



Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism

Spyridon Orestis Palermos¹

Received: 1 July 2019 / Accepted: 1 April 2020
© The Author(s) 2020

Abstract

Strong epistemic anti-individualism—i.e., the claim that knowledge can be irreducibly social—is increasingly debated within mainstream and social epistemology. Most existing approaches attempt to argue for the view on the basis of aggregative analyses, which focus on the way certain groups aggregate the epistemic attitudes of their members. Such approaches are well motivated, given that many groups to which we often ascribe group knowledge—such as juries and committees—operate in this way. Yet another way that group knowledge can be generated is on the basis of *epistemic collaborations*, such as scientific research teams and Transactive Memory Systems. To produce knowledge, epistemic collaborations rely heavily on the mutual interactions of their group members. This is a distinctive feature of epistemic collaborations that renders them resistant to aggregative analyses. To accommodate this kind of group knowledge, the paper combines virtue reliabilism with the hypothesis of distributed cognition in order to introduce the hybrid approach of *distributed virtue reliabilism*. On this view, (1) beliefs produced by epistemic collaborations entertain positive epistemic standing (i.e., they are both reliable and epistemically responsible) in virtue of the mutual interactions of their group members; (2) this positive epistemic standing is a collective property; (3) epistemic collaborations qualify as epistemic group agents; (4) collaborative knowledge is a special kind of group knowledge, motivating strong epistemic anti-individualism in a distinctive way.

Keywords Group knowledge · Strong epistemic anti-individualism · The hypothesis of distributed cognition · Epistemic group agents · Distributed virtue reliabilism

✉ 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

1 Epistemic Collaborations

Knowledge is many times attributed to entire groups such as committees, corporations, intelligence agencies and scientific research teams. Interpreting such collective knowledge ascriptions literally amounts to accepting strong epistemic anti-individualism—the increasingly debated claim that knowledge, under the right conditions, can be irreducibly social.¹

Several epistemologists attempt to motivate this view by focusing on the way certain groups (such as juries and committees) aggregate the epistemic attitudes of their members in order to form an epistemic attitude on the part of the collective. For example, Tuomela (2004), Gilbert (1994, 2007a, b, 2010) and Rolin (2008) focus on cases where group members jointly believe or accept individually derived justified true propositions as the representative view of their group. In a similar fashion, Goldman's 'social process reliabilism' (2014) is cashed out in terms of reliable processes for aggregating the justificatory statuses of group members' true beliefs and List and Pettit (List 2005; List and Pettit 2011) focus on processes for aggregating individual judgments so as to track certain truths.² Arguing for the existence of group knowledge by focusing on processes for aggregating the epistemic attitudes of group members is, indeed, well-motivated, given that many groups to which we often ascribe group knowledge—such as juries and committees—operate in this way.

Yet another way that group knowledge may be generated is on the basis of *epistemic collaborations*, such as the ones taking place between the members of scientific research teams, intelligence agencies and medical teams. This appears to be a special kind of group knowledge that is not amenable to aggregative analyses. The reason is that the final product of epistemic collaborations—i.e., true beliefs with positive epistemic standing—is not generated by adding together the individually derived epistemic attitudes of their members. Instead, in epistemic collaborations, knowledge is produced on the basis of the participants' collaborative interactions. Focusing on procedures for merely aggregating the epistemic attitudes of isolated individuals cannot therefore capture the process by which epistemic collaborations operate.³ To account for this special kind of group knowledge, mainstream and social epistemology need an alternative, non-aggregative approach.

In the following sections, the paper offers such an account that consists in the combination of virtue reliabilism with the hypothesis of distributed cognition. The

¹ See, for example, Hardwig (1985), Tollefsen (2006, 2015), Bird (2010, 2014), List (2005), Lackey (2014a, 4b, 2016), Goldman (1999, 2004), (de Ridder 2014) and (Wray 2007).

² List (2005) and List and Pettit (2011) follow Nozick's (1981) 'truth tracking' approach to knowledge.

³ Goldman (2014) notes that, sometimes, the members of the groups he is interested in may arrive at their individual epistemic attitudes—which will then be aggregated—on the basis of interacting with each other. As he repeatedly notes in his essay, however, his approach to group knowledge focuses exclusively on the process of aggregating the epistemic attitudes of individuals (however these are derived). As it stands, therefore, Goldman's account does not aim to explain (aspects of) collective knowledge that is collaboratively produced. Its focus is specifically on group knowledge that is aggregatively produced.

motivation for this hybrid approach is threefold and it centers around the distinctive feature of epistemic collaborations—i.e., their members' collaborative interactions. First, according to the hypothesis of distributed cognition, appropriately interacting individuals can give rise to distributed cognitive systems that manifest collective properties. Such properties may be collective skills such as sport-team strategies (Cooke et al. 2013; Duarte et al. 2012, 2013a, b), collective foraging (Theiner et al. 2010) and collective memory (Wegner et al. 1985). The paper argues that, in the case of epistemic collaborations, the collective cognitive property is the resulting beliefs' positive epistemic standing. To make sense of this claim from an epistemological perspective leads to the second reason for employing the suggested hybrid approach. Ideally—from the point of view of mainstream epistemology—a group's positive epistemic standing should be modeled after an account of the positive epistemic standing of individuals. The claim here is that the positive epistemic standing of epistemic collaborations results from their members' collaborative interactions. Consequently, we need to draw on an account of individual knowledge, according to which the positive epistemic standing of an individual agent's beliefs arises out of collaborative interactions between different parts of their cognitive system. Virtue reliabilism puts forward precisely such an account. Accordingly, the third and final reason for this hybrid approach is that virtue reliabilism can be read along the lines suggested by the hypothesis of distributed cognition. On such a reading, it will be argued, epistemic collaborations collectively manifest—on the basis of their members' interactions—the property of positive epistemic standing. Since this is a collective property, knowledge produced by epistemic collaborations can be used to motivate strong epistemic anti-individualism.

The paper proceeds in the following stages. To illustrate the special character of collaborative knowledge, Sect. 2 turns to the fields of cognitive psychology and philosophy of science to present two examples of epistemic collaborations: (1) Transactive Memory Systems and (2) scientific research teams. Both examples demonstrate that, within each field, collaborative knowledge is taken to be irreducibly social, because it arises out of the members' mutual interactions. Section 3 introduces the hypothesis of distributed cognition (Sect. 3.1) and virtue reliabilism (Sect. 3.2), and then combines the two theories to explain how epistemic collaborations collectively manifest the property of positive epistemic standing (Sect. 3.3). Finally, to clarify the account on offer, Sect. 4 contrasts it with the alternative of distributed reliabilism (Michaelian and Arango-Muñoz 2018). By way of juxtaposing the two views, it will also become clear how the notion of collective epistemic standing can motivate the existence of epistemic group agents. As it will become apparent, motivating the existence of such epistemic group entities is necessary, if we are to avoid the claim that collaborative knowledge is knowledge without a knowing subject (Giere 2007; Popper 1968).

