

The subject of knowledge in collaborative science

Duygu Uygun Tunç

Abstract

The epistemic subject of collective scientific knowledge has been a matter of dispute in recent philosophy of science and epistemology. Following the distributed cognition framework, both collective-subject accounts (most notably by Knorr-Cetina, in *Epistemic Cultures*, Harvard University Press, 1999) as well as no-subject accounts of collective scientific knowledge (most notably by Giere, *Social Epistemology* 21:313–320, 2007; in Carruthers, Stich, Siegal (eds), *The Cognitive Basis of Science*, Cambridge University Press, 2002a) have been offered. Both strategies of accounting for collective knowledge are problematic from the perspective of mainstream epistemology. Postulating genuinely collective epistemic subjects is a high-commitment strategy with little clear benefits. On the other hand, eliminating the epistemic subject radically severs the link between knowledge and knowers. Most importantly, both strategies lead to the undesirable outcome that in some cases of scientific knowledge there might be no individual knower that we can identify. I argue that distributed cognition offers us a fertile framework for analyzing complex socio-technical processes of contemporary scientific knowledge production, but scientific knowledge should nonetheless be located in individual knowers. I distinguish between the production and possession of knowledge, and argue that collective knowledge is collectively produced knowledge, not collectively possessed knowledge. I propose an account of non-testimonial, expert scientific knowledge which allows for collectively produced knowledge to be known by individuals.

Introduction

Large research collaborations constitute an increasingly prevalent form of social organization of research activity in many scientific fields. In the last decades, the concept of distributed cognition has provided a suitable basis for thinking about collective knowledge in the philosophy of science. Karin Knorr-Cetina's and Ronald Giere's analyses

of high energy physics experiments are the most prominent examples. Although they both refer to distributed cognition in describing the processes of knowledge production in these experiments, their accounts regarding the epistemic subject of knowledge thus produced are quite different. While Knorr-Cetina argues for an irreducibly collective subject, Giere argues for eliminating the epistemic subject and opting for using the passive voice in describing collectively produced knowledge. Neither of these views are easy to assimilate within a mainstream epistemological account. The collective subject view postulates a new, supra-individual epistemic subject and denies knowledge to individuals when the processes of knowledge production are distributed. Both moves create tension with epistemology's traditionally individualist framework. The no-subject view envisions that we can divorce knowledge and knowledge-production from knowers, which clashes both with the intuitive assumption that knowledge implies a knower and the traditional association between knowledge and intellectual autonomy. Both views entail that in dealing with the phenomenon of distributed cognition we can do without individual knowers.

I will argue that epistemology should be extended in a way that can accommodate collectively produced knowledge, but that we will have a serious problem if this involves denying scientific knowledge to individuals. If the members of a large research collaboration cannot be said to know the collectively produced epistemic outcomes, we would have to accept the absurd conclusion that either no one or only a supra-individual entity learns from the most successful research collaborations we have. I will both advance skepticism against the collective subject view and counter the skepticism towards individual knowledge in the context of distributed cognition. To this aim, I will argue for conceiving research collaborations in terms of a cognitive system that produces (not possesses) knowledge, which can eventually be possessed (though not produced) by constituent individuals when certain conditions are met. Firstly, the distributed research process should be reliable in producing scientific evidence and secondly, there should be a reliable distributed process of criticism for scrutinizing the reliability of the scientific evidence that is collectively produced. I will analyze both conditions in terms of distributed first-order and second-order justification, where I put forward a virtue reliabilist account of justification that is compatible with epistemic dependence. I will

conclude that the notion of justified epistemic dependence enables us to attribute knowledge to individuals when knowledge production is irreducibly social.

The structure of the paper is as follows. In Sect [1](#) I briefly outline how the distributed cognition framework has been applied to collaborative research and the radically divergent accounts of collective knowledge it inspired; namely, the collective subject account and the no-subject account. In Sect [2](#) I elaborate on the problems both the collective-subject account and the no-subject account face. In [2.1](#) I examine several examples of the strongly anti-individualist perspective on collective scientific knowledge which treats groups as genuine epistemic subjects. In [2.2](#) I examine the eliminativist strategy which opts for impersonal or subjectless knowledge. In Sect [3](#) I present my original account of collectively produced, individually possessed scientific knowledge. In [3.1](#) I characterize research collaborations as distributed cognitive systems for producing objective knowledge (i.e., a system of scientific propositions). In [3.2](#) I present my account of knowledge that allows for attributing non-testimonial, expert scientific knowledge of collectively produced epistemic outcomes to individual scientists. In Sect [4](#) I conclude the paper with some thoughts on potential further directions that can be taken by a social epistemology of science that does not investigate collective epistemic subjects.

1. Distributed cognition model of collaborative research and the elusive subject of knowledge

Scientific inquiry is at bottom a highly structured cognitive process. Cognitive processes are generally thought to occur exclusively within organismic boundaries, so as a cognitive process scientific inquiry is intuitively something that happens in the head of the individual scientist. But we rarely find that such a complex form of cognition as scientific inquiry is realized without substantial reliance on scientific instruments and other experts, past and present. Various kinds of factors external to the individual agent seem to play not only supportive but constitutive roles in the production of scientific knowledge. Such epistemic dependence comes into full relief in large research collaborations, where individual agents coordinate their diverse expertise, cognitive effort and interactions with various epistemic artifacts in ways that give rise to what we may call complex cognitive systems. Research collaborations are formed to realize overly complex cognitive tasks, or “big questions,” that typically surpass the bounds of individual expertise and cognitive

capacity, thus can be said to produce knowledge at the supra-individual or epistemic system level.

The concept of distributed cognition, which originated in cognitive science, is grounded in the non-individualist or externalist premise that cognition is not necessarily an intracranial process but can extend to external epistemic sources such as scientific instruments as well as incorporate the cognitive activities of multiple agents (Hutchins, 1995; also, Clark, 1996; Clark & Chalmers, 1998). Distributed cognition provides a useful framework for analyzing collective knowledge production in terms of division of cognitive labor, and it has already been employed in the philosophy of science to describe collaborative research processes in certain fields. On the basis of his observations at the Indiana University Cyclotron Facility, Ronald Giere (2002a) describes the collaborative research activity thus:

In thinking about this facility, one might be tempted to ask, who is gathering the data? From the standpoint of distributed cognition, that is a poorly framed question. A better description of the situation is to say that the data is being gathered by a complex cognitive system consisting of the accelerator, detectors, computers and all the people actively working on the experiment. Understanding such a complex cognitive system requires more than just enumerating the components. It requires also understanding the organization of the components. And [...] this includes the social organization.

Giere (2002b) also provides a more general description of distributed cognition (which he does not intend as a definition): We speak of distributed cognition where two or more individuals reach a cognitive outcome by combining un-shared individual knowledge and by interacting with epistemic artifacts. Karin Knorr-Cetina (1999) similarly depicts the High Energy Physics experiments she observed during her field research stay at CERN in terms of distributed cognition:

The point is that no single individual or small group of individuals can, by themselves, produce the kind of results these experiments are after—for example, vector bosons or the long “elusive” top quark or the Higgs mechanism. It is this impossibility which the authorship conventions of experimental HEP exhibit. They signify that the individual has been turned into an element of a much larger unit that functions as a collective epistemic subject (p. 167-8).

...reflexivity is turned into an instrument of knowledge, machines are redefined and recruited into the social world, and the subjectivity of participants is put on the line – and quite successfully replaced by something like distributed cognition (p. 25).

Although distributed cognition presents a particularly useful model for examining the epistemic structure of collaborative science, it also raises serious doubts about whether we can still conceive scientific knowledge as a state of the traditional subject of epistemology – the individual. While Giere and Knorr-Cetina offer similar descriptions of how knowledge is *produced* in collaborative experiments in terms of distributed cognition, their accounts differ significantly when it comes to identifying the *epistemic subject* of collectively produced knowledge.

For Knorr-Cetina, the epistemic subject in the case of HEP experiments is the experiment itself. The whole collaboration, together with the instruments it employs and all the communicative and practical activities and interactions that weave the people and the instruments into a unitary entity, presents a novel epistemic subject:

The HEP experiments studied, in continually integrating over themselves (to put it in mathematical terms), continually assemble the collaboration into a community reflexively bound together through self-knowledge. The medium that brings this assemblage about is the conversation a collaboration holds with itself. This conversation, I maintain, replaces the individual epistemic subject, which is so prominent in other fields. It construes, and accounts for, a new kind of epistemic subject, a procurer of knowledge that is collective and dispersed. No individual knows it all, but within the experiment's conversation with itself, knowledge is produced (*Op. Cit.*, p. 178).

