Collective scientific knowledge without a collective subject

Duygu Uygun Tunç[1]

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 conceive the processes of knowledge production in these experiments in terms of distributed cognition, 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 an epistemological account, since epistemology traditionally operates within an individualist framework. They both entail that we should deny knowledge to individuals when the processes of knowledge production are distributed. I will argue that epistemology should be extended in a way that can accommodate collectively produced knowledge, but that we would have a serious problem if we deny scientific knowledge to individuals. If the members of a large collaboration cannot be said to know, we 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 argue instead 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 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.

1. Distributed cognition model of collaborative research

Scientific inquiry is at bottom a highly structured cognitive process. Cognitive processes are generally though to occur exclusively within organismic boundaries, so as a cognitive process scientific inquiry is intuitively something that happens in the head of the

[1] Research Fellow, Middle-East Technical University. This work is funded by the European Union under the Horizon 2020 Marie Skłodowska-Curie grant 801509 TÜBİTAK Co-Funded Brain Circulation Scheme 2.

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 highly 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; see also "extended cognition," 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).

While presenting a very useful model for examining the epistemic structure of collaborative science, distributed cognition raises serious doubts about whether we can still conceive scientific knowledge as a state of the traditional subject of epistemology—the individual.

2. The subject of knowledge in research collaborations

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).[2]

Giere (2002b, 2007), on the other hand, finds such an ascription of collective subjectivity to research collaborations too much of an ontological commitment.[3] He argues that we can view certain research collaborations as distributed "cognitive systems" because they realize a cognitive task, 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.

---

[2]In 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 who still remember Durkheim (1933: chap. 3), the conversation produces a version of his much-rebuffed 'conscience collective'."

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

He maintains, instead, that we should characterise 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 rather than folk-psychological term, 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).

3. 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. Distributed cognition provides us with a framework in which we can reconsider this core individualistic assumption and talk about distributed or collective knowledge, as it is increasingly being done in social epistemology. I maintain, however, that this extension or revision of traditional epistemology (cf. Palermos and Pritchard, 2013) should not go as far as postulating distributed or collective epistemic subjects or endorsing an exclusively impersonal view of knowledge in case of distributed cognition. Both these strategies are problematic, and collective production of scientific knowledge does not present us with a forced choice between these two.

3.1. 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 faciae 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 has to show at least how distributed cognition implies distributed mental states. Such an account has to go beyond joint actions and argue for irreducibly collective subjective properties.[4]

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 (2014) and Orestis Palermos (2020) argued for genuinely or irreducibly collective scientific knowledge. Bird,

[4]There 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, 2004) and Margaret Gilbert (1987, 2004). These and similar accounts can possibly suffice in explaining collective knowledge in terms of joint acceptance of propositions or systems of propositions on the basis of collectively acquired or shared evidence, without recourse to collective mental/subjective states. A detailed analysis of the joint acceptance accounts of collective belief or knowledge go beyond the narrow scope of the present paper.

like Knorr-Cetina, invokes Durkheim's concept of "organic solidarity" in grounding distributed cognitive systems as genuine epistemic subjects. 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). He 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 a state. Thus, there can be scientific knowledge (of the group) without any individual knowing. Bird does not even restrict this account to distributed cognitive systems with clearly defined tasks, but extends to wider science on the basis of epistemic interdependence of the scientific community, calling it a single entity.

A core concern here is obviously that Bird's account is actually not able to differentiate between unified cognitive systems and loosely organized epistemic communities, and 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.[5]

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 argues for the further conclusion that the emergent system is an irreducible group entity, which can also

[5]A 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.

be seen as a group mind.[6] The reason is that emergent distributed cognitive systems exhibit, for Palermos, socio-cognitive properties that do not belong to any individual member (2016b).[7] Palermos' account is clearly free from the the kind of inflation of the 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 characterised 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 system as a whole not in an aggregative or summative but irreducible sense, even if no constituent member manifests them.