Before starting, it is worth clarifying a point regarding the paper's dialectics. Knowledge is largely thought to be a composite notion, whose main two components are (1) belief and (2) its associated positive epistemic standing (usually referred to as justification or warrant). Many of the existing attempts to motivate strong epistemic anti-individualism focus on the belief component—i.e., they attempt to argue that instances of knowledge are irreducibly social, because the relevant propositions are,

in some appropriate sense, collectively believed or accepted.⁴ While such approaches to collective knowledge merit serious consideration, it is not clear whether they are successful in showing that the group's beliefs (and thus the pieces of knowledge they form part of) are, indeed, irreducibly social.⁵ Whatever the outcome of this debate may be, the present account does not hinge on it. This is because, as noted above, the approach on offer argues for the irreducibly social nature of collaborative knowledge by focusing, instead, on the component of positive epistemic standing. That is, collaborative knowledge, it will be argued, can motivate strong epistemic anti-individualism, not because the relevant propositions are collectively believed, but because of the collective nature of the belief-generating process that confers positive epistemic standing to them. Of course, if the resulting propositions turn, on some appropriate construal, to also be held by the group in an irreducibly social manner, that would be a welcome result. But if the following account is correct, arguing for this additional claim is not necessary for demonstrating that collaborative knowledge can motivate the truth of strong epistemic anti-individualism.

2 Examples of Collaborative Knowledge

To get a grip on the phenomenon of collaborative knowledge, it is useful to consider two examples that are available in the literature. The first one concerns Transactive Memory Systems and it comes from the field of cognitive psychology. The second example is the case of scientific research teams and it comes from the field of philosophy of science. In both cases, collaborative knowledge is considered to be irreducibly social, because it arises out of members' collaborative interactions.

2.1 Transactive Memory Systems

Memory is often cited as a justificatory process on the basis of which individuals know propositions that they have encountered in the past. While memory is standardly associated with individual knowers, recent studies within cognitive psychology suggest that memory can be instantiated by groups of two or more individuals. Known as Transactive Memory Systems (TMS), such groups can be used in order "to conceptualize how people in close relationships may depend on each other for acquiring, remembering, and generating knowledge" (Wegner et al. 1985, p. 253).

"To build a transactive memory is to acquire a set of communication processes whereby two minds can work as one" (Wegner et al. 1985, 263). This involves skillful, interactive, and simultaneous coordination between people (Bernier et al. 2008). Collaboration allows the group to reliably recall pieces of information in a manner that would be unavailable to its members were they to recollect on their own. Consider the following example:

⁴ See, for example, Gilbert (2007a, b, 2010), and Tuomela (2004).

⁵ For an overview of the existing debate, see Tollefsen (2004).

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)

Wegner et al. explain that, during the discussion between Rudy and Lulu, the various ideas they exchange lead them through and elicit their individual memories. “In a process of interactive cueing, they move sequentially toward the retrieval of a memory trace, the existence of which is known to both of them. And it is possible that without each other, neither Rudy nor Lulu could have produced the item” (1985, p. 257).

Moreover, transactive processes do not only take place during the retrieval of memories. They may also occur during the processes of encoding as well as storage (Wegner et al. 1985). When partners discuss about an event as they are experiencing it, their discussion, far from being a mere rehash of the original event, can lead them to encode memories that are qualitatively different from the memories that they would have acquired, had they been on their own. Similarly, if the dyad engages in transactive processes after the event has taken place, it may qualitatively alter the members’ initial, private memories of the event, thereby also having an effect on the process of storage.

Accordingly, Wegner et al. claim, “the 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).

2.2 Scientific Research Teams

In a series of papers, Giere (Giere 2002a, b, 2006, 2007; Giere and Moffat 2003) has proposed that collaborative experiments give rise to knowledge that is irreducibly social.

To motivate this claim, Giere places emphasis on the scientists’ coordinated interactions in generating scientific knowledge. In much of scientific experimentation, “completing the task requires coordinated action by several different people” (2006, p. 711). This coordinated action “makes possible the acquisition of knowledge that no single person, or a group of people without instruments, could possibly acquire” (2003, p. 305). Thus, to “understand the workings of the big cognitive system one has to consider the human–machine interactions as well as the human–human interactions” (Giere 2002b, p. 292). The examples that Giere proposes are High Energy Physics (HEP) experiments (2002a), the Hubble Space Telescope (2006), and Latour’s (1999) (Giere and Mofat 2003) example of a scientific investigation that seeks to determine whether the Amazonian rainforest is encroaching on the adjacent savannah or whether the savannah is encroaching on the rainforest.

Knorr-Cetina's (1999) ethnographic study of HEP experiments puts forward similar remarks. She notes that large collaborations are not run by any individual alone and no individual is responsible for their management and organization. Such experiments are managed, instead, by *discourse*. Discourse “channels individual knowledge in the experiment, providing it with a sort of distributed cognition [...], which flows from the astonishingly intricate webs of communication pathways” (ibid. 173). These ongoing interactions between participants form a grid of discourse spaces, which “was and is today perhaps the most important vehicle of experimental coordination and integration” (ibid. 174). Knorr-Cetina lists, among others, research and development meetings, panel meetings, institute meetings, steering group meetings, collaboration meetings, referee meetings, accelerator meetings, fixed committee meetings, special workshops dedicated to detector complexes, ‘submeetings of some of the former, and the very important “meetings after the meeting” (the informal exchanges that occur after scheduled events)’ (ibid., 174).

It is evident that, just like TMSs so in the case of collaborative science, the participants’ ongoing mutual interactions play a central role in coordinating the activities that make up the overall process of scientific experimentation. The above authors take this as evidence that the resulting knowledge is not the product of any individual alone or the aggregate of the knowledge possessed by the individual scientists. Instead, knowledge acquired on the basis of collaborative experiments is the product of a collective socio-technical process that arises out of the scientists’ coordinated interactions.

3 Distributed Cognition and Virtue Reliabilism

The previous examples indicate that referring to collaborative knowledge as the knowledge possessed by a group of people as a whole is a *bona fide* phenomenon advocated within academic disciplines such as philosophy of science and cognitive psychology. They also demonstrate that the reason why collaborative knowledge is taken to be irreducibly social is that it arises out of the group members’ collaborative interactions.

To understand collaborative knowledge from the point of view of mainstream epistemology, the following sections build an account of the phenomenon by focusing on the distinctive feature of epistemic collaborations—i.e., the group members’ collaborative interactions. Arriving at such an account requires a hybrid approach that consists of the combination of virtue reliabilism with the hypothesis of distributed cognition. Section 3.1 lays out the main contours of the hypothesis of distributed cognition and Sect. 3.2 focuses on virtue reliabilism. Section 3.3 synthesises the two theories in a hybrid view in order to accommodate collaborative knowledge.

3.1 Distributed Cognition

Previously, we saw that Giere (2002b) has proposed to understand scientific experiments in terms of distributed cognition. Likewise, TMSs are exemplars of

distributed cognition within philosophy of mind and cognitive science (e.g., Sutton 2008; Sutton et al. 2010; Theiner 2013b).