For Knorr-Cetina the subjectivity of the individual subject is erased, and through distributed cognition the experiment not only becomes a supra-individual entity (e.g., a system) but an epistemic subject tout court, as it acquires “a stream of (collective) *self-knowledge*” (p. 171-173), “a sort of consciousness” (p. 178).¹

¹In her portrayal even the instruments become organismic entities by virtue of the way in which researchers interact with them and are integrated into an organismic whole that is the experiment – which she models along the lines of Durkheimian collective consciousness. The above quoted paragraph continues: “For those

Giere (2002b, 2007), on the other hand, finds such an ascription of collective subjectivity to research collaborations too much of an ontological commitment.² He argues that we can view certain research collaborations as distributed “cognitive systems” because they realize a cognitive task,³ not because they exhibit as a whole cognitive properties that imply agency. Thus, we do not need to postulate distributed cognitive agents in order to speak about distributed cognitive systems. In particular we do not need to endow such systems with mental states such as knowledge or (its prerequisite) belief. Giere maintains, instead, that we should characterize them in a depersonalized or impersonal way, “so that we would say things like ‘This experiment has shown that. . . .’ or ‘This experiment leads to the conclusion that. . . .’” He envisions that the developing science of cognition could allow us to redefine cognition as a technical scientific concept (which does not correlate with mindedness or agency) rather than a folk-psychological one, and to leave behind the assumption that “if knowledge is being produced, there must be an epistemic subject, the thing that knows what comes to be known” (2007, p. 316).

2. Why both no-subject and the collective-subject accounts of scientific knowledge are problematic

Both the strategy of conceiving collective knowledge in a non-subjective or impersonal way and that of postulating collective epistemic subjects conflict with the individualistic perspective of traditional epistemology, according to which knowledge is a cognitive/epistemic state of the individual. The collective subject account of collaborative scientific knowledge is premised on the idea that the subject of knowledge should be whomever that produces it, which in this case eliminates the individual as a candidate. It postulates a novel kind of epistemic subject, the research collaboration or the experiment,⁴ in its stead. The no-subject account maintains, on the other hand, that knowledge production does not necessarily imply subjecthood and in the particular case of large

who still remember Durkheim (1933: chap. 3), the conversation produces a version of his much-rebuffed ‘conscience collective’.”

²Kitcher (1994) and Thagard (1997) similarly argue against the view that knowledge can be possessed by a collective subject.

³For taking the “task” as the individuating factor for distributed cognitive systems, see also Magnus, 2007.

⁴Knorr-Cetina’s supra-individual subject, the experiment, comprises not only the human members of a research team but also the technical instruments the research team relies on in producing knowledge.

research collaborations such an implication would be mistaken – which again disqualifies the individual along with any other candidate entity.⁵

As both Knorr-Cetina and Giere admit, it is clear that some modern forms of scientific inquiry substantially challenge some of our core epistemic intuitions, starting with our traditional individualism about knowledge. Distributed cognition provides us with a framework in which we can reconsider this core individualistic assumption of epistemology and talk about collective epistemic states and achievements, as it is increasingly being done in social epistemology (E.g., Gilbert, 1987; 2004; Goldman, 2014; List and Petit, 2011; Tuomela, 1992; 2004). I maintain, however, that this extension or revision of traditional epistemology (cf. Palermos and Pritchard, 2013) should not go as far as postulating collective epistemic subjects or endorsing an exclusively impersonal view of knowledge. Both these strategies are problematic, as they entail the possibility that no one, strictly speaking, learns from the most successful research collaborations we have. However, as I will argue in the third section, collective production of scientific knowledge does not present us with a forced choice between these two. What we need to acknowledge is only that some epistemic processes whereby individuals come to acquire knowledge can require for their realization complex cognitive systems that comprise other agents and possibly epistemic artifacts.

2.1. Irreducibly collective knowledge

The collective-subject account is problematic primarily due to the unnecessarily high degree of ontological commitment it has to make. Firstly, research collaborations do not seem *prima facie* to manifest subjective properties such as consciousness, reflectivity, care or self-knowledge. Knorr-Cetina attributes the HEP experiments precisely such subjective properties but does so without putting forward an explicit ontological argument that would warrant such an attribution. In order to warrant the postulation of collective subjects, one has to demonstrate that collective accomplishment of a cognitive task entails a collective mind. To put this in terms of distributed cognition, one must show at least

⁵ It may be argued neither account necessarily denies individual knowledge but only argues for the possibility of knowledge without an individual knower. Such a weaker reading would not make a significant difference for the present argument.

how distributed cognition implies distributed mental states. Such an account must go beyond joint epistemic actions and argue for irreducibly collective mental properties.⁶

For justification of such an inference from distributed cognition to irreducibly collective (or social) epistemic subjects, we can turn to other accounts that similarly advocate high-commitment positions. More recently Alexander Bird (2010, 2014) and Orestis Palermos (2020) argued for genuinely or irreducibly collective scientific knowledge. Bird (2014), like Knorr-Cetina, invokes Durkheim's concept of "organic solidarity" in grounding distributed cognitive systems as genuine epistemic subjects.⁷ Scientists in a research collaboration, for Bird, compose a genuine social entity on the basis of their mutual interdependence due to the division of scientific labor, which implies a distribution of cognitive sub-tasks not merely in a quantitative but also qualitative manner (i.e., in accordance with the heterogeneity of the expertise required). Bird then goes from division of scientific labor to irreducibly collective epistemic states via a functionalist argument: The collective entity realizes a cognitive function, which consists in cognitive activity geared towards a certain goal, and we can explain a cognitive function the best by attributing intentional states to the target system. The system as a whole can be said to have a cognitive/epistemic state on the basis of accomplishing a cognitive function even if no individual member of the system is in that state. Thus, the group can have scientific knowledge that all individual members lack. Bird (2010) gives the example of an imaginary research team consisting of a physicist and a mathematician, where one establishes that if p then q and the other the truth of p without interacting with one another and the conclusion that q is published by an assistant per pre-arrangement. Bird argues that in this case the *research team alone* knows that q .⁸ However, it is not clear what exactly binds the two researchers into a research *team*, as nothing would change in the example if one or both were dead. In line with the lack of any principle of individuation for epistemic groups, Bird does not restrict this account to distributed cognitive systems with clearly

⁶There are other accounts of collective epistemic states which do not make the ontological commitment in the second step, such as the joint commitment or acceptance accounts of group belief by Raimo Tuomela (1992, 1995, 2004) and Margaret Gilbert (1987, 2004, 2014).

⁷For a similar invocation of Durkheim in regard to group knowledge, see Wray, 2007.

⁸Actually, this thought example can testify to the no-subject or impersonal knowledge account much better than it does to the collective-subject account, as there is hardly any reason to presume a collective belief that q .

defined tasks. He extends it to wider science on the basis of the epistemic interdependence of the scientific community, calling it a single entity.

A core concern here is obviously that Bird's account is not able to differentiate between unified cognitive systems and loosely organized epistemic communities (see also Wray, 2007). For this reason, the framework of distributed cognition loses its conceptual role in accounting for collective knowledge. Epistemic interdependence in the broad sense can be said to characterize all human epistemic endeavors and we can clearly not speak of an epistemic subject who is absolutely autonomous in producing knowledge. In this regard, he is not in a position even to delineate an actively interacting epistemic community from its long past contributors, since findings, theories and inventions live much longer than their originators. This directly leads to the worry that the subject of scientific knowledge is inflated to the point of meaninglessness.⁹

Palermos (2020) offers a similarly strong definition of distributed cognitive systems, which nonetheless delineates distributed cognitive systems from broader communities of knowledge. His account draws on Dynamic Systems Theory and can be summarized as follows:

Emergent dynamic system view of distributed cognition: There is a distributed cognitive system if and only if *continuous and reciprocal interactions* between constituent members give rise to an *integrated system with novel, non-aggregative properties*.

For Palermos, collective knowledge that arises in such a distributed cognitive system is a special kind of group knowledge, one that is not summative. Palermos' argument contains the premise that the emergent system is an irreducible group agent, which can also be seen as a group mind. The reason is that emergent distributed cognitive systems exhibit, for Palermos, socio-cognitive properties that do not belong to any individual member (2016b).¹⁰ Palermos' account (2020) is clearly free from the kind of inflation of the

⁹A similar objection directed at the extended (or distributed) cognition thesis is known as the "cognitive bloat" (see e.g., Rupert, 2004). I am not concerned with this argument in this paper, since I assume that distributed cognitive systems can be meaningfully individuated although I argue against attributing them subjective or agentive states.