Besides the costly ontological commitment to collective epistemic subjects, these and similar accounts explicitly acknowledge the probability of a scenario where we can rightly attribute knowledge of a scientific discovery to literally no scientist. This undesirable conclusion, I think, rests 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.[8] 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 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.

Palermos' argument in particular proceeds from collective performances to emergent collective properties, such as epistemic responsibility. I think one can convincingly argue that distributed cognitive systems have weakly emergent collective

[6]See also "extended mind."

[7]Against the possible objection that the atrribution 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.

[8]I have to note that Hutchins himself is more symphatetic to the idea of a distributed mind than I am.

properties. In the case of research collaborations, the required "expertise" for implementing the collectively agreed research design, data collection and analysis methods, manipulation and coordination of instruments and so on is a property of the system as a whole, as well as properties such as the "reliability" and "efficiency" of the research process in yielding credible empirical evidence. Such weakly emergent properties could be among the determinants of whether accepting a scientific proposition counts as knowledge. However, it is not clear what would be gained by attributing strongly emergent subjective/agentive properties such as collective intentionality, consciousness, motivations or beliefs to research collaborations. The distributed research process realized by a collaboration is primarily one of establishing scientific evidence for a proposition by implementing a methodological plan; it is not a process of belief-formation. The epistemic status of the scientists' belief in the scientific proposition collectively asserted by a research collaboration could depend on various weakly emergent properties of the distributed cognitive system such as expertise and reliability, as I will explicate further in the following sections, but we do not need to invoke collective mental states to account for this.

3.2 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 clearly unlike 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 cognitive phenomenon—specifically, a form of belief.

Although he does not specify what he means by impersonal knowledge beyond suggesting that we reformulate knowledge attribution statements in passive form, Giere's impersonal knowledge can lend itself to be interpreted in a way quite similar to Popper's objective knowledge (see esp. Giere, 2007). But the concept of objective knowledge does not tell us by itself anything about the processes of 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 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).

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 be known that the universe is expanding even if the world enters another dark age and nobody is left who understands theoretical physics? The no-subject account of collectively produced knowledge leads us, just like the collective subject account, to the absurd conclusion 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 argues that it is the individual experts who evaluate the outcomes and draw conclusions on the basis of the experiments, and indirectly the lay person through their testimony. Although this kind of knowledge cannot be produced by individuals, it can be known by them (2002b, p. 643).

However, the traditional epistemological concept of knowledge, despite all variety in its analysis, is that of subjective knowledge: a mental (cognitive) 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 agentive constituents of the system might come to entertain true beliefs by accepting the outcome (if the distributed process is successful in yielding true propositions), but they are never 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.[9]

If 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: 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).[10] 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). I think exploring the second (in my opinion more conservative) option is a better strategy in accounting for collective knowledge. But 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.

4. 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 4.1), which can eventually be possessed (though not produced) by constituent individuals when certain conditions are met (section 4.2).

4.1 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 which stand in inferential relations to the reported data given the documented methodological procedures. Thus, as far as we see the product as

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

[10]Freiman and Miller (2020) and Palermos (2016a) call this problem "Hardwig's dilemma."

"knowledge," it is knowledge only in the objective, non-mental sense.[11] We can alternatively say that the distributed cognitive process is only one of evidence-generation in support of collectively made assertions. Either way, the outcome is not knowledge in the subjective sense.

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, agents or subjects, because, as I argued, the research process as a whole is not a mental, agentive or subjective activity like belief formation, but a process of knowledge production in the objective sense or, still more narrowly, one of evidence generation.

4.2 Individual collaboration members as the proper subjects of knowledge

I believe that the force of the collective-subject argument rests on the implicit intiution that epistemic dependence is not compatible with knowledge.[12] Strong anti-individualist perspectives on collective knowledge, such as those of Bird and Palermos, 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 the 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 would lead us to postulate collective subjects only if we assume that knowledge requires sufficient justification on the basis of cognitive agency.