According to the hypothesis of distributed cognition,⁶ when the members of a group collaboratively interact with each other, they give rise to novel, cognitive properties, over and above the sum of the properties possessed by its individual members. These are supposed to be the *collective* properties of a *distributed cognitive system* that consists of all the participating individuals. So, what is a distributed cognitive system and in what sense are its properties collective?

Within the literature, there are several ways to motivate the existence of distributed cognitive systems with collective properties.⁷ However, for the present purposes, I shall restrict the focus on what is perhaps the most succinct and informative approach on offer—one that relies on the mathematical field of Dynamical Systems Theory (DST).⁸ Though mathematically informed, the approach is fairly simple. It starts with DST's formal understanding of what a property is and it ends with capitalizing on DST's notion of a *coupled system*.

According to DST, properties are (actual or counterfactual) *regularities* in the behavior of a system as a whole.⁹ When a system functions in isolation, or it merely receives the effects of other systems, it is disposed to manifest a specific range of regular behaviors, as these are dictated by the system's background conditions and the equation governing its operation. These regularities of behavior are the system's properties. As it happens, however, when two (or more) systems engage in continuous, *reciprocal* interactions with each other—such that the effects of each system are continuously fed back to itself—*novel* properties arise. The reason these properties are supposed to be novel is because they refer to regularities of behavior that the contributing systems won't manifest either in isolation or aggregatively (i.e., when they are combined or added together, but in such a way that they do not continuously and reciprocally interact with each other).¹⁰ Properties, however, are always properties *of a given system*. Thus, to avoid leaving these regularities of behavior unaccounted for, it is necessary to postulate a corresponding *coupled* (i.e., distributed) system comprising of all the interacting (sub)systems.

⁶ See for example (Heylighen et al. 2004; Hutchins 1995; Sutton et al. 2010; Sutton 2008; Theiner et al. 2010; Theiner 2013a, b; Tollefsen and Dale 2012).

⁷ For an overview of the available argumentative approaches to the hypothesis of distributed cognition, see Theiner (2017).

⁸ DST is the branch of theoretical mathematics that deals with the study of complex systems that change over time—i.e., dynamical systems. It is widely used within nearly every scientific discipline (including physics, chemistry, biology and cognitive science) to model the behaviour of dynamical systems.

⁹ The qualification 'counterfactual' is necessary for accommodating the fact that no behaviour is going to be regular in the absence of the appropriate background and initial conditions. Only if the right conditions occur regularly, will the relevant properties also manifest regularly.

¹⁰ Aggregative properties are reducible to the sum of the properties of the contributing systems. Consider, for example, the property of the weight of a group. A group of five people will regularly weigh more than one person. However, this regularity of behaviour fails to qualify as a collective property, because it is due to the *additive* effect of individual-level properties (i.e., the masses of the individual members).

According to DST, then, when and only when systems engage in continuous reciprocal interactions with each other, do they give rise to an integrated, distributed system.¹¹ In practice, this means that when individuals interact loosely and in a largely unidirectional way—such as in cases of receiving testimony from a stranger or in the court of law—they do not, by the lights of DST, give rise to a distributed cognitive system. In such cases, while there may be a minimal amount of interpersonal communication involved, the speaker’s informational output is not the product of sustained reciprocal interactions with the hearer. On the contrary, in TMSs and collaborative scientific research teams, the completion of the relevant cognitive task involves ongoing reciprocal interactions between the participating individuals. Therefore, according to DST, in such cases we can talk of an overall distributed cognitive system that consists of all the participating individuals. The DST approach to distributed cognition thus provides a clear way for distinguishing between distributed cognition and cognition that is merely socially embedded. The difference is that only in the case of distributed cognition, does there exist an integrated system that arises out of the members’ *continuous and reciprocal interactions*.

By relying, in this way, on DST to draw a sharp distinction between groups of people that form genuine distributed cognitive systems and cases where two or more people engage in mere social interactions is a considerable advantage of the present approach to collaborative knowledge. If any group of individuals who socially interacted with each other were to count as a distributed cognitive system, then it would seem that there is something fundamentally wrong with the idea of collaborative knowledge as a special kind of group knowledge. Indeed, standard cases of testimonial knowledge would run the risk of qualifying as such. Of course, this is intuitively incorrect, and through its reliance on DST, the present account has the resources to explain why. As noted above, acquiring knowledge via testimony may involve *some* social interaction—perhaps the hearer asks a question and, depending on the answer they receive, they may also ask for some clarification. But the piece of information that constitutes the knowledge transmitted from the speaker to the hearer is not produced on the basis of *continuous reciprocal interactions* between the two.¹² Thus standard cases of testimonial exchange fail, by the lights of DST, to

¹¹ For a detailed explanation of why the existence of reciprocal interactions between agents and their artifacts ensures the existence of distributed cognitive systems, see (Palermos 2016a). To pre-empt a possible worry, the relevant reciprocal interactions need only be continuous during the operation of the distributed cognitive system and the unfolding of any processes related to it. For example, if, as part of her job and during normal working hours, individual *S* participates in distributed cognitive system *X*, *S* does not need to continuously interact with the other members of *X*, when she is at home. However, whenever *X* is in operation, *S* must continuously and reciprocally interact with the rest of the *X*-members.

¹² If the resulting piece of knowledge were the product of ongoing reciprocal interactions between individuals, as in cases of brainstorming, or, better, Transactive Memory Systems, then it would be a candidate for collaborative group knowledge. But I take it that when it comes to brainstorming one is more inclined to allow for this possibility, and as noted in Sect. 2.1, cognitive scientists already think of TMSs as “knowledge-acquiring, knowledge-holding and knowledge-using [systems] that [are] greater than the sum of [their] individual member systems” (Wegner et al. 1985, p. 256). Indeed, as an anonymous referee also notes, many long-term intimate relationships and close friendships would be plausible candidates for qualifying as TMSs capable to give rise to collaborative knowledge. So, although testimonial exchanges do not give rise to collaborative knowledge, epistemic collaborations in the form of TMSs and the collaborative knowledge they produce may be more common than what one may initially assume.

give rise to distributed cognitive systems and so, in agreement with intuition, would not, on the present approach, be candidates for giving rise to collaborative group knowledge either.¹³

Additionally, DST can help us clarify the sense in which certain properties are *collective* properties. First, we must not forget that, according to DST, properties are (actual or counterfactual) *regularities* in the behavior of a system as a whole. Though highly unlikely, this means that a token instance of a behavior—which were to be performed regularly would qualify as a distributed system's property—can, in principle, be *irregularly* performed by a random collection of individuals that falls short of forming an integrated group.¹⁴ According to DST, however, such irregular behavior won't count as the manifestation of a system's property. Consider the following example, with respect to the behavior involved in performing a collaborative scientific experiment.