¹⁰Against the possible objection that the attribution of a mind implies attribution of consciousness, which groups lack, Palermos (2016b) states that consciousness may not be necessary for mindedness. In particular, he considers it plausible that groups manifest specific cognitive processes such as memory, decision-making and knowing. See n.1.

epistemic subject, since his criterion of inclusion is continuous and reciprocal interaction. This criterion, for Palermos, applies to distributed cognitive systems in the same way it does to individual (biological) cognitive systems. Individual cognitive systems are characterized by cooperative interactions between the (functionally parsed) constituent parts and sub-parts of the system (e.g., memory, motor control). Distributed cognitive systems are organized through the coupling of multiple cognitive systems through continuous and reciprocal interactions and by virtue of functional equivalence they also deserve the status of cognitive systems. Further, in case distributed systems can accomplish the same cognitive functions as biological systems, such as decision-making or belief-formation, the resulting cognitive/epistemic states are those of the entire system in a non-summativ, irreducible sense. Thus, they can be manifest at the system level even if no constituent member manifests them.

Being developed within a performance-based virtue framework originating from Ernest Sosa (2007, 2009, 2011), Palermos' argument proceeds from collective *performances* to emergent collective *properties*. He argues that the positive epistemic standing¹¹ of the beliefs produced by research collaborations is an irreducibly collective property. A similar conclusion is advanced by Jesper Kallestrup (2016) and Adam Carter (2020), who also proposed virtue epistemological perspectives on collective knowledge. Both Kallestrup and Carter argue for the irreducibility of collective knowledge on the basis of the irreducibility of a particular epistemic property of group belief; namely, its *aptness*. Aptness, Sosa's criterion for epistemic virtue, refers to epistemic success that manifests competence. An apt belief is a belief that is true because it is formed via the exercise of a reliable cognitive skill or ability. Similarly to Palermos, Kallestrup and Carter argue that a group can form an apt belief while no individual member of the group can.¹²

I think one can convincingly argue that distributed cognitive systems have weakly emergent collective properties, which do not compel us to invoke collective subjective states. In the case of research collaborations, the "reliability" or the "efficiency" of the distributed research process in yielding credible empirical evidence are such weakly emergent properties which cannot be obtained by simply adding up the corresponding

¹¹ This positive epistemic standing of group beliefs consists, according to Palermos, in their being reliable and epistemically responsible.

¹² Using Carter's (2020) formulation, this group of virtue epistemological accounts of collective knowledge endorse a symmetrical conception of aptness, which they apply equally to individual and group apt belief.

properties of constituent sub-processes with disregard to the organizational structure of the system. Similarly, the required complex “expertise” or “competence” for implementing the collectively agreed research design, data collection and analysis strategy, and for the manipulation and coordination of the heterogeneous set of scientific instruments is a property of the research collaboration as a whole. It is perfectly possible and oftentimes true that no member of a research collaboration individually manifests this complex competence manifested at the system level. Moreover, the system-level competence can comprise certain “skills” that no member of the research collaboration exercises; namely, those that are due the scientific instruments which function as epistemic artifacts in extending (or replacing) human cognitive capabilities (see e.g., Palermos, 2011). However, the collective epistemic competence of a research collaboration is not an irreducible property. There is nothing in this complex competence that cannot be analyzed in terms of constituent skills and the way in which they are organized and coordinated (cf. Carter, 2020;¹³ Pino, 2021¹⁴). As Kallestrup similarly maintains, “a group’s innermost competence is reducible to a summation of innermost competences of its individual members and their manner of arrangement within the group,” hence, “novel competences of groups do not spring into existence or mysteriously emerge when conjoining existing individual ones” (2016, p. 10).¹⁵ No single member of a research collaboration manifests

¹³ Carter maintains that not only the aptness of group beliefs but also group competences are genuinely or irreducibly collective. His argument is based on the case of so-called Mandevillian intelligence; i.e., “some dispositions that are unreliable and thereby are not individual-level competences can [...] lead to knowledge conducive dispositions at the collective level” (2020, p.24). There is one minor and one major problem with this argument. The minor one is that the case for Mandevillian intelligence (see Smart, 2018) lacks convincing real-world examples of collaborative research and rests largely on computer simulations. For this reason, its external validity remains to be explored. The major one is that the concept, assuming that it has some external validity, is applicable to loosely organized scientific communities rather than research collaborations. Research collaborations have clear cognitive goals and research strategies to achieve them, both of which are often set in advance. Thus, they do not give the individual researchers sufficient elbow room to engage in “deviant” epistemic behaviors which would be incompatible with the collective aims. We can see this more clearly if we take some applications of the idea to scientific inquiry, such as Zollman’s (2010), who argues that individuals who are intellectually dogmatic can possibly bring about epistemic benefits at the collective level by exploring areas of the “epistemic landscape” that would be left uncharted by those who are motivated by “truth.” While such independent exploration can take place in a loosely organized scientific community such as a research field, it is unrealistic to expect it in a research collaboration.

¹⁴ Pino (2021) defends that group epistemic competence should be seen as an irreducible property, as “group normative status that guides towards knowledge.” However, his argument rests on an anti-intellectual conception of group know-how. Research collaborations are typically much more fitted for an intellectualist interpretation, according to which a “group is guided by explicit norms that result from the member’s joint acceptance.”

¹⁵ Kallestrup conjoins this deflationist account of group competence, however, with a non-deflationist account of group knowledge on the basis of group apt belief, which I reject as I see belief as an unfit category

fully the complex competence that we see at the level of the distributed cognitive system, but the latter consists in a particular organization of various individual competences.

These and similar weakly emergent properties of research collaborations can be among the determinants of whether accepting a scientific proposition asserted by the collaboration counts as knowledge. As Sosa acknowledges, the epistemic success of an individual belief might be due to a “complex social competence only partially seated in that individual believer” (2007, p. 97). The epistemic status of the individual scientist’s belief in the findings established by a research collaboration depends to a substantial extent on various weakly emergent properties of the distributed cognitive system such as the complex competence it manifests, as I will explicate further in the following sections.

However, it is not clear what in particular would be gained by attributing *epistemic subjecthood*, let alone *beliefs* to research collaborations in any literal sense, which would imply strongly emergent properties such as collective intentionality or collective epistemic responsibility¹⁶ (if not collective consciousness). The distributed cognitive process realized by a collaboration is primarily one of establishing scientific evidence for a (set of) proposition(s) by implementing a methodological plan; *it is not a process of belief-formation*.¹⁷ It is of secondary importance that these statements are also affirmed by a particular collective body. Here I must further note that when we take the research collaborations at CERN as real-world and prominent examples of a distributed research process, we see that while they can be said to accept and assert propositions, it is much more difficult to characterize them as having group beliefs – except maybe in a summative

to describe the collective epistemic output of research collaborations. I elaborate on this in the rest of the paper.

¹⁶ Unlike subjecthood, I have no objection to the application of the concept of collective *accountability* to research collaborations. Collectives can sensibly be treated as agents for specific purposes, such as attributing due credit or blame when they make assertions in the form of publications or public announcements, similar to the category of juridical persons in law. But it is important to note that when we treat groups as agents who are subject to praise and blame, a strongly social view of knowledge is particularly problematic. As Lackey (2014) maintains, when a group knows that p without any individual member knowing that p , then it will be epistemically irrational for the group to act on p . Since a public assertion (e.g., publication) is an action, even if we grant strongly emergent epistemic states such as belief or knowledge that p to a research collaboration, it would be epistemically irrational for the collaboration to assert that p – which would undermine its purpose.

¹⁷ I should note that Palermos (2020) does not premise his argument on the metaphysics of collective belief. His focus is on the justificatory properties of processes of belief-formation rather than beliefs themselves. Nonetheless, he treats distributed research processes as processes of belief-formation, which I think is not necessary. See also my next note.

sense, which does not imply genuinely collective mental states.¹⁸ Beliefs are generally considered to be automatic or involuntary (Hakli, 2006). However, in these and similar cases in science the goal (e.g., investigating whether p) is typically determined beforehand in panels and committees, and the results do not automatically turn into epistemic outputs as the collaborations still have to decide whether to make the findings public or examine them further. Moreover, a collaboration may always revise the conclusions drawn on the basis of the findings or decide to retract a publication upon finding a mistake in the calculations. Thus, while the epistemic output of a research collaboration, typically a publication or public announcement, can be regarded as involving a propositional attitude for certain intends and purposes, such as attributing due credit or blame, such propositional attitudes are better characterized as acceptances and assertions rather than beliefs (see also Hakli, 2007; Wray, 2001). Consequently, the distributed cognition framework is pertinent for conceiving how research collaborations realize the cognitive task of evidence-generation, but we do not have sufficient reason to regard collective research processes as processes of belief-formation (cf. Carter, 2020).

