate in. Thus, when individuals agree to engage in a collaborative project, they

¹¹ In fact, Bird (2010, p. 21; 2014 fn. 6) recognises that common knowledge is necessary for scientific cooperation, even within the wider scientific community.

need to ensure that, throughout its duration at least, there will be agreement on certain standards, values, assumptions, attitudes toward certain kind of evidence, a given body of knowledge and other propositions, which were they to be called into question by the participants, it would significantly hinder their epistemic collaboration.

This is where the joint commitment model, with its focus on the right to rebuke, comes into play. Joint commitments consolidate the group by preventing members of the collaboration from *epistemically diverging*—i.e., from easily changing their minds with respect to beliefs and knowledge, whose acceptance is crucial for communicating, avoiding disagreement and thereby sustaining their coordinated epistemic performance.¹²

Gilbert's account of collective beliefs may thus be necessary for capturing the manner in which members of epistemic collaborations, such as a scientific research teams, jointly commit to accepting certain propositions stand as the view of their group. Gilbert (2000) herself has argued for using her model to account for what we may call the scientific community's 'consensus' and Kristina Rolin (2008) has expanded Gilbert's model to give an account of what we may refer to as the scientific community's 'background assumptions'—for example assumptions "required to establish the relevance of empirical evidence to a theory or a hypothesis" (2008, p. 119). I wish to remain silent as to whether the wider scientific community collectively knows or at least merely believes certain propositions. As far as I can see, however, Gilbert's model is ideally placed to account for the manner in which members of scientific research projects are committed to certain metaphysical and background assumptions, the rejection of which would amount to having abandoned their research project.¹³ Specifically, following Imre Lakatos (1970), I have in mind the kind of assumptions that inhabit the 'hard core' of scientific research programmes and which are deemed "irrefutable by the methodological decision" of the programme's participants (*ibid.*, p. 132). These are fundamental assumptions for which *no, or at least no widely acceptable or sufficiently articulated, justification and/or evidence exists*, but which nevertheless provide the research programme with its distinctive identity, such that changing them simultaneously amounts to having abandoned the research programme. What might such assumptions look like? Rolin (2008) mentions beliefs regarding the relevance of certain evidence with respect to hypotheses—consider, for example, a behaviorist's commitment to rely on overt behavior alone. Other examples could be commitment to naturalism, materialism, a set of logical axioms, etc.

Therefore, the commitment model as propounded by Gilbert does appear to play a significant role in the operation of epistemic collaborations.¹⁴ What, then, of Tuomela's

¹² Joint commitments seem to also play a similar—though not identical—role in (collaborative) Group Know-How (which is distinct from the kind of collaborative *propositional* knowledge we are here interested in): In the case of collaborative Group Know-How, joint commitments may prevent members from practically (as opposed to epistemically) diverging. For more details, see section V of (Palermos & Tollefsen, 2018).

¹³ In a footnote, Bird (2014, fn. 6) notes in passing that he himself is also open to the idea that Gilbert's model could be applicable in the case of collaborative research teams.

¹⁴ Strictly speaking, I have only argued for the relevance of Gilbert's account to epistemic collaborations such as scientific research teams. I am positive, however, that, by extrapolation, the same can be said about epistemic collaborations in general. Here I am focusing on scientific research teams, because it is easier to provide specific examples of joint commitments to believe.

stronger version? Do joint commitments to reasons as well as beliefs play a role within epistemic collaborations? The answer, I believe, is yes. For the members of scientific research teams are not only committed to accepting certain assumptions. Often, to carry out their projects, they are also committed to specific propositions along with the associated reasons. For example, these might be commitments to accept certain yet-to-be-replicated (or yet-to-be-peer-reviewed) reasons in support of some cutting-edge scientific theory, a body of newly acquired evidence or the validity of an innovative method for conducting experiments. Such commitment to the reasons in support of certain propositions can play a crucial role in carrying out collaborative research projects, in a way that rejecting them would stall the collaboration or even cause it to dissolve. So what might this role be? Why think that in such cases joint commitment to accept a proposition *as well as* reasons in its support is more appropriate than merely accepting the proposition alone? Tuomela's analysis, I believe, is highly suggestive.