When they work on their own, individual scientists are free to undertake a number of activities including interacting with equipment, performing calculations, taking breaks, writing papers, supervising students, refereeing articles and so on. Now, imagine that, one day, a few isolated scientists who work on their individual projects happen to perform a sequence of actions. Surprisingly, this sequence of actions is identical to a process that would routinely take place during a collaborative scientific experiment. Yet, intuitively, these scientists' activity does not qualify as performing a collaborative scientific experiment. The reason is that their actions, though identical to those of an experiment, are highly unlikely to occur in a consistent manner. On the contrary, actual experiments consist of sustained behaviors that are (or at least can be) regularly repeated. Accordingly, it would be incorrect to claim that the above scientists' one-off behavior can qualify as performing a collaborative scientific experiment.

The crucial question to ask, then, is what would it be required of those scientists to turn their fleeting behavior into regular behavior? The answer is a rather obvious one. To routinely perform the required sequence of actions, these scientists need to coordinate; which, in turn, requires that they engage in continuous reciprocal interactions with each other. Put another way, the behavioral regularities that constitute scientific experiments can only manifest in the presence of collaborative activity between the participants.

With that in mind, we can now see the sense in which certain properties are collective. As the example demonstrates, the regular behavior that underlies a scientific experiment cannot be manifested by any of the participating scientists alone. And neither can it be achieved by their mere aggregate—such an aggregate would at best

¹³ Indeed, elsewhere (Palermos 2016b), I have argued that testimonial knowledge can only motivate weak epistemic anti-individualism—i.e., the significantly weaker claim that, occasionally, knowledge can be both social and individual at the same time.

¹⁴ In connection with footnote 9, 'irregularly' should be here understood as 'irregularly even if the appropriate background and initial conditions are regularly in place.'

manifest the relevant behavior only fleetingly. Accordingly, performing the experiment is a property that cannot be attributed to either the underlying individuals or their sum. As the regularity of behavior that it is, however, it must still be accounted for as the property of *some* system. So, what might this system be? Given that only the scientists' collaborative interactions can support the relevant regularity of behavior, this is, by the lights of DST, the property of the *entire* distributed cognitive system, consisting of all the collaborating participants. In other words, such regularities of behavior that arise out of group members' continuous and reciprocal interactions are *collective* (as opposed to individual or aggregative) properties, because they can only be attributed to a corresponding group entity.

3.2 Virtue Reliabilism

It is widely accepted that in order to know, one must have a true belief with positive epistemic standing. Moreover, broadly speaking, epistemologists assume that positive epistemic standing consists in the agent's true belief being (1) produced by a process that is de facto reliable and/or (2) formed in a way that is epistemically responsible.

Traditionally, the second condition of epistemic responsibility has been approached in terms of justification, where justification is understood as the ability to provide positive explicit reasons in support of one's beliefs or in support of the reliability of one's beliefs, by reflection alone.¹⁵ While this is an intuitive way to think about epistemic responsibility, a significant worry concerns several belief-forming processes, such as our perceptual faculties and memory. Such belief-forming processes are generally held to be knowledge-conducive, but most epistemic agents do not know how they work or why they are reliable. Accordingly, when we acquire knowledge on their basis, it seems incorrect to require that we be in a position to offer, by reflection alone, explicit positive reasons in their support.¹⁶

To solve this problem, process reliabilists have suggested that we stop thinking about epistemic responsibility in terms of justification, and perhaps abandon the requirement of epistemically responsible true beliefs altogether.¹⁷ Instead, positive epistemic standing consists only in arriving at true beliefs on the basis of processes that are de facto reliable. According to process reliabilism, to know, one only needs to form a true belief on the basis of a belief-forming process that is in fact reliable—no matter whether one is aware that this is the case. This is widely accepted to be a promising strategy for sidestepping the above problem for epistemic responsibility when it is understood in terms of justification. Nevertheless, it invites a host of problems of its own.

¹⁵ This is known as the access internalist approach to justification. For classic defenses of the view see (BonJour 1985; Chisholm 1977; Steup 1999).

¹⁶ For more on this problem for access-internalism, see Greco's (1999) discussion, which draws on Hume's problem of induction.

¹⁷ The locus classicus of process reliabilism is (Goldman 1979). For an overview see (Goldman & Beddor 2016).

First, process reliabilism does not guarantee the right direction of fit between truth and belief. One may reliably form true beliefs on the basis of a process, not because the relevant process gets the world correct. Instead, there might be some secret scheme in play (say a benevolent demon) that manipulates the world so as to correspond to the agent's randomly formed beliefs, as soon as they are formed.¹⁸ Secondly, process reliabilism fails to rule correct in cases where true beliefs are generated on the basis of reliable, but fleeting (or strange) processes. Intuitively, despite being reliable, unstable and abnormal processes should not count as knowledge-conducive.¹⁹ Finally, by focusing solely on de facto reliability, process reliabilism misses an important dimension of our epistemic nature. While it is true that in order to know we need the way of forming our beliefs to be in fact reliable, intuitively, this is not sufficient on its own. What we further need is that we be epistemically responsible in the way we form our beliefs by somehow being sensitive to the reliability of our evidence. While process reliabilists are correct to point out, for the reasons mentioned above, that this cannot always be achieved on the basis of justification (i.e., by having reflective access to positive explicit reasons in support of our beliefs), it seems incorrect to ignore this dimension of our epistemic nature altogether.

To solve these problems for process reliabilism, virtue reliabilists have proposed to supplement process reliabilism with the *ability intuition on knowledge*.²⁰ According to this intuition, in order for one's true beliefs to qualify as knowledge, they must be the product of a belief-forming process that counts as a cognitive ability. This is an appealing move for three reasons. First, if the relevant process is a cognitive ability, the direction of fit between truth and belief is guaranteed to be the correct one. Secondly, cognitive abilities are stable intellectual dispositions that are neither strange nor fleeting. Finally, cognitive abilities seem to be, at least intuitively, the sort of reliable processes that can generate beliefs in an epistemically responsible manner, even if one has no explicit positive reasons to offer in their support. For example, no one needs to explain why their vision or hearing is reliable when they come to acquire knowledge on their basis.

Nevertheless, if this is the way to approach knowledge and epistemic responsibility, two central questions present themselves: (1) When does a process count as a cognitive ability (and thereby as knowledge-conducive); and (2) in what sense can one generate epistemically responsible beliefs on the basis of their cognitive abilities, without being in a position to offer explicit reasons in their support? In answering these two important questions, virtue reliabilists turn to the concept of cognitive integration.

¹⁸ For such a case see Pritchard's (2009) Truetemp.

¹⁹ For an example of a fleeting but reliable process that should fail to count as knowledge-conducive, see Greco's (2010) Careless Math Student. For examples of strange but reliable processes that should also fail to count as knowledge-conducive, see Bonjour's (1980) case of clairvoyance and Greco's (2010) Serendipitous Brain Lesion.