The chief complication for any strongly anti-individualist view of collective knowledge such as the foregoing is obviously the close association between epistemic states and subjective states. An alternative strategy may be to distance epistemological concerns from psychological ones and give terms like “believes” a novel, non-mentalistic interpretation which need not imply consciousness.¹⁹ Such a move would be permissible of course, but it is not so clear why it should be desirable. Ethics of terminology would require that there should be a clear gain from changing the ordinary meanings of terms to render them applicable to a broader class of objects. I do not think that this is the case with group knowledge, as it can be conceived in a way that “saves the phenomenon” to the same extent without undertaking a substantial redefinition of ordinary epistemological and psychological terms. Moreover, there may be good reasons to treat group or system

¹⁸ A research collaboration can be said to summatively believe that p when most or all collaboration members believe that p (see Quinton, 1976). The summative interpretation of group belief may still be difficult to apply to most research collaborations, because it is perfectly possible that a collaboration refrains from publicly affirming that p (for instance due to high epistemic or other kinds of risks involved) while all collaboration members believe that p . A notion of group acceptance can thus be more suitable in most cases. For accounts of group acceptance, see Hakli (2007), Wray (2001).

¹⁹ Similarly to how cognition has largely been de-psychologized in turning into a cognitive scientific concept. There are also widely acclaimed proposals for de-psychologizing the concept of mind. See “extended mind,” Clark & Chalmers, 1998.

level epistemic properties as being qualitatively distinct from individual ones, because there are properties that can be exhibited by social organizations and not by individuals (no matter with or without subjective states) and vice versa. I will come back to this point in the last section of the paper.

Most importantly, besides their costly ontological commitment to collective epistemic subjects, these and similar accounts explicitly acknowledge the possibility of a scenario where we can rightly attribute knowledge of a scientific discovery to literally no scientist.²⁰ For instance, Palermos (2020) states that his “distributed virtue reliabilism denies that collaboratively produced knowledge belongs to any particular individual subject.” Similarly, Bird (2014) maintains that “there can be scientific knowledge without any individual knowing.” This is a highly undesirable conclusion, as it entails that knowledge does not supervene on individual cognitive states.²¹ It rests, I think, partly on a conflation of collective processes and their properties with the outcomes of such processes. Sometimes a task consists merely in a “performance,” but in many other cases there is an output distinct from the performance that brought it about. Let us think of Hutchins’s example of ship navigation, through which he greatly popularized the concept of distributed cognitive system.²² A typical task on a ship can be bringing the ship to a dry dock, the outcome of which is only that the ship has been dry-docked. It is accomplished by a system, where instruments and people co-constitute a vast network of mutual computational and representational dependencies, as Hutchins describes. The task is massively distributed, such that we can point to no one who indeed docks the ship. A sub-task such as determining the relative position of the ship vis-a-vis the dock, however, has a specific output: the calculated relative position of the ship. While the task of determining it is a genuinely collective cognitive effort, the position of the ship can be known in principle by anyone. In this regard, collaborations ultimately produce empirical evidence for

²⁰ Palermos’ motivation for postulating epistemic group agents, as he explicitly states, is to avoid an impersonal construal of collective knowledge as proposed by Giere (2020, pp. 117). As I argue in the next section, I agree with him that such a position “significantly departs from mainstream epistemology, which has always assumed that knowledge is knowledge of a subject S.” However, his proposal (or any other similarly anti-individualist proposal) equally departs from mainstream epistemology for its rejection of its individualist framework) and is not necessary to avoid the impersonal knowledge view. Moreover, I believe that “knowledge without an individual subject” does not fare much better in terms of desirability than “knowledge without a subject.”

²¹ For further epistemological problems with this conclusion, see Lackey, 2014.

²² I have to note that Hutchins himself is more sympathetic to the idea of a distributed mind than I am.

scientific propositions, and I doubt that it is an appealing conclusion to say that some scientific propositions are not known by anyone but a supra-individual entity.

2.2 Impersonal knowledge

To turn to the no-subject account, we can admit that conceiving scientific knowledge as impersonal knowledge, or knowledge without a subject has some conceptual advantages and a certain appeal. Scientific knowledge, arguably unlike mundane knowledge-that and knowledge-how, is at a fundamental level a system of statements that are interwoven via logical operations and methodological rules. In this respect, scientific knowledge can be regarded as “objective knowledge” in Popper’s sense (1968) in contradistinction to “subjective knowledge” which is a mental phenomenon – specifically, a form of belief:²³

knowledge or thought in the subjective sense, consisting of a state of mind or of consciousness or a disposition to behave or to react, and knowledge in an objective sense, consisting of problems, theories, and arguments as such. Knowledge in this objective sense is totally independent of anybody’s claim to know; also it is independent of anybody’s belief, or disposition to assent; or to assert, or to act. Knowledge in the objective sense is knowledge without a knower: it is knowledge without a knowing subject.

Although Giere does not specify what he means by impersonal knowledge beyond suggesting that we reformulate knowledge attribution statements in passive form, his account can lend itself to be interpreted in a way quite similar to Popper’s notion of objective knowledge (see esp. Giere, 2007). However, the concept of objective knowledge lacks any reference to acts of thinking and practices of inquiry. For this reason, it does not tell us by itself anything about the processes of scientific knowledge production, which establish the empirical justification for the targeted system of statements, or where this kind of knowledge resides – in individual minds, groups of minds, or in books, articles, databases? It merely refers to the *outcome* of an epistemic process, which in turn can be regarded as mental content as well as a material system of external signs. Thus, the concept of objective knowledge does not imply any commitment to any epistemic subject either in its production or its possession. Consequently, we still have to ask the question of what exactly is collective in collective scientific knowledge, to which we can in principle give two answers: We can say that it is collectively *produced* knowledge or that it is

²³ For a similar point, see Faulkner, 2018. Faulkner further remarks that Popper’s distinction between objective and subjective knowledge can be likened to the distinction between the justification of a proposition and the justification of a belief, which does not have the ontological commitment to a third world of intelligibles as does Popper’s.

collectively *possessed* knowledge (or both). The way Giere analyzes research collaborations through the concept of distributed cognition leads us to the first option: Research collaborations produce objective knowledge (e.g., a scientific finding) by realizing *collectively* the complex cognitive processes that are required for its establishment, where these processes involve combining various kinds of background knowledge (i.e., expertise), interacting with various scientific instruments (i.e., epistemic artifacts), and organizing various cognitive activities into a coherent procedure (e.g., analyzing data, drawing inferences, comparing calculations).

Collective production of knowledge (through distributed cognition) is also a feature of Knorr-Cetina's, Bird's and Palermos' analyses. The core difference between these two perspectives is how they answer the question as to the epistemic subject of the knowledge thus produced. This question addresses, as I have said, the *seat* of knowledge. For Giere we do not need to answer this question; we do not have to assume an epistemic subject that knows "what comes to be known" (i.e., objective knowledge). For others, the subject that knows is "the experiment," "the scientific community," or "the collaboration" – an irreducibly collective subject.

While scientific knowledge is in one respect clearly objective knowledge, which can "reside" in systems of material, external signs (e.g., printed in books), it would be a far-fetched conclusion to say that it can reside *solely* in this manner. Can we say that it will still be *known* that the universe is expanding even if the world enters another dark age, and nobody is left who understands physical cosmology? It is reasonable to say, with Popper (1968), that the following two scenarios would not be the same: There is no living person who has sufficient knowledge in physical cosmology, but (i) all scientific publications are preserved in libraries, or (ii) all scientific publications are also destroyed. In the first case it is highly probable that one day some people who will have trained themselves in physical cosmology using the materials in the libraries will read the relevant publications and be able to learn that the universe is expanding. Nevertheless, we can say without much hesitation that until that happens it would not be *known* that the universe is expanding. Thus, it is difficult to say that objective knowledge can exist without furnishing the content of subjective knowledge.

The no-subject account of collectively produced knowledge leads us, just like the collective subject account, to the absurd possibility that nobody comes to know what is established in some of the most successful cases of scientific research, such as the empirical confirmation of the Higgs boson. I think a much more commonsensical position is to say that objective knowledge implies subjective knowledge. Tuomela (2004) also hints at such an implication by saying that “such knowledge is not an abstract entity floating around in some kind of Platonic ‘third world’. Rather it is knowledge that some actual agent or agents actually have or have had as contents of their appropriate mental states.” Thus, we should be able to say that research collaborations produce knowledge in a distributed manner, but it is the individual scientists that come to know the outcomes of the distributed cognitive process. Giere actually has a suggestion in a similar direction, though he does not specify it in a way that would satisfy the epistemologist. He remarks that it is the individual scientists who evaluate the outcomes and draw conclusions on the basis of the experiments, and indirectly the lay person through (a chain of) testimony. He maintains that while this kind of knowledge cannot be produced by individuals, the final result can be known by individuals in the ordinary sense of the term (2002b, p. 643). Although Giere’s suggestion is completely in line with one of my two main points, namely that individuals can come to know collectively produced knowledge, he takes lightly the challenge posed by distributed knowledge production for epistemology. Giere does not offer any specification for how the scientists can be said to know the final results of collaborative experiments, and testimonial knowledge (which he ascribes to all others who learn about the results) is already not part of the challenge. The individual scientist does not come to know collectively produced knowledge in any “ordinary” sense of the term. Let me first explicate the challenge in more detail, and then provide a suggestion for how to meet it.