As we previously saw, Tuomela thinks that jointly accepting reasons in support of a proposition provides an additional layer of cohesion to the group, increasing the group's ability and resolve to successfully carry out its current and future epistemic pursuits. This additional joint acceptance of reasons might be inapplicable in the case of metaphysical assumptions or unnecessary with regards to well-established assumptions that only very few people, if anyone at all, would conceive of rejecting. Nevertheless, epistemic collaborations often rely on evidence, hypotheses and other assumptions and pieces of information at the forefront of their discipline. Due to their vanguard status, such propositions may still be considered controversial, even if some justification exists in their support. Perhaps the available justification is itself somewhat contentious or it is intelligible only to very few (or even to no) members in the group. In cases like this, jointly accepting the existing justification for a proposition in addition to the proposition itself can help the group stay focused on its inquiry without becoming easily distracted by outsiders' dissenting voices, doubts or criticisms, essentially allowing them to avoid disagreement and thereby keep up with their project. Thus, Tuomela's version of the joint commitment model also seems to often play a central role with respect to epistemic collaborations' cohesion.

Therefore, as the above indicates, in order to fully understand the workings of epistemic collaborations, it often seems necessary to appeal to all common knowledge, joint beliefs and joint knowledge. To avoid confusion, however, it is useful be clear about what I think is each model's distinctive contribution. As I have previously argued, the best explanation of why a certain piece of collaboratively produced propositional knowledge, p , counts as irreducibly social knowledge is most likely offered by (some appropriate version of) the distributed model. Thus, if we only care about finding out why a certain piece of collaboratively produced knowledge p amounts to *collective knowledge*, we need look no further than the distributed model; in other words, joint commitments with regards to p are not necessary to make knowledge that p irreducibly social (though their existence could of course add to its irreducibly social nature). Nevertheless, common knowledge of and joint commitments to other propositions q , r , z appear to be prerequisites for the efficient operation of the epistemic collaboration. Thus, to fully understand how epistemic collaborations produce collective knowledge, it is necessary to employ both the distributed and the joint commitments models.

Another point that also needs clarification concerns the modal strength of this claim. Do we *always* need to complement the distributed model with some version of the joint commitment model? I am inclined to think that the answer is positive, at least with regards to joint beliefs (and common knowledge). Appealing to joint knowledge, by contrast, might not always be necessary. It all depends on the target collaboration and proposition at hand.

To explain what I have in mind, it appears that even the simplest epistemic collaborations—e.g., TMSs made up of old couples, or teams of colleagues and students engaging in brainstorming—need to possess some common knowledge to communicate with and understand each other so as to efficiently collaborate. Likewise, the efficient collaboration of such simple teams seems to also always require the presence of joint commitments towards certain basic beliefs, values and/or assumptions that entitle participants to rebuke their teammates when they breach them.

Imagine, for example, an old philosophers' couple who have been married for several years: One day, in the process of attempting to engage in collaborative remembering, one of them discovers that their partner is (perhaps playfully) not committed to the claim that the world is more than five minutes old or the claim that there is an external world. In such a case, one would clearly be disposed and have the right to rebuke their partner for breaching their joint commitment to let the denial of these hypotheses stand, outside skeptical philosophical discourse. Or imagine that the members of a work group dedicated to identifying the best strategy for maximising the company's profit discovered that one of their colleagues does not care about the company's reputation. Clearly, upon finding out, they would not be just surprised; rather they would feel disposed to chastise their colleague for not sharing this value with them. This is a value that they are expected—not just assumed—to uphold.

The upshot is that such and similar joint commitments to values, assumptions and beliefs seem to be necessary for the workings of even the simplest epistemic collaborations. By contrast, joint knowledge in Tuomela's sense might not always be necessary for engaging in epistemic collaborations. Previously, I claimed that joint knowledge of certain propositions is useful to epistemic collaborations when their members need to rely on controversial propositions at the forefront of their discipline or on propositions the justification for which is contentious or unintelligible to most members. In such cases, jointly accepting not just the proposition at hand but also existing reasons in its support can add an extra layer of cohesion required for the efficient operation of the collaboration. Now, my reason for doubting joint knowledge is always necessary for the efficient operation of epistemic collaborations is that relying on such controversial propositions (steadfastness to which can be significantly facilitated by also jointly committing to reasons in their support) seems to be rarely the case for simple collaborations such as TMSs.

Nevertheless, even if joint knowledge is not always necessary for the efficient operation of epistemic collaborations (though, of course, I am happy to be convinced otherwise), the necessity of the commitment model in understanding the efficient operation of epistemic collaborations seems to be still well-motivated by the apparent need for joint commitments to beliefs, values and assumptions in all kinds of epistemic collaborations—from the simplest to the most complicated ones. If that much

is correct, then the commitment and distributed models are, indeed in an important sense, complementary to one another.¹⁵

4 Conclusion

In line with the distributed model, I have assumed that the distinctive feature of collaborative knowledge and what best explains its collective nature is group members' interdependence and dense interactivity, resulting from the group's division of labour (i.e., its differentiated structure). And yet, I have argued, within epistemic collaborations, joint commitments also play a role too crucial to ignore: To interact in a coordinated manner, group members need to both agree on and hold fast to certain assumptions or even a specific body of knowledge. To this end, they jointly commit to certain beliefs and perhaps also to reasons in their support.