²⁰ The idea that knowledge must be grounded in cognitive abilities can be traced back to the writings of (Sosa 1988, 1993) and Plantinga (1993). For more recent approaches to this intuition, see Greco (1999, 2004, 2007) and Pritchard (2009, 2010a, b, 2012).

According to Greco, in order for a reliable process to count as a cognitive ability (and thereby as knowledge-conducive) it must have been cognitively integrated, where “cognitive integration is a function of cooperation and interaction, or cooperative interaction with other aspects of the cognitive system” (2010, p. 152). The answer to the first question, then, is that a belief-forming process counts as a cognitive ability if and only if it is cognitively integrated on the basis of processes of mutual interactions with other aspects of the cognitive system.²¹

Answering the first question in this way provides the means to respond to the second question too; conceiving of cognitive abilities in terms of cognitive integration can explain how one is epistemically responsible in forming beliefs on their basis, without being able to offer reasons in their support (Palermos 2014). If an agent’s belief-forming process cooperatively interacts with other aspects of their cognitive system, then it can be continuously monitored, in the background, by their cognitive system, such that *if* there is something wrong with it, *then* the agent will be able to notice this and respond appropriately. Otherwise, if the agent’s background monitoring generates no negative beliefs against their belief-forming process, they can be epistemically responsible in employing the relevant process and accepting its results *by default*—even if they have no reasons to offer in their support.

On virtue reliabilism, then, provided that one’s belief-forming process is cognitively integrated by cooperatively interacting with other aspects of their cognitive system, they can be epistemically responsible in holding the resulting beliefs merely by lacking any negative reasons against them. If this cognitively integrated belief-forming process is also *de facto* reliable, then the agent can entertain beliefs with full positive epistemic standing.

3.3 Virtue Reliabilism and Distributed Cognition

With the preceding points in place, we can now attempt a hybrid account of collaborative knowledge that consists in the combination of virtue reliabilism with the hypothesis of distributed cognition. The cornerstone of this hybrid approach is the concept of cognitive integration. Even though independently developed within distinct philosophical disciplines, both theories ascribe central role to cognitive integration, which they both understand in terms of continuous reciprocal interactions between different parts of the cognitive system. Just as proponents of distributed cognition claim that a group of people is cognitively integrated when its individual members engage in reciprocal interactions, so virtue reliabilism claims that cognitive integration of a belief-forming process is a matter of cooperative interactions with other parts of the cognitive system. Moreover, virtue reliabilism holds that a

²¹ Empirically, this is not an implausible demand. Scientific observations of the human and of many primates’ brain indicate that brain areas hardly ever operate in isolation. Instead, several parts of the brain always work in parallel. Take the case of reaching for and grasping a mug sitting on the desk before you. Even this simple cognitive task requires the delicate use of visually received information by functionally and neuro-anatomically distinct subsystems operating together. Similarly, it is widely accepted that there is constant, bidirectional relationship between our basic perceptual processes, memory systems and other inference mechanisms.

cognitively integrated process qualifies as a cognitive ability that is suitable for generating knowledge. This suggests that a knowledge-conducive cognitive ability can be distributed across an integrated group of collaborating individuals.

Nevertheless, to adequately motivate this claim, we need to understand how the integrated nature of epistemic collaborations contributes to their positive epistemic standing; in what way is their cognitively integrated nature involved in the production of reliable and epistemically responsible beliefs? To answer, we need to unpack the way the members' interactions give rise to two epistemic group properties: The group's (1) *epistemic self-organization* and (2) *epistemic self-regulation*.

Self-organization is a diachronic process, during which members of the collaborative group mutually interact until they evolve into a stable configuration. Normally, when the group has achieved this stable configuration, it means that its component parts have mutually adapted by restricting their interactions to those that allow them to regularly accomplish their end.²² In the case of epistemic collaborations, self-organisation contributes, over time, to the reliability of the collective belief-forming process: If the dynamics of the group do not produce a preponderance of true over false beliefs, the collective does not regularly accomplish its end; in result, it keeps reconfiguring into different structures until it stabilizes into a successful (i.e., sufficiently reliable) one. Alternatively, if this is not possible, eventually, the collective will (most likely) dissolve.²³ So, overall, by mutually interacting, group members participate in a process of self-organisation, which is conducive to shaping a group structure that can reliably produce true beliefs. It is this self-organized structure that Knorr-Cetina draws attention to when she notes that HEP experiments are not managed and organized by any individual alone. Instead, within HEP experiments, the participants' interactions create a grid of discourse spaces, which "was and is today perhaps the most important vehicle of experimental coordination and integration" (1999, p. 174).

Mutual interactions between individual members also explains how collaborative groups self-regulate synchronically, at the time of performance, to produce epistemically responsible true beliefs. The continuous interactivity between the members of the group allows them to keep monitoring each other's performance, such that if there is something wrong with the overall process, then the problem will be spotted by at least one member of the group, allowing the group to respond appropriately. For example, if the members of a TMS are sufficiently coordinated, they will be able to notice whether there is a problem with their exchange. For instance,

²² For more on this process of self-organisation, see (Heylighen et al. 2004).

²³ In cases where the group cannot reconfigure into a sufficiently reliable stable structure, it will *most likely*, but not necessarily dissolve for the following reason. It is possible to imagine that an epistemic collaboration would persist (i.e., manage to preserve itself) despite continuous failures. For example, though rather implausibly, there could be a group of people that have the financial capacity and drive to keep their epistemic collaboration going, despite consistently failing to deliver a preponderance of true over false beliefs. Therefore, the mere existence of a self-organising epistemic collaboration does not guarantee that, eventually, either it will dissolve or it will become reliable (or reliable enough for epistemic purposes)—though, in most cases, these are the most likely outcomes. Be that as it may, the claim is that *if* an epistemic collaboration establishes a *stable* structure that is reliable enough to deliver knowledge, this is going to be (at least to a significant extent) due to its capacity to self-organise.

they may notice that their partner's non-verbal or para-verbal cues suggest that they misremember. Similarly, in a scientific experiment, if the experimentalist sends in data with weak peaks, the mathematician will notice this and will call the physicist's attention to it. Such synchronic monitoring processes are essentially driven by the individuals' coordinated activity. If on their basis, no member expresses negative reasons against the overall process, as it unfolds over time, the group can count as epistemically responsible for its performance by default—even if no one in the group, or even the group as a whole, fully understands how its performance works.

Overall then, the ongoing interactions between the members of epistemic collaborations contribute to their positive epistemic standing in both a diachronic and a synchronic manner. During the development of the group, the ongoing mutual interactions of the members allows them to self-organize. This process of self-organization contributes to shaping the overall system into a structure that can reliably bring about true beliefs. Then, on the basis of this self-organized structure, the members of the group efficiently interact with each other, providing the group with the ability to self-regulate during the actual time of performance.²⁴ The upshot is that the epistemic collaboration's overall ability to generate beliefs with full positive epistemic standing is not the product of merely adding together the cognitive abilities of the members of the group. Instead, it arises out of their synergetic cooperation. According to DST, this means that the epistemic collaboration's positive epistemic standing is a collective property that belongs to the group as a whole.