The traditional epistemological concept of knowledge, despite all variety in its analysis, is that of subjective knowledge: a mental phenomenon and more specifically a particularly valued form of belief. It is generally the qualities of the belief-forming process that raises it to the level of knowledge, in addition to the qualities of the belief’s content. From a virtue reliabilist perspective, for instance, a true proposition or a system of true propositions is not knowledge; it is the belief in a true proposition (or a system thereof) that is formed via the exercise of a reliable cognitive competence. From an internalist

perspective, it is a true belief which is supported by consciously available good reasons. In any case, the processes whereby knowledge is produced cannot be divorced from it, as they are the source of its justification. But this is exactly what happens in distributed cognitive systems: The agentic constituents of the system might come to entertain true beliefs by affirming the outcomes (if the distributed process is successful in yielding true propositions), but they are typically not sufficiently justified in doing so. The problem with distributed processes of scientific justification for the epistemologist stems thus from the fact that the traditional individualistic view of knowledge involves epistemic autonomy: Epistemic subjects can be said to know if they are solely or primarily responsible in the production of this knowledge.²⁴

When we admit that objective knowledge implies subjective knowledge, the traditional individualism of epistemology leads us directly to a problem in the case of distributed cognition, which is often referred as Hardwig's dilemma: we either have to postulate a collective epistemic subject who solely has the justification (i.e., scientific evidence) for accepting a system of propositions (i.e., a scientific claim), or we have to provide an account of how the individual scientist can be said to know without having the justification to do so (See Hardwig, 1985, p. 348-9). In either case we ironically end up going radically against the individualist premise (by denying either the individuality of the epistemic subject or the requirement for epistemic autonomy). Many authors, including Hardwig, have opted for taking the first horn of this dilemma and offered increasingly robust accounts of collective knowledge (e.g., Carter, 2020; de Ridder, 2014; Faulkner, 2018; Kallestrup, 2016; Palermos, 2020; Pino, 2021). I think exploring the second (in my opinion more conservative) option is a better strategy in accounting for collectively produced, individually possessed knowledge. Hardwig was motivated to avoid the second option, which he called "vicarious" knowing, as he wanted to save the intuition that "knowing a proposition requires understanding the proposition and possessing the relevant evidence for its truth" (1985, p. 349). I propose a more nuanced account which allows that individuals can have *sufficient* justification *non-autonomously*, which grounds my position that scientific knowledge that can be collectively produced and individually known. Thus, we can indeed reject both horns of Hardwig's dilemma.

²⁴See also Palermos, 2016a. Palermos formulates epistemic autonomy in terms of autonomous possession of justification.

3. A third way: Collectively produced, individually known

The most parsimonious and plausible way to save both subjective knowledge of scientific propositions and the premise that the proper epistemic subject is the individual goes through reconsidering the requirement for epistemic autonomy and updating our view of knowledge to accommodate epistemic dependence. We can then be in a position to formulate an alternative account of collective scientific knowledge by conceiving research collaborations as distributed cognitive systems that produce (not possess) knowledge (section 3.1), which can eventually be possessed (though not produced) by constituent individuals when certain conditions are met (section 3.2).

3.1 Research collaborations as distributed cognitive systems for production of objective knowledge

In research collaborations the “output” is not a collective mental state such as belief but a system of scientific propositions (i.e., a scientific claim) which stand in inferential relations to the reported data given the documented methodological procedures.²⁵ We can more particularly say that the distributed cognitive process is one of evidence-generation in support of collectively made assertions. Thus, as far as we see the product as “knowledge,” it is not knowledge in the subjective sense but only in the objective, non-mental sense.²⁶

The construal of a research collaboration as a “cognitive” system means, in line with Giere, that it is a socio-technological system of various activities that serve the fulfillment of a cognitive task. A significant portion of these activities are also cognitive in nature, while the rest can be primarily practical, social or instrumental. The implementation of a research plan through distributed cognition in research collaborations does not compel us beyond this to postulate distributed minds or subjects, because, as I

²⁵ Cf. Vaesen (2011), who argues that if one rejects distributed knowledge, then the only way to apply the distributed cognition framework to science is endorse a deflationary view and regard the proper “output” of distributed cognitive systems in science as information or data. However, this alternative ignores that research collaborations typically put forward a system of scientific statements, not merely the empirical results of the research. On this basis, they make collective assertions (e.g., scientific publications) for which they can be held accountable. If, on the other hand, a distributed cognitive system was to merely produce data, we would have little reason to regard it as anything else than a giant socio-technical data collector.

²⁶ According to Palermos (2020), in the case of epistemic collaborations, the collective cognitive property is the resulting beliefs’ positive epistemic standing. But we do not have to accept that “positive epistemic standing” implies a collective agent, since it is not even a cognitive property. For instance, a high “degree of corroboration” of a scientific claim can ground positive epistemic standing, although it is an objective, formal property.

argued, the research process as a whole is not a mental, agentic or subjective activity like belief formation, but a process of knowledge production in the objective sense or, still more clearly, one of evidence-generation.

The process of evidence-generation, whether in individual research or in a research collaboration, can be analyzed in terms of its two central objectives: (i) producing evidence E for scientific claim p , and (ii) producing higher-order evidence E' for E . Any successful process of scientific inquiry produces not only first-order evidence that justifies the acceptance of a scientific claim, but also some satisfactory amount²⁷ of higher-order evidence that the first-order evidence is genuine (i.e., it is not a fluke or an artifact of the research procedure) and is not defeated. Higher-order evidence, or evidence of evidence, is necessary for establishing an evidential connection between the findings and a scientific claim, and it can range from error or uncertainty estimation, validation of measurement tools, testing of alternative hypotheses, investigation of potential confounding factors to analysis of coherence with background knowledge and other well-established theories. In research collaborations the objectives (i) and (ii) are typically met in a distributed manner.

A research collaboration implements a complex research plan that requires the effective coordination of various research activities that are globally geared towards a unitary goal, such as establishing evidence in support of a scientific theory. These activities or sub-tasks typically require diverse expertise, simultaneous manipulation of multiple scientific instruments, or data collection at different times and places. Thus, the evidence towards the truth of a scientific proposition is established in a distributed manner. We can call the process whereby this first-order evidence is produced the *distributed research process*. It is distributed, since producing such complex scientific evidence exceeds the cognitive ability and capacity of individual researchers and requires a distributed cognitive system. As an integral part of the research plan, a research collaboration also engages in various activities for scrutinizing the evidential value of first-order scientific evidence, whereby it produces higher-order evidence. In a distributed cognitive system, higher-order evidence is typically also generated in a distributed manner, where different collaboration members provide diverse kinds of higher-order evidence in accordance with their expertise. The activities of higher-order evidence-generation constitute collectively a

²⁷ I.e., to a degree that meets the evidential standards of the scientific community or the research field.

distributed higher-order regulative mechanism, which we can call the *distributed process of criticism*. The distributed research process and the distributed process of criticism together constitute the *epistemic performance* of the research collaboration. Thus, the epistemic performance of a research collaboration is *distributed*, while this distributed performance has a distinct *output*; namely, first-order and higher-order evidence in support of a scientific claim. This output constitutes what we referred as objective knowledge.

3.2 Individual collaboration members as the proper subjects of knowledge

Knowledge simpliciter is subjective knowledge, since objective knowledge can be said to exist only as its content. I have argued so far that research collaborations cannot be said to possess knowledge. In the following I will argue that individuals can be said to possess collectively produced knowledge.

Following Sosa's (2007, 2009, 2010, 2011, 2015) twofold distinction between animal and reflective knowledge, we can conceive scientific knowledge of the expert as a species of reflective knowledge; that is, a case of knowledge which not only implies that one reaches true beliefs through the exercise of reliable cognitive skills or abilities (i.e., epistemic competence), but also that one has a positive judgment regarding the reliability of the skills or abilities in question. In order to attain reflective knowledge, this judgment should also manifest competence; namely, a meta-competence for evaluating the reliability of one's epistemic performance (Sosa, 2010; 2011). In other words, animal knowledge can enjoy merely first-order justification, while reflective knowledge requires both first-order and second-order justification. Generally speaking, while epistemic support for the proposition p constitutes first-order justification, epistemic support for the reliability of the processes whereby one's belief that p is formed constitutes second-order justification. Evidence for the proper functioning of my visual system constitutes second-order justification for my perceptual belief that p , good calibration of the astronomer's telescope constitutes second-order justification for the accuracy of the measurements made with it, or my reasons for believing that A 's testimony that p is based on A 's knowledge that p constitute second-order justification for p .