Pritchard's (2015) formulation of positive epistemic dependence gives us a conception of knowledge that commits to a weaker form of 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.

From the perspective of a weak epistemic anti-individualism, one can be said to know in a way that is dependent on enabling external factors if one's agency plays a significant, but not necessarily a primary role in one's epistemic success. One formulation of weak epistemic anti-individualism can be found in an earlier work by Palermos (2015), where he

[11]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.

[12]For a similar interpretation, see Pritchard (2015).

argues that in certain cases knowledge can be creditable to social factors as well as to the individual and in Pritchard's (2010) weak cognitive ability condition on knowledge:

> $COGA_{WEAK}$: One knows that *p* only if one's epistemic success is due to a significant degree to one's manifestation of relevant cognitive agency.

In the following I will go into how we can conceive knowledge-enabling external factors with respect to distributed cognitive systems in science.

4.2.1 Distributed first-order justification and reliability of distributed research processes

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 evidence is established *distributed first-order justification*. It is distributed, since producing such complex scientific evidence exceeds the cognitive ability and capacity of individual researchers and requires a distributed cognitive system.

The constituent members of a research collaboration do not have this kind of complex first-order justification. What they typically have is *partial* first-order justification. However, they can reliably form true beliefs by accepting a scientific proposition that is empirically established through a distributed research process of which they implemented a part. The reliability of such an epistemically dependent belief-formation process refers in significant part to the reliability of the distributed research process, which constitutes one of the enabling external factors we are looking for.

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 they cohere into a unified body of scientific evidence necessary for asserting the scientific claim put forward by the collaboration. That is, on the one hand the organization of the distributed cognitive system should realize an efficient division of scientific labor and reliable flow of information, and on the other the research should manifest theoretical, methodological and experimental virtues such as valid inferential structure, good research design and reliable scientific-technical infrastructure. The former pertain to the properties of the distributed cognitive system that creates and implements a research plan.[13] The latter are related to the properties of

[13]It 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.

objective knowledge the cognitive system is set to generate. Together these two factors constitute the *epistemic competence* of the cognitive system as a whole to produce epistemically valuable outputs such as true empirical propositions. This distributed epistemic competence gives us the complete first-order justification for the (system of) scientific propositions put forward by a research collaboration.

4.2.2 Distributed second-order justification and reliability of criticism

Following Sosa's (2007) twofold distinction between animal and reflective knowledge, we can conceive scientific knowledge (of the scientist) 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 dispositions (i.e., epistemic virtues), but also that one has a positive judgment regarding the reliability of the skills or dispositions in question. 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 gives the astronomer 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 my second-order justification for *p*.

In this regard, the objective reliability of the research process, namely the epistemic competence of the distributed cognitive system[14] 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 thereby the reliability of the distributed research process. This evaluation gives us second-order justification for the (system of) scientific propositions put forward by a research collaboration. In the scientific context, second-order justification concerns the assessments of reliability regarding the data, methods, instruments, or the track-record of other experts as informants. The whole body of such assessments constitute second-order justification that the resulting (system of) scientific propositions are the outcome of a reliable process of scientific justification.

A research collaboration has to aim for effective control of various sources and kinds of error. In a distributed cognitive system, second-order justification may also be distributed; namely, when the required reliability-assessments are made via a distributed social process, where different collaboration members realize different parts of the whole reliability-assessment task. This comprises a wide range from the calibration of instruments to comparison of independent calculations and nested review committees. A

[14]While I extend Sosa's notion of epistemic virtue (i.e., reliable competence) to distributed cognitive systems, I do not extend either of his two levels or grades of knowledge beyond the individual agentive components of the system (cf. Palermos, 2020).

distributed cognitive system can have distributed higher-order regulative mechanisms (based on social practices) to achieve this, which we can call the *distributed social process of criticism*. The reliability of the social process of criticism implies that the collaboration actively monitors sources of error and has the necessary social and technological means at its disposal to detect and fix errors when they are present. 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 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 accept 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 social process of criticism thus gives us the other enabling external factor we were looking for: The reliability of the distributed research process and the reliability of the distributed process of criticism together determine whether the acceptance of an individual member of a collaboration of the scientific proposition(s) put forward by the collaboration (if the proposition is true) counts as knowledge. The cognitive agencies of collaboration members still play a significant part in the explanation of their individual knowledge, since they both significantly contribute to the distributed process of research and its criticism, and are well-informed about the reliability of the scientific justification for the propositions they come to accept. Thus, they also satisfy a weak cognitive ability condition on knowledge.