An interesting consequence follows from this. We have here reached the conclusion that, in the case of epistemic collaborations, positive epistemic standing is the collective property of a group of individuals, constituting an integrated distributed cognitive system; since the collective property under consideration is the property of *positive epistemic standing*, the corresponding distributed cognitive system may qualify as an *epistemic group agent*. The following section elaborates on this point.

4 Distributed Virtue Reliabilism

To sum up, in cases of epistemic collaborations, agents continuously and reciprocally interact with each other. The result is a group that epistemically self-organizes and self-regulates so as to form a belief-forming process that can reliably and responsibly generate true beliefs—i.e., beliefs with positive epistemic standing. DST points out that this positive epistemic standing is a collective property, because it cannot be accounted for without appealing to an overall distributed cognitive system, consisting of all the collaborating individuals. At the end of last section, it was noted that, since the epistemic collaboration's collective property is the property of

²⁴ It must be noted, however, that the distinction between diachronic reliability and synchronic responsibility, as well as the corresponding distinction between the underlying mechanisms of self-organization and self-regulation, are largely theoretical artifacts. In practice, every instance of epistemic self-regulation contributes to the process of epistemic self-organization and vice versa: The diachronic structure of the group is continuously shaped by and at the same time shapes its ongoing performance.

positive epistemic standing, there are good reasons for thinking that the underlying distributed system qualifies as an epistemic group agent.

To unpack this point further, it is helpful to contrast the present view with an alternative account of collaborative knowledge. Drawing on process reliabilism, rather than virtue reliabilism, Michaelian and Arango-Muñoz (2018) have put forward *distributed process reliabilism* in order to account for knowledge that is specifically produced by TMSs.

The common ground between the present approach and distributed process reliabilism is that they both hold that reliable belief-forming processes can be distributed between several agents and their artifacts. The main point of departure is that the two theories are motivated on the basis of different approaches to knowledge. Even though Michaelian and Arango-Muñoz concede that rejecting virtue reliabilism and the associated ability intuition on knowledge is a theoretical sacrifice, they opt to develop their approach in terms of process reliabilism. Their methodological choice is due to the fact that virtue reliabilism requires a rich notion of group agency that entails epistemic responsibility. This is problematic, according to Michaelian and Arango-Muñoz, because they doubt that collaborative groups such as TMSs are likely to manifest epistemic responsibility. To the contrary, process reliabilism does not require that the agent be epistemically responsible, but only *de facto* reliable in the way they form their beliefs. Accordingly, they hold, contrary to process reliabilism, virtue reliabilism is not in a good position to account for collaborative knowledge, because of its strong demands on cognitive agency. In their words:

[Virtue reliabilism requires] what we might refer to as responsible cognitive agency, where responsible cognitive agency requires cognitive agency, plus responsibility. There is no clear sense in which a TMS, for example, might be assigned responsibility for its cognitive success and failures, and it is in this sense we have suggested that extended and distributed memory systems do not qualify as cognitive agents: they may be cognitive agents simpliciter, but they are not responsible cognitive agents (idid., 242).

Thus, Michaelian and Arango-Muñoz distinguish between two notions of cognitive agency, one of which is richer, because it also involves responsibility. In this richer sense, they claim, “distributed memory systems do not qualify as cognitive agents.” While this may initially sound unproblematic, it leads to a rather negative consequence of their view. If a group does not qualify as a cognitive agent, then it cannot qualify as an epistemic group agent either. Especially when the missing property—i.e., epistemic responsibility—is one of the two components of positive epistemic standing.

Contrary to the above, the present approach demonstrates that epistemic collaborations can qualify as both reliable *and* epistemically responsible cognitive agents. This is possible on the present view, because of the way it captures the ability intuition on knowledge, which Michaelian and Arango-Muñoz abandon. According to virtue reliabilism, a process counts as a cognitive ability, capable of generating epistemically responsible beliefs, if and only if it is cognitively integrated by cooperatively interacting with other aspects of the cognitive system. In the case of epistemic collaborations, group members continuously interact with each other in order to

generate an integrated distributed belief-forming process. The continuous interactions between the members of the group allow them to continuously monitor each other's performance. In result, if there is something wrong with their distributed process, the group will be alerted to it and respond appropriately. When no alert signals are communicated, the group can responsibly accept the deliverances of its distributed cognitive ability by default. This sense of epistemic responsibility arises out of the members' collaborative interactions and belongs to the group as a whole.

Therefore, in contrast to distributed process reliabilism, the present approach—*distributed virtue reliabilism*—has the resources to understand groups as epistemically responsible cognitive agents, and thus as epistemic group agents with full positive epistemic standing (i.e., as agents that can be both reliable and epistemically responsible). Furthermore, it does so by focusing on the group members' reciprocal interactions, which is epistemic collaborations' distinctive feature.

Being in a position to account for epistemic responsibility in the case of collaborative knowledge provides distributed virtue reliabilism with three advantages. First, it preserves the ability intuition on knowledge, even in the case of epistemic collaborations. While process reliabilists may wish to undermine its theoretical import, in Sect. 3.2, it was noted that the ability intuition on knowledge is well motivated, precisely because it can help address several problems with process reliabilism.

Secondly, focusing on epistemic responsibility can help clarify the sense in which epistemic collaborations can motivate strong epistemic anti-individualism in a distinctive way. The present approach does not claim that collaborative knowledge is irreducibly social, because the relevant proposition is collectively believed (or 'accepted'). As noted in the introductory section, this is an alternative approach to group knowledge (see Gilbert (2007a, b, 2010), Rolin (2008), Tuomela (2004)) that is not necessary to the present approach. Independently of whether the relevant proposition may count, in some appropriate sense, as collectively believed, it can still qualify as irreducibly collective knowledge. According to distributed virtue reliabilism, the resulting knowledge is irreducibly social, because the positive epistemic standing of the relevant belief consists in its being the product of a responsible and reliable process that is irreducibly social.²⁵

Finally, accounting for epistemic responsibility consolidates the claim that epistemic collaborations qualify as epistemic group agents, which is crucial for avoiding the idea of knowledge without a knowing subject. This is a theoretical move that has been recently advocated by Giere with respect to knowledge generated by scientific research teams. Even though Giere accentuates the distributed nature of the scientific process that produces knowledge, he denies the existence of a distributed agent that knows. Giere suggests that, in such cases, we should opt for an "epistemology

²⁵ Of course, group knowledge, just as any other type of knowledge, involves belief in the proposition known, and the relevant belief must, on some appropriate construal, qualify as the belief of the group—if only in an instrumentalist sense. Nevertheless, whether the relevant belief (or acceptance) counts as the belief of the group, because it is of the summative or non-summative type, the belief of some operative members of the group or merely the belief of a single representative (for an overview, see Tollefsen 2004) is an issue we do not here need to take a stance on. On the present account, any of these possibilities with respect to group belief may give rise to collaborative group knowledge.

without a knowing subject” [(Giere 2007), quoting (Popper 1968)], according to which, we may only claim that “it has been scientifically established that *p*” (Giere 2002b) or that ‘it is scientifically known that *p*.’