In individual scientific inquiry, expertise in the relevant field typically enables scientists to evaluate (i) whether a given body of evidence E confers empirical support to a scientific proposition p , and (ii) the total available higher-order evidence for E so as to

form a judgment on the reliability of the research process whereby E is generated, hence on the evidential quality of E . So, in the relevant field of expertise, the individual scientist's belief whether p can be justified on both the first and the second order, thus constitute knowledge.

Collaborative research is typically interdisciplinary. This implies that individual scientists with relevant field expertise can competently evaluate (i) but not (ii). If a member of a research collaboration lacks basic cognitive access to the first-order scientific evidence produced by the collaboration as its output, hence cannot evaluate (i), then that member of the collaboration is not a candidate for non-testimonial, expert scientific knowledge of the research question. This can often be the case with collaboration members who offer technical support, but do not make a significant agential, non-instrumental contribution to the research process. Such members often also do not have any commitment to the collective scientific claim, for instance as co-authors of publications. At best, they would competently suspend judgment on the truth of the scientific claim of the collaboration. For a scientific investigation to qualify as an interdisciplinary research collaboration, we have to assume that at least two scientists involved in it qualify as candidates for non-testimonial, expert scientific knowledge. If there is only one such candidate, the investigation would be better characterized as an individual research project.²⁸ Typically, several collaboration members will satisfy the basic cognitive access condition. Additionally, most collaboration members would have cognitive access to *some* higher order evidence in light of their own expertise, but typically none will have the adequate competence to evaluate (ii) individually. Taking (i) for granted, then, we can say that in the case of interdisciplinary, collaborative research the individual scientist's belief whether p can have first-order justification but will individually have only partial second-order justification. This is precisely where endorsing epistemic individualism will lead us to deny non-testimonial, expert knowledge to individual scientists.

I believe that the force of the collective-subject argument rests on the implicit intuition that epistemic dependence is not compatible with knowledge.²⁹ Strong anti-individualist perspectives on collective knowledge, such as those of Bird and Palermos,

²⁸ There is no need to dwell on the possibility of there being no candidate for expert scientific knowledge in the research collaboration, as it would fit only an impersonal, objective knowledge account.

²⁹For a similar interpretation, see Pritchard (2015).

arguably still conceive epistemic justification in traditional individualistic terms. They seem to assume, namely, that attributing a belief the status of knowledge, or any other valuable epistemic standing requires that the processes of justification that underly or support the belief should be autonomous. In other words, they should be the primary target of epistemic credit or blame. Since no individual scientist in a research collaboration is not primarily creditable with the success of the distributed research process, there should be a collective subject or agent who is thus creditable. Thus, epistemic dependence will lead us to postulate collective subjects only if we assume that knowledge requires sufficient justification on the basis of autonomous cognitive agency.

From the perspective of a weaker form of anti-individualism, one can be said to know in a way that is dependent on knowledge-enabling external factors if one's agency plays a significant, but not necessarily a primary role in one's epistemic success. Pritchard's (2015) formulation of positive epistemic dependence and his (2010) weak cognitive ability condition on knowledge give us a suitable conception of knowledge that commits to weak anti-individualism:

(Positive) Epistemic Dependence: An epistemic subject can come to know that p by exercising a degree of cognitive agency that is not sufficient for knowing that p through enabling factors that are external to the subject's cognitive agency.

$COGA_{WEAK}$: One knows that p only if one's epistemic success is due to a significant [not necessarily sufficient] degree to one's manifestation of relevant cognitive agency.

Sosa already leaves the door open for knowledge-enabling epistemic dependence by saying that knowledge does not require that the relevant epistemic competence is exclusively seated in the individual (2007, p. 97; 2011, p. 87-88). Sosa's account allows for the case that a subject A comes to know that p through the exercise of a complex social competence that is partially seated in A . A similar formulation of weak epistemic anti-individualism can also be found in an earlier work by Palermos (2015), where he argues

that in certain cases knowledge can be creditable to social factors as well as to the individual.

In the following I will go into the knowledge-enabling external factors and the weak cognitive ability condition with respect to distributed cognitive systems in science, which in conjunction give us an account of epistemically dependent expert scientific knowledge.

3.2.1 First-order justification and reliability of distributed research processes

Any individual's affirmation of a (set of) scientific proposition(s) asserted by a research collaboration will have first-order justification if (i) the individual has basic cognitive access to the evidence provided in support of the scientific proposition(s) asserted by the research collaboration, and (ii) the distributed research process whereby the evidence is generated is a reliable one for investigating the scientific proposition(s) in question. This basic cognitive access comprises the competence to affirm that there is an evidential connection between the evidence and a (set of) scientific proposition(s). Any scientist with a common level of expertise in one of the core constituent fields of the interdisciplinary research would have this competence. The requirement of basic cognitive access gives us the first constituent of the weak cognitive ability condition on expert scientific knowledge, and it is typically satisfied by all members of a research collaboration who make an agential contribution to the research process and are thereby partially accountable for the collectively made assertion. The reliability of the distributed research process, on the other hand, constitutes the first of the two knowledge-enabling external factors we are looking for. Let me elaborate on what is required for a distributed research process to be reliable in terms of the notion of a complex, system level competence.

The reliability of a distributed research process implies that the individual pieces of information (including data, results, other testimony) contributed by the members of the collaboration are true sufficiently often and manifest suitable kinds of scientific expertise, and they cohere into a unified body of scientific evidence necessary for asserting the scientific claim put forward by the collaboration. In order to achieve this, (i) the organization of the distributed cognitive system should realize an efficient division of scientific labor and reliable flow of information, and (ii) the research process should manifest theoretical, methodological and experimental virtues such as valid inferential connections between theory, hypotheses and data, good research design, and proper

choice and application of data analysis tools. The former (i.e., efficient division of scientific labor and reliable flow of information) pertain to the more general or constitutional properties of the distributed cognitive system that enable it to investigate certain kinds of research questions.³⁰ They give us the general *epistemic competence* of the distributed cognitive system as a whole to produce epistemically valuable outputs, such as true empirical propositions, in a certain field. In many research collaborations this general competence also comprises well-calibrated and suitable scientific-technical infrastructure. The latter (i.e., theoretical, methodological and experimental virtues) are the kind of properties one expects to see in the methodology section of a scientific publication and pertain to the particular research process that sets and implements a specific research plan. These constitute the *manifestation* of the general epistemic competence of the distributed cognitive system in the realization of its particular cognitive goal. The appropriate manifestation of the epistemic competence of the distributed cognitive system is a substantial determinant of the evidential quality of the research outputs and the extent of empirical support they confer to the scientific propositions asserted by the research collaboration. In the case of epistemic failure, such as false or highly uninformative empirical findings, both the general epistemic competence of the research collaboration (e.g., lacking sufficient statistical expertise or employing unreliable scientific instruments) and its manifestation (e.g., flaws in the data collection strategy) can be found responsible for the failure.

Together with the condition of a basic cognitive access to the research outcomes, this system-level or distributed epistemic competence and its appropriate manifestation constitute complete first-order justification for affirming the (set of) scientific proposition(s) put forward by a research collaboration.

³⁰It is possible to draw an analogy here to Hardwig's (1991) analysis of trust in a testifier in terms of trust in the epistemic and moral character of the testifier. The epistemic character of the testifier can be replaced by the efficient division of scientific labor in a research collaboration, and the moral character can be replaced by successful (i.e., sufficiently free from error and noise) internal communication. However, instead of trust I prefer to speak of justification, in the reliabilist sense, since a research collaboration has to plan, implement and constantly monitor the performance of its epistemic and social organization.

3.2.2 Distributed second-order justification and reliability of criticism

The objective reliability of the research process, namely the epistemic competence of the distributed cognitive system³¹ to produce *objective* knowledge is only a necessary condition for acquiring *subjective* scientific knowledge through reliance on the distributed research process. A further requirement is that one can positively evaluate the epistemic competence of the distributed cognitive system and its successful exercise, hence the reliability of the distributed research process. This evaluation gives us second-order justification for affirming the (set of) scientific proposition(s) put forward by a research collaboration. In the scientific context, second-order justification concerns all assessments of reliability regarding the data, methods, instruments, or the track-record of other experts as informants. The whole body of such assessments, which we called the process of criticism, constitute second-order justification that the resulting (set of) scientific proposition(s) are the outcome of a reliable process of scientific justification. In research collaborations the process of criticism is necessarily distributed, because the total higher order evidence is a highly heterogeneous set in regard to the expertise required to establish it.