The above goes to show, then, that the commitment and the distributed model are both required to account for collaborative knowledge. Despite what the debate on collective knowledge may have so far seemed to suggest, therefore, there is at least one kind of collective knowledge the account of which requires treating the two models not as competitors but as complementary to each other.

Could this be true of more instances of collective knowledge?

Acknowledgements I am thankful to the audience of the Social Ontology 2021 conference for their useful feedback. I would also like to thank two anonymous referees for *Synthese* whose insightful comments helped me improve the paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bird, A. (2010). Social knowing: The social sense of “scientific knowledge.” *Philosophical Perspectives*, 24, 23–56.
- Bird, A. (2014). When is there a group that knows? Distributed cognition, scientific knowledge, and the social epistemic subject. In J. Lackey (Ed.), *Essays in collective epistemology* (pp. 42–63). Oxford University Press.
- De Ridder, J. (2014). Epistemic dependence and collective scientific knowledge. *Synthese*, 191(1), 37–53.
- Giere, R. (2002a). Discussion note: Distributed cognition in epistemic cultures. *Philosophy of Science*. <https://doi.org/10.1086/344627>
- Giere, R. (2002b). Scientific cognition as distributed cognition. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *Cognitive bases of science*. Cambridge University Press.
- Giere, R. (2006). The role of agency in distributed cognitive systems. *Philosophy of Science*, 73, 710–719.

¹⁵ I would like to thank an anonymous referee for prompting me to expand on the modal strength of my claim regarding the necessity of the commitment model for understanding the workings of epistemic collaborations.

- Giere, R. (2007). Distributed cognition without distributed knowing. *Social Epistemology*, 21(3), 313–320.
- Gilbert, M. (2000). “Collective Belief and Scientific Change,” In: *Sociality and Responsibility: New Essays in Plural Subject Theory*, Lanham, MD: Rowman & Littlefield, pp. 37–49.
- Gilbert, M. (1987). Modelling collective belief. *Synthese*, 73(1), 185–204.
- Gilbert, M. (2004). Collective epistemology. *Episteme*, 1(2), 95–107.
- Gilbert, M. (2013). *Joint commitment: How we make the social world*. Oxford University Press.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago press.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge*. Cambridge University Press.
- Lewis, D. (1972). ‘Language and Language’, in *Minnesota Studies for the Philosophy of Science*, VII, pp. 3–35.
- Lewis, D. (1969). *Convention*. Harvard University Press.
- Palermos, S. O., & Pritchard, D. (2016). The distribution of epistemic agency. *Social epistemology and epistemic agency: De-centralizing epistemic agency*, 109–26.
- Palermos, S. O. (2015). Active externalism, virtue reliabilism and scientific knowledge. *Synthese*, 192(9), 2955–2986.
- Palermos, S. O. (2017). Social machines: A philosophical engineering. *Phenomenology and the Cognitive Sciences*, 16(5), 953–978.
- Palermos, S. O. (2020). Epistemic collaborations: distributed cognition and virtue reliabilism. *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00258-9>
- Palermos, S. O., & Pritchard, D. (2013). Extended knowledge and social epistemology. *Social Epistemology Review and Reply Collective*, 2(8), 105–120.
- Palermos, S. O., & Tollefsen, D. (2018). Group know-how. In J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos & D. Pritchard (Eds.), *Socially extended epistemology*. Oxford University Press.
- Quinton, A. (1975). Social objects. *Proceedings of the Aristotelian Society*, 76, 1–27.
- Rolin, K. (2008). Science as collective knowledge. *Cognitive Systems Research*, 9(1–2), 115–124.
- Schiffer, S. (1972). *Meaning*. Oxford University Press.
- Tuomela, R. (2004). Group knowledge analyzed. *Episteme*, 1(2), 109–127.
- Vanderschraaf, P., & Sillari, G. (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/>
- Wegner, D. M., Giuliano, T., & Hertel, P. T. (1985). Cognitive interdependence in close relationships. In D. W. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253–276). Springer.
- Wegner. (1986). *Theories of group behavior*. New York: Springer-Verlag
- Wray, K. B. (2007). Who has scientific knowledge? *Social Epistemology*, 21(3), 337–347.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.