Initially, this might sound like a viable option. The problem, however, is that it significantly departs from mainstream epistemology, which has always assumed that knowledge is knowledge of a subject *S*. The present account can help us resist Giere’s (and Popper’s) impersonal image of collaborative knowledge. While distributed virtue reliabilism denies that collaboratively produced knowledge belongs to any particular individual subject, *S*, it insists that it still is knowledge of a subject, *G*—the epistemic group agent, whose positive epistemic standing results from its members’ continuous and reciprocal interactions.²⁶

5 Conclusion

The paper introduced the hybrid approach of distributed virtue reliabilism, in order to account for the special case of group knowledge that is produced on the basis of epistemic collaborations. Collaborative knowledge is resistant to aggregative analyses of group knowledge, because of the defining feature of epistemic collaborations—i.e., the mutual interactions of their members. The present account focused on this special feature of epistemic collaborations in order to explain how the resulting true beliefs can entertain positive epistemic standing. It was argued that the continuous reciprocal interactions between the members of epistemic collaborations contribute both to the reliability and epistemic responsibility of the resulting

²⁶ This section compared distributed virtue reliabilism to distributed process reliabilism. To my knowledge, this is the only alternative account of knowledge produced by epistemic collaborations that combines distributed cognition with a mainstream approach to knowledge (i.e., process reliabilism). Bird (2010, 2014) has also argued that scientific knowledge is a form of group knowledge, by appealing to the notion of distributed cognition. I do not, however, consider Bird’s account to be an alternative to the present approach for the following two reasons. First, Bird is interested in developing an account of ‘social knowing’ as it applies to ‘wider science’, where the epistemic subject is not a well-defined group of collaborating individuals. Rather on Bird’s view, the epistemic subject is the whole scientific (international) community, which might partly comprise of individuals that are not involved in the production of knowledge (some of them may simply be consumers of it, for example). Secondly, Bird does not cash out his view in terms of any of the existing mainstream accounts of knowledge. Therefore, both Bird’s explanatory target and his approach are significantly different to the present ones. A further difference concerns our understanding of distributed cognition and thus what may count as part of a distributed cognitive system. Bird interprets the idea of distributed cognition by invoking Durkheim’s notion of organic solidarity. Following this line, Bird concludes that, in the case of scientific knowledge, the epistemic group agent can consist of individuals that only indirectly (i.e., non-epistemically) contribute to or benefit from a scientific claim. On the contrary, the present approach understands distributed cognition in terms of DST. According to DST, in order for something (be it an artefact or an individual) to count as a constitutive part of the epistemic group agent, it must continuously and mutually interact with the group in the way the group produces knowledge. While I am sympathetic to Bird’s account and his idea of ‘social knowing’ as it might apply to ‘wider science’, I worry that his underlying understanding of distributed cognition is vulnerable to the ‘causal-constitution’ fallacy and the ‘cognitive bloat’ worry as these are widely discussed in the literature on extended and distributed cognition. On the contrary, the present approach can satisfactorily deal with these worries due to its reliance on the notion of a coupled system from DST (for more details on these issues, see (Palermos 2016a)).

beliefs. This collaboratively generated positive epistemic standing is unique to epistemic collaborations and amounts to a collective property that belongs to the corresponding group as a whole. Thus, on distributed virtue reliabilism, knowledge produced by epistemic collaborations qualifies as a special—albeit widespread—kind of group knowledge that lends support to strong epistemic anti-individualism in a distinctive way.

Acknowledgements For their feedback on the ideas presented in this paper, I am thankful to Andy Clark, Sandy Goldberg, Jesper Kallestrup, Duncan Pritchard and Mark Sprevak. I am also thankful to the audiences of the research seminars organized by the Cogito Group in Glasgow and the Epistemology Research Group in Edinburgh, as well as the audience of the Epistemic Subjects Beyond Individuals workshop held at Vrije University in Brussels. I would also like to thank two anonymous referees for *Erkenntnis*.

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

- Barnier, A. J., Sutton, J., Harris, C. B., & Wilson, R. A. (2008). A conceptual and empirical framework for the social distribution of cognition: The case of memory. *Cognitive Systems Research*, 9(1–2), 33–51. <https://doi.org/10.1016/j.cogsys.2007.07.002>.
- Bird, A. (2010). Social knowing: The social sense of scientific knowledge. *Philosophical Perspectives*, 24(1), 23–56. <https://doi.org/10.1111/j.1520-8583.2010.00184.x>.
- Bird, A. (2014). When is there a group that knows? Distributed cognition, scientific knowledge, and the social epistemic subject. *Essays in Collective Epistemology*. <https://doi.org/10.1093/acprof:oso/9780199665792.003.0003>.
- BonJour, L. (1980). Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, 5, 53–73.
- BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge, MA: Harvard University Press.
- Chisholm, R. M. (1977). *Theory of Knowledge* (2nd ed.). Englewood Cliffs, NJ: Prentice-Hall.
- Cooke, N. J., Gorman, J. C., Myers, C. W., & Duran, J. L. (2013). Interactive team cognition. *Cognitive science*, 37(2), 255–285.
- de Ridder, J. (2014). Epistemic dependence and collective scientific knowledge. *Synthese*, 191(1), 37–53.
- Duarte, R., Araújo, D., Correia, V., & Davids, K. (2012). Sports teams as superorganisms. *Sports Medicine*, 42(8), 633–642.
- Duarte, R., Araújo, D., Correia, V., Davids, K., Marques, P., & Richardson, M. J. (2013a). Competing together: Assessing the dynamics of team–team and player–team synchrony in professional association football. *Human Movement Science*, 32(4), 555–566.
- Duarte, R., Araújo, D., Folgado, H., Esteves, P., Marques, P., & Davids, K. (2013b). Capturing complex, non-linear team behaviours during competitive football performance. *Journal of Systems Science and Complexity*, 26(1), 62–72.
- Giere, R. N. (2002a). Discussion note: Distributed cognition in epistemic cultures. *Philosophy of Science*, 69(4), 637–644.
- Giere, R. N. (2002b). Scientific cognition as distributed cognition. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *Cognitive bases of science* (pp. 285–299). Cambridge: Cambridge University Press.
- Giere, R. N. (2006). The role of agency in distributed cognitive systems. *Philosophy of Science*, 73(5), 710–719.