The distributed process of criticism itself can be reliable to differing extents in providing credible assessments of the reliability of the distributed research process. The reliability of the distributed process of criticism implies that the collaboration actively monitors various sources of error and has the necessary social and technological means at its disposal to detect and fix errors when they are present. Using Sosa's terminology, we can call the reliability of the process of criticism the *meta-competence* of the distributed cognitive system for evaluating the reliability of its first-level epistemic performance, which is partially seated in all collaboration members. A reliable socially distributed process of criticism would be organized so as to make use of available expertise and resources in the most efficient and effective way and can do so by relying on the already established social organization of a research collaboration. In HEP experiments the distributed process of criticism involves horizontally organized cross-checking and monitoring tasks, validation mechanisms such as sister experiments (e.g., ATLAS and

³¹While I extend Sosa's notion of epistemic virtue (i.e., competence or reliable skill) to distributed cognitive systems, I do not extend either of his two levels or grades of knowledge beyond the individual agentic components of the system (cf. Carter, 2020; Kallestrup, 2016; Palermos, 2020).

CMS) as well as vertically organized review processes realized by nested work groups, panels and committees. Together with the high transparency and ongoing record-keeping of all aspects of the research process, the distributed process of criticism gives the collaboration members second-order justification to affirm the findings and conclusions. Individual members of a collaboration do not have to scrutinize *all* aspects of the research process when this task of scientific scrutiny or criticism can be realized as a reliable distributed process.

Following Pritchard's formulation of positive epistemic dependence, the reliability of the distributed process of criticism thus gives us the other enabling external factor we were looking for.

3.2.3 Epistemically dependent knowledge

The reliability of the distributed research process and the reliability of the distributed process of criticism together co-determine whether the affirmation of an individual researcher of the scientific proposition(s) put forward by the collaboration (if true) counts as knowledge. In Sosa's (2011) terms, the proper explanation of the epistemic success of an individual researcher's judgment that p features the epistemic competence of the distributed cognitive system for empirically investigating whether p , and its meta-competence for evaluating the reliability of this empirical investigation, both of which are partially seated in the individual researcher.

Both of the two knowledge-enabling external factors, the reliability of the distributed research process and the reliability of the distributed process of criticism, require that the social process of criticism is *spatiotemporally connected* to the research process. Complex distributed research processes require constant monitoring and calibration in order to be and remain reliable. The social process of criticism accordingly should fulfil the functions of both evaluating and maintaining the reliability of research, but without a spatiotemporal connection it cannot fulfil the latter.

The weak cognitive ability condition on knowledge as applicable to non-testimonial, expert scientific knowledge is met when (i) the individual researcher has basic cognitive access to the evidential connection between the empirical findings and the (set of) scientific proposition(s) put forward by the collaboration, and (ii) make an agential epistemic contribution to the distributed process of research and/or its criticism, and (iii)

the outcomes of distributed process of criticism are available to them. When these three conditions are met, the cognitive agencies of individual researchers play a significant part in the explanation of their knowledge.

The cognitive ability condition can be differentially satisfied by different epistemic subjects. The requirement of *cognitive access* (i) is what essentially distinguishes expert from lay scientific knowledge. The requirement of *agential epistemic contribution* (ii) is what essentially distinguishes non-testimonial from testimonial scientific knowledge from a virtue-epistemological perspective. If an individual researcher makes a significant agential, non-instrumental epistemic contribution to the research process and/or its criticism, then the individual epistemic competence of the researcher plays a part in the explanation of the epistemic success of her epistemically dependent belief-forming process. When such contribution is lacking, individual competence will feature in the explanation of the existence and arguably the rationality of individual belief but not in the explanation of its success (see also Kallestrup, 2016). This makes the members of a research collaboration the primary candidates for the possession of collaboratively produced knowledge, but this is not to say that external experts cannot qualify. Even if contribution to the research process itself will be more or less confined to the actual members of the research collaboration, all scientists with relevant expertise will have cognitive access to some part of the higher-order evidence. For instance, an external expert reviewing the published evidence can notice an error or inconsistency in the results or notice a potential defeater such as a possible confounding factor in light of her own background knowledge. Thereby she can be even in a better position than some collaboration members to judge the reliability of the research process. The requirement of *availability* (iii), on the other hand, is dictated by the nature of reflective knowledge itself. For a subject A to come to know that *p* by relying on the research process X, X should not only be objectively reliable, but also A should be well-informed of the evidence of its reliability. However, it is a quite realistic research scenario that the reliability of a certain method, instrument or some other aspect of the research procedure cannot be conclusively assessed at the time it is conducted, but technological or theoretical developments enable a conclusive positive assessment at a much later date. In such cases the researchers would not be in a position to know their scientific conclusions, though they may have good reasons to tentatively accept them and pursue their research project.

When both the two knowledge-enabling external factors and the weak cognitive ability condition on expert scientific knowledge obtain, the individual researcher will have a reliable judgment about the reliability of the scientific justification for the propositions they come to affirm.³² Thus we can combine the two knowledge-enabling external factors and the weak cognitive ability condition under an account of non-testimonial, expert scientific knowledge that is collectively produced and justified:

Collectively produced individual knowledge: An epistemic subject A can come to know that p by relying on the distributed cognitive process X of which evidence E for p is the outcome only if (i) X is a reliable process for establishing the evidence that would be sufficient for knowing that p , (ii) there is a reliable distributed process of criticism Y for evaluating and maintaining the reliability of X that is spatiotemporally connected to X and available to A, (iii) A has basic cognitive access to the evidential connection between E and p , and (iv) A makes a significant agential contribution to X and/or Y.

Abandoning the requirement of epistemic autonomy for knowledge will yield a substantial gain. We can avoid the possibility that no human being learns even from our most successful research collaborations by granting that individual knowledge can be collectively produced and justified; that is, through efficient and reliable social mechanisms for scrutinizing the reliability of the complex body of evidence produced by a research collective. The resulting knowledge will not be “vicarious,” as Hardwig (1985) worried, but non-testimonial, expert scientific knowledge. It will, instead, be suitable to anchor chains of testimonial knowledge. On the other hand, if we deny knowledge even to the researchers who make significant epistemic contributions to a distributed research process, non-contributing scientists who are working in the same discipline, let alone other scientists and lay people, can in no way be said to have any adequate justification to accept the results and thus to be in a position to know. This would place distributed cognitive systems in rather unfit position as social mechanisms for scientific knowledge generation.

³² In Sosa’s (2015) terms, this would amount to a *fully apt* epistemic performance.

Conceiving collective scientific knowledge as collectively produced objective knowledge allows us to accommodate truly distributed cognitive processes of scientific justification, and the concept of socially extended knowledge allows us to retain the commonsensical intuition that objective knowledge implies subjective knowledge. Thus, collective scientific knowledge fruitfully prompts us to reconsider processes of scientific justification without necessarily leading to a dilemma regarding its epistemic subject.

4. Where to? Collective epistemology of science without collective subjects

The account I presented may be said to rule out collective knowledge in an irreducible or strong sense. But such a conclusion does not radically shrink the subject matter of a collective epistemology of science. While I defended here a rather conservative view of the epistemic subject in collaborative research, there is a plethora of directly related issues on which a break with the tradition opens novel, fertile avenues of thought.

My account is potentially compatible with any non-deflationist view of group assertion, group testimony, or group justification. Actually, my two reliabilist conditions for collectively produced knowledge, the reliability of the distributed research process and the reliability of the distributed process of criticism, partly sketch out an account of the justification of assertions made by research collaborations. Intimately related to this is the question of group accountability, which has implications for how we approach issues such as authorship, credit, scientific integrity in regard to research collaborations. However, we conceive group knowledge, group accountability in regard to research collaborations presents a huge task for both theoretical and empirical research (see e.g., Huebner & Bright, 2020; Winsberg, Huebner & Kukla, 2014).

Secondly, I believe that the phenomenon of distributed cognition compels us to develop a non-individualist concept of epistemic competence, and this is a highly fertile topic on its own. This point is also at the heart of several virtue-epistemological perspectives of collective scientific knowledge. But beyond a non-individualist view of epistemic competence, many virtue epistemologists also postulate genuinely collective epistemic subjects in accounting for collective scientific inquiry (e.g., Carter, 2020; Kallestrup, 2016; Palermos, 2020). This strongly anti-individualist position is an attractive option for the virtue epistemologist, arguably because then she does not need to be apologetic in developing a virtue perspective in collective epistemology. By adopting a

strongly anti-individualist stance, the virtue epistemologist does not need to care about the problem of how to individuate competences. She can apply largely the same epistemological toolset in all knowledge situations independently of the particular configuration of epistemic subjects. But this attractiveness comes at a cost: If we make the collective level isomorphic to the individual, we are not really solving anything but just increasing metaphysical commitments. We should instead respect the complexity of collective epistemic processes and be able to recognize the qualitative differences in virtues belonging to distinct levels, which brings me to my last point.