4.2.3 Epistemically dependent knowledge

In conclusion we can combine the two knowledge-enabling external factors under an account of *epistemically dependent knowledge* in distributed cognitive systems:

> *Epistemically dependent knowledge:* An epistemic subject A can come to know that *p* by relying on the distributed cognitive process X of which evidence for *p* is the outcome if (i) X is a reliable process for establishing the evidence that would be sufficient for knowing that *p*, and (ii) there is a reliable distributed process of criticism for evaluating and maintaining the reliability of X that is available to A.

When the conditions (i) and (ii) are satisfied we can minimally talk about *justified epistemic dependence*, where the individual members of a research collaboration would be justified for accepting the scientific proposition(s) put forward by the research collaboration. If it is

further the case that the scientific proposition(s) are true, we can reasonably attribute epistemically dependent knowledge to the individual members of a research collaboration.

Both (i) and (ii) 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 requirement of availability, on the other hand, is dictated by the nature of reflective knowledge itself. 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 accept them and pursue their research project. Both these requirements, namely spatiotemporal connection and availability, are fulfilled ideally by integrating the social process of criticism into the research process itself in the form of internal criticism.

Lastly, in relation to my criticism of the no-subject and collective-subject accounts of collective scientific knowledge which resonate in their rejection of individual subjective knowledge of (a system of) scientific propositions established via a distributed research process, I would like to reiterate my concern that this rejection leads us to an absurd or undesirable conclusion. Namely, if the members of a large collaboration cannot be said to know, even in the presence of efficient and reliable social mechanisms for scrutinizing the reliability of the complex body of evidence, scientists outside of the collaboration 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. But this would lead to the absurd conclusion that no one learns from the most successful research collaborations we have.

Conceiving collective scientific knowledge as collectively produced objective knowledge allows us to accommodate truly distributed cognitive processes of scientific justification, and the concept of epistemically dependent 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.

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.

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.

Giere, R. 2007. Distributed Cognition without Distributed Knowing. *Social Epistemology, 21*(3), 313–320. doi: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. doi:10.1017/CBO9780511613517.016.

Giere, R. 2002b. Discussion Note: Distributed Cognition in Epistemic Cultures. *Philosophy of Science*, *69*(December 2002), 637–644.

Gilbert, M. Modelling collective belief. *Synthese*, *73*, 185–204 (1987). doi:10.1007/BF00485446

Gilbert, M. 2004. Collective Epistemology. *Episteme*, *1*(2), 95–107.

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.

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.

Miller, B., Freiman, O. 2020. Trust and Distributed Epistemic Labour. In J. Simon (Ed.), *The Routledge Handbook on Trust and Philosophy*. New York: Routledge.

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 (2016). https://doi.org/10.1007/s11023-016-9402-5

Palermos, O., Pritchard, D. 2013. Extended Knowledge and Social Epistemology. *Social Epistemology Review and Reply Collective, 2*(8),105-120.

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.

Rupert, R. 2004. Challenges to the hypothesis of extended cognition. *Journal of Philosophy, 101*(8), 389–428.

Sosa, E. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford University Press.

Thagard, P. 1997. Collaborative Knowledge. *Noûs, 31*(2), 242-261

Tuomela, R. 2004. Group Knowledge Analyzed. *Episteme*, *1*(2), 109-127.

Tuomela, R. 1992. Group beliefs. *Synthese*, *91*(3), 285-318.