- Giere, R. N. (2007). Distributed cognition without distributed knowing. *Social Epistemology*, 21(3), 313–320.
- Giere, R. N., & Moffatt, B. (2003). Distributed cognition: Where the cognitive and the social merge. *Social Studies of Science*, 33(2), 301–310. <https://doi.org/10.1177/03063127030332017>.
- Gilbert, M. P. (1994). Remarks on collective belief. In F. Schmitt (Ed.), *Socializing epistemology: The social dimensions of knowledge* (pp. 235–256). Lanham: Rowman & Littlefield.
- Gilbert, M. P. (2007a). *Collective epistemology*. Cambridge: Cambridge University Press.
- Gilbert, M. P. (2007b). Modeling collective belief. *Synthese*, 73, 185–204.
- Gilbert, M. P. (2010). Belief and acceptance as features of groups. *Protosociology: An International Journal of Interdisciplinary Research*, 16, 35–69.
- Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: Springer.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Clarendon Press.
- Goldman, A. I. (2004). Group knowledge versus group rationality: Two approaches to social epistemology. *Episteme*, 1(01), 11–22. <https://doi.org/10.3366/epi.2004.1.1.11>.
- Goldman, A. I. (2014). Social process reliabilism. In J. Lackey (Ed.), *Essays in collective epistemology*. Oxford: Oxford University Press.
- Goldman, A., & Beddor, B. (2016). Reliabilist Epistemology. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. Oxford: Oxford University Press.
- Greco, J. (1999). Agent reliabilism. In J. Tomberlin (Ed.), *Philosophical perspectives 13: Epistemology* (pp. 273–296). Atascadero, CA: Ridgeview Press.
- Greco, J. (2004). Knowledge as credit for true belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology* (pp. 111–134). Oxford: Oxford University Press.
- Greco, J. (2007). The nature of ability and the purpose of knowledge. *Philosophical Issues*, 17, 57–69.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
- Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy*, 82(7), 335–349.
- Heylighen, F., Heath, M., & Van, F. (2004). The emergence of distributed cognition: A conceptual framework. In *Proceedings of collective intentionality IV*.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge: MIT Press.
- Knorr, K. C. (1999). *Epistemic cultures: How the sciences make knowledge*. Cambridge, Mass: Harvard University Press.
- Lackey, J. (2014a). Socially extended knowledge. *Philosophical Issues*, 24, 282–298.
- Lackey, J. (2014b). A deflationary Account of Group Testimony. In J. Lackey (Ed.), *Essays in collective epistemology* (pp. 65–95). Oxford: Oxford University Press.
- Lackey, J. (2016). What is justified group belief? *Philosophical Review*, 125(3), 341–396.
- Latour, B. (1999). *Pandora's hope: An essay on the reality of science studies*. Cambridge: Harvard University Press.
- List, C. (2005). Group knowledge and group rationality: A judgment aggregation perspective. *Episteme*, 2(01), 25–38.
- List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford: Oxford University Press.
- Michaelian, K., & Arango-Muñoz, S. (2018). Collaborative memory knowledge: A distributed reliabilist perspective. In M. Meade, A. Barnier, P. Van Bergen, C. Harris, & J. Sutton (Eds.), *Collaborative remembering: Theories, research, and applications* (pp. 231–247). Oxford: Oxford University Press.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge/MA: Harvard University Press.
- Palermos, S. O. (2014). Knowledge and cognitive integration. *Synthese*, 191(8), 1931–1951.
- Palermos, S. O. (2016a). The dynamics of group cognition. *Minds and Machines*, 26(4), 409–440.
- Palermos, S. O. (2016b). Spreading the credit: Virtue reliabilism and weak epistemic anti-individualism. *Erkenntnis*, 81(2), 305–334.
- Plantinga, A. (1993). *Warrant and proper function*. New York: Oxford University Press.
- Popper, K. R. (1968). Epistemology without a knowing subject. In B. van Rootselaar & J. F. Staal (Eds.), *In Proceedings of the third international congress for logic, methodology and philosophy of science*. Amsterdam: North-Holland Publishing Company. (Reprinted from *Objective knowledge: An evolutionary approach*, by K. R. Popper, 1972. Oxford: Oxford University Press).
- Pritchard, D. H. (2009). *Knowledge*. London: Palgrave Macmillan.

- Pritchard, D. (2010a). Knowledge and understanding. In A. Haddock, A. Millar, & D. H. Pritchard (Eds.), *The nature and value of knowledge: Three investigations* (pp. 3–90). Oxford: Oxford University Press.
- Pritchard, D. (2010b). Cognitive ability and the extended cognition thesis. *Synthese*, 175, 133–151.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Rolin, K. (2008). Science as collective knowledge. *Cognitive Systems Research*, 9(1–2), 115–124. <https://doi.org/10.1016/j.cogsys.2007.07.007>.
- Sosa, E. (1988). Beyond Skepticism, to the Best of our Knowledge. *Mind, New Series*, 97(386), 153–188.
- Sosa, E. (1993). Proper Functionalism and Virtue Epistemology. *Nous*, 27(1), 51–65.
- Steup, M. (1999). A defense of internalism. In L. Pojman (Eds.), *The Theory of Knowledge: Classical and Contemporary Readings* (3rd ed., pp. 310–321). Belmont, CA: Wadsworth.
- Sutton, J. (2008). Between individual and collective memory: Interaction, coordination, distribution. *Social Research*, 75(1), 23–48.
- Sutton, J., Harris, C. B., Keil, P. G., & Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, 9(4), 521–560. <https://doi.org/10.1007/s11097-010-9182-y>.
- Theiner, G. (2013a). Onwards and upwards with the extended mind: From individual to collective epistemic action. In L. Caporael, J. Griesemer, & W. Wimsatt (Eds.), *Developing scaffolds* (pp. 191–208). MIT Press: Cambridge.
- Theiner, G. (2013b). Transactive memory systems: A mechanistic analysis of emergent group memory. *Review of Philosophy and Psychology*, 4(1), 65–89. <https://doi.org/10.1007/s13164-012-0128-x>.
- Theiner, G. (2017). Group-sized distributed cognitive systems. In M. Jankovic & K. Ludwig (Eds.), *The Routledge handbook of collective intentionality*. London: Routledge.
- Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research*, 11(4), 378–395. <https://doi.org/10.1016/j.cogsys.2010.07.002>.
- Tollefsen, D. P. (2004). Collective intentionality. In: J. Fieser, B. Dowden, (eds) *Internet encyclopedia of philosophy*.
- Tollefsen, D. (2006). Group deliberation, social cohesion, and scientific teamwork: Is there room for dissent? *Episteme A Journal of Social Epistemology*, 3(1), 37–51. <https://doi.org/10.1353/epi.0.0008>.
- Tollefsen, D. (2015). *Groups as agents*. Cambridge: Polity Press.
- Tollefsen, D., & Dale, R. (2012). Naturalizing joint action: A process-based approach. *Philosophical Psychology*, 25(3), 385–407. <https://doi.org/10.1080/09515089.2011.579418>.
- Tuomela, R. (2004). Group knowledge analyzed. *Episteme*, 1(2), 109–127.
- Wegner, M., Giuliano, T., & Hertel, P. (1985). Cognitive interdependence in close relationships. In W. J. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253–276). New York: Springer.
- Wray, K. B. (2007). Who has scientific knowledge? *Social Epistemology*, 21(3), 337–347. <https://doi.org/10.1080/02691720701674288>.

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