I believe that a collective epistemology of science will be the most fruitful when it investigates how collectives qualitatively differ from individuals rather than how they resemble individuals. Such an epistemological investigation cannot be done with the notions crafted originally for individualist epistemology. If we once more take virtue epistemology as an example, a collective virtue epistemology of science can make a substantial contribution by investigating how individual, group and scientific community level epistemic virtues differ, and how an epistemic virtue at one level may complement, contribute to, or hinder the realization of one at another level. Some examples of community-level (or epistemic system-level, see Goldman, 2011) epistemic virtues can be cognitive diversity (Kitcher, 1993), organized skepticism (Merton, 1973), openness (Vicente-Saez, R., & Martinez-Fuentes, C., 2018) or communalism (Merton, 1973), none of which are sufficiently meaningful properties at the individual level. On the other hand, scientific groups and communities cannot manifest epistemic virtues such as intelligence, open-mindedness, scientific integrity, inquisitiveness or intellectual humility – except metaphorically. A multi-level approach can also greatly benefit the study of epistemic vices in the scientific context. To illustrate, scientific misconduct such as data fabrication or questionable research practices such as data dredging and selective reporting of results or studies may arise not only due to individual factors such as lack of scientific integrity but also due to factors that belong to the structure of the scientific community, for instance through publication bias (Sterling, 1959; Sutton, 2009) or misaligned incentives (Heesen, 2018). Thus, I think that collective epistemology of science would do better if it complements individual epistemology rather than modelling itself on it, and such a collective epistemology of science can do without collective epistemic subjects.

References

- 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*. Oxford University Press.
- Bird, A. (2010). Social knowing. *Philosophical Perspectives*, 24, 23–56
- Clark, A. (1996). *Being There: Putting Brain, Body, and World Together Again*. MIT Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Faulkner, P. (2018). Collective Testimony and Collective Knowledge. *Ergo*, 5.
<https://doi.org/10.3998/ergo.12405314.0005.004>
- Giere, R. (2007). Distributed Cognition without Distributed Knowing. *Social Epistemology*, 21(3), 313–320. <https://doi.org/10.1080/02691720701674197>
- Giere, R. (2002a). Scientific cognition as distributed cognition. In P. Carruthers, S. Stich, and M. Siegal (Eds.), *The Cognitive Basis of Science*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511613517.016>
- Giere, R. (2002b). Discussion Note: Distributed Cognition in Epistemic Cultures. *Philosophy of Science*, 69(December 2002), 637–644.
- Gilbert, M. (2013). *Joint Commitment: How We Make the Social World*. Oxford University Press.
- Gilbert, M. (2004). Collective Epistemology. *Episteme*, 1(2), 95–107.
- Gilbert, M. (1987). Modelling collective belief. *Synthese*, 73, 185–204.
<https://doi.org/10.1007/BF00485446>
- Goldman, A. (2014). Social Process Reliabilism: Solving Justification Problems in Collective Epistemology. In Lackey (Ed.), *Essays in Collective Epistemology*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199665792.003.0002>
- Goldman, A. I. (2011). A guide to social epistemology. In A. I. Goldman & D. Whitcomb (Eds.), *Social Epistemology: Essential Readings*. Oxford University Press.
- Raul Hakli. (2007). On the Possibility of Group Knowledge without Belief, *Social Epistemology*, 21:3, 249-266. <https://doi.org/10.1080/02691720701685581>

- Hakli, Raul. (2006). Group Beliefs and the Distinction Between Belief and Acceptance. *Cognitive Systems Research*, 7 (2-3), 286–97.
- Hardwig, J. (1991). The role of trust in knowledge. *The Journal of Philosophy*, 88(12), 693-708.
- Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy*, 82(7), 335–349.
- Heesen, R. (2018). Why the reward structure of science makes reproducibility problems inevitable. *The Journal of Philosophy*, 115(12), 661-674.
<https://doi.org/10.5840/jphil20181151239>
- Huebner, B., & Bright, L. K. (2020). Collective responsibility and fraud in scientific communities. *The Routledge Handbook of Collective Responsibility*, 358.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Kitcher, P. (1994). Contrasting conceptions of social epistemology. In Schmitt (Ed), *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham, MD: Rowman & Littlefield.
- Knorr-Cetina, K. (1999). *Epistemic Cultures*. Harvard University Press.
- Lackey, J. (2018). Group Assertion. *Erkenntnis*, 83, 21–42. <https://doi.org/10.1007/s10670-016-9870-2>
- Lackey, J. (2014). Socially Extended Knowledge. *Philosophical Issues*, 24(1), 282–298.
<https://doi.org/10.1111/phis.12034>
- List, C., & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Magnus, P. (2007). Distributed cognition and the task of science. *Social Studies of Science*, 37, 297-310
- Merton, R. K. (1973). The Normative Structure of Science. In R. K. Merton (Ed.), *The Sociology of Science: Theoretical and Empirical Investigations* (pp. 267–278). Chicago: University of Chicago Press.

- Palermos, O. (2020). Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism. *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00258-9>
- Palermos, O. (2016a). Spreading the Credit: Virtue Reliabilism and Weak Epistemic Anti-Individualism. *Erkenntnis*, 81, 305–334. <https://doi.org/10.1007/s10670-015-9741-2>
- Palermos, O. (2016b). The Dynamics of Group Cognition. *Minds & Machines*, 26, 409–440. <https://doi.org/10.1007/s11023-016-9402-5>
- Palermos, O. (2011). Belief-Forming Processes, Extended. *Rev.Phil.Psych.*, 2, 741–765. <https://doi.org/10.1007/s13164-011-0075-y>
- Palermos, O., & Pritchard, D. (2013). Extended Knowledge and Social Epistemology. *Social Epistemology Review and Reply Collective*, 2(8),105-120.
- Pino, D. Group (epistemic) competence. (2021). *Synthese*. <https://doi.org/10.1007/s11229-021-03294-2>
- Pritchard, D. (2015). Epistemic Dependence. *Philosophical Perspectives*, 29, 305-324. <https://doi.org/10.1111/phpe.12067>
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175, 133-151. <https://doi.org/10.1007/s11229-010-9738-y>
- Popper, K. R. (1968). Epistemology without a knowing subject. In B. van Rootselaar and J. F. Staal (Eds.), *Proceedings of the Third International Congress for Logic, Methodology and Philosophy of science*. Reprinted in Popper, K. R. (1972). *Objective knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.
- Quinton, A. (1976). Social objects. *Proceedings of the Aristotelian Society*, 76, 1–27.
- de Ridder, J. (2014). Epistemic dependence and collective scientific knowledge. *Synthese*, 191(1), 37–53. <https://doi.org/10.1007/s11229-013-0283-3>
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.
- Sosa, E. (2015). *Judgment and Agency*. Oxford University Press, USA.
- Sosa, E. (2011). *Knowing Full Well*. Princeton University Press.

- Sosa, E. (2010). How Competence Matters in Epistemology. *Philosophical Perspectives*, 24(1), 465–475. <https://doi.org/10.1111/j.1520-8583.2010.00200.x>
- Sosa, E. (2009). *Reflective Knowledge*. Oxford: Oxford University Press.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford University Press.
- Sterling, T. (1959). Publication decisions and their possible effects on inferences drawn from tests of significance—or vice versa. *Journal of the American Statistical Association*, 54(285), 30-34.
- Sutton, A. J. (2009). Publication bias. *The handbook of research synthesis and meta-analysis*, 2, 435-452.
- Thagard, P. (1997). Collaborative Knowledge. *Noûs*, 31(2), 242-261
- Tuomela, R. (2004). Group Knowledge Analyzed. *Episteme*, 1(2), 109-127.
- Tuomela, R. (1995). *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford University Press.
- Tuomela, R. (1992). Group beliefs. *Synthese*, 91(3), 285-318.
- Vaesen, K. (2011). Giere's (In)Appropriation of Distributed Cognition. *Social Epistemology*, 25(4), 379–391. <https://doi.org/10.1080/02691728.2011.604444>
- Vicente-Saez, R., & Martinez-Fuentes, C. (2018). Open Science now: A systematic literature review for an integrated definition. *Journal of business research*, 88, 428-436.
<https://doi.org/10.1016/j.jbusres.2017.12.043>
- Winsberg, E., Huebner, B., & Kukla, R. (2014). Accountability and values in radically collaborative research, *Studies in History and Philosophy of Science Part A*, 46, 16-23.
<https://doi.org/10.1016/j.shpsa.2013.11.007>.
- Wray, K. B. (2007). Who has Scientific Knowledge? *Social Epistemology*, 21(3), 337–347.
- Wray, K. B. (2001). Collective belief and acceptance. *Synthese*, 129(3), 319-333.

Zollman, K. J. S. (2010). The Epistemic Benefit of Transient Diversity. *Erkenntnis* 72 (1), 17.

<https://doi.org/10.1007/s10670-009-9194-6>