

We Should Redefine Scientific Expertise: An Extended Virtue Account

Duygu Uygun Tunç¹

Abstract

An expert is commonly considered to be somebody who possesses the right kind of knowledge and skills to find out true answers for questions in a domain. However, this common conception that focuses only on an individual's knowledge and skills is not very useful to understand the epistemically interdependent nature of contemporary scientific expertise, which becomes increasingly more relevant due to the rise of large interdisciplinary research collaborations. The typical scientific expert today relies substantially on complex scientific instruments and numerous other experts in forming expert judgment. Moreover, we have research collaborations where multiple scientists interact in a way that gives rise to distributed cognitive systems, which can act as a single informant. Accordingly, our concept of scientific expertise should not consist only in individual knowledge and skills, but also accommodate epistemic dependence and collective knowledge production. To this aim, this paper proposes a reconstruction of the concept of scientific expertise as informant reliability by building on the virtue-epistemological account of epistemic competences and theories of extended and distributed cognition. Considered in reference to the social epistemic function of expertise, a scientific expert should be conceived as a reliable informant in a scientific domain, which implies that when consulted on matters in that domain they assert competently, honestly, and completely. Competent expert assertion involves the epistemic responsibility to draw on nothing but the highest degree of epistemic competence relevant to the given context. Thus, being a reliable informant may require one to draw on an extended epistemic competence that goes beyond one's individual competence, or to form supra-individual or group-informants that manifest collectively the kind of complex epistemic competence required for the investigation of certain research questions.

Keywords: scientific expertise; virtue epistemology; epistemic competence; extended cognition; norms of assertion

¹ Postdoctoral Fellow, Eindhoven University of Technology.

Introduction

An expert is commonly considered as a special sort of “knower:” somebody who is deeply knowledgeable in a domain and possesses the pertinent kind of epistemic skills for applying this knowledge to find out true answers for questions in that domain. However, this common concept of an expert is not very useful to understand those contexts of scientific inquiry where the epistemically (inter-)dependent nature of knowledge production is prominent.

Firstly, the possession of knowledge and skill does not tell us much about how the social epistemic role or function of the scientific expert should be realized. Hence, we usually must specify additional characteristics that are demanded of the scientific expert which are not directly implied by knowledgeability and skill, such as intellectual honesty, sincerity, and transparency. These latter kinds of qualities are more closely related to the epistemic responsibility of a scientific expert towards those who consult expert testimony. Where these are absent, the scientific expert does not fail in the sense of lacking the relevant knowledge and skill but fails in the sense of realizing the due epistemic performance of a scientific expert as a member of the scientific and broader epistemic community. The due epistemic performance of an expert is to act as a *reliable informant*. We are interested in and value expertise because we need to rely on some members of the society as the best and dedicated sources of information over a certain field. Within the scientific community we need domain experts because scientific knowledge is a collective enterprise which requires specialization and division of cognitive labor. If this trust in experts is somehow violated too often, the social epistemic responsibility of an expert comes under the spotlight. The way we usually treat such cases of violation of trust in an expert is by engaging in a moral evaluation of the expert’s status but refraining from any epistemic evaluation of it. Thus, the expert in question might be judged to be a bad expert, but still an expert, nonetheless.² I will argue in this paper that we should better incorporate the social epistemic dimension of expertise already in our concept of an expert. Doing this requires that we take as basis of expertise ascriptions not intrinsic qualities like one’s knowledge-base or technical skills, which in practice are often identified through the expert’s credentials, but one’s

² I would like to note that I do not claim or propose here that the moral and epistemic evaluations of expertise are always clearly distinguished. In many cases moral evaluations of experts merge into their epistemic evaluation and vice versa. My intention is rather to point out that our commonplace practices of evaluating expertise epistemically are too restrictive and often exclude equally relevant social epistemic dimensions of expertise. As a result, the social epistemic aspects of expertise are often regarded as a topic for ethics, and even serious violations of the social epistemic responsibilities of an expert might eventually not influence their expert status defined in epistemic terms.

epistemic performance as an informant. This way, a purely epistemic evaluation of an expert's epistemic performance would also reflect its social epistemic implications.

Secondly, maximizing one's reliability as an informant may sometimes require of a scientist to rely on an expertise that is not reducible to their individual knowledge and skill, or to become part of a supra-individual epistemic agent such as a large interdisciplinary collaboration which functions collectively as a single informant. This creates a tension that must be resolved at the expense either of our common conception of expertise or of the social function and value of experts. The common conception the concept of expertise as knowledgeability and skill in a domain commits strongly to the traditional epistemological association between knowledge and epistemic autonomy. For this reason, it cannot easily accommodate the challenge posed by contemporary forms of scientific knowledge production to our conception of knowledge.

Contemporary scientific inquiry depends ever increasingly on technologically enhanced research skills and multidisciplinary research collaborations. The ordinary situation in which we refer to scientific expertise today is not a linear transfer of information from generalists who are in firm command of their instruments and methods. It is rather situations of mutual interdependence between collaborating specialists each of whom has incomplete understanding of the overarching phenomena they are studying as well as the tools they are relying on and methods they are using. The typical scientific expert today relies substantially on scientific instruments, digital tools, and other experts in forming expert judgment. Not rarely expert judgment comes from a group of individuals who act as a single body in making a scientific assertion. We commonly identify such collective epistemic agents as merely groups of experts, although the expertise that is relevant to the collectively made scientific assertion is the one that belongs to the epistemic collective as a whole and irreducibly so. A more useful concept of scientific expertise would accordingly not consist only in individual knowledge, skills and traits, but also accommodate technological and social forms of epistemic interdependence, and extended and distributed epistemic competences.³

Thus, it is worthwhile to reconstruct the concept of expertise in a non-individualist way. I will propose such a reconstruction in terms of the concept of "informant reliability," which is partly inspired by Sosa's epistemic virtue reliabilism (2007, 2009, 2010, 2011, 2015), recently proposed accounts of extended knowing and

³ In this paper I address only those aspects and dimensions of expertise that have to do with epistemic normativity. In this regard, I deal with a rather abstract and ideal notion of expertise. The social phenomenon of expertise clearly is multifaceted, and questions related to social, cultural, moral and political normativity as well as situational factors that influence expertise ascriptions would have significant bearing on a more comprehensive investigation of expertise (see e.g., Collins, 2014; Collins & Evans, 2007; Eyal, 2019). These questions, however, fall outside of the limited scope of this paper.

collective knowledge (e.g., Bird, 2014; 2010; Carter, *forthcoming*; Kallestrup, 2020; Palermos, 2011; 2020; Palermos and Pritchard 2013) and Craig's notion of a good informant (1990). A reliable informant would exercise a kind and degree of epistemic competence that is fit to the investigation of a particular question, which sometimes can imply substantial reliance on external epistemic sources. Moreover, some research questions are better investigated by group agents rather than individuals, in which case a reliable informant would not be an individual but a supra-individual agent such as an extended or distributed cognitive system involving multiple agents as well as technological artifacts. For this reason, instead of conceiving a scientific expert by default as an epistemically autonomous individual and defining expertise in a derivative way, as the cognitive qualities of that individual, we should reverse the conceptual order and start with identifying an individuating principle for scientific expertise. The individuating principle for scientific expertise should be the right kind and degree of epistemic competence as demanded by a given research question, regardless of wherever this competence might be seated. Once we answer the question of expertise regarding the question at hand, we can go on with identifying the most proper host of that kind of expertise: an unaided individual agent, a cognitively extended (i.e., enhanced) agent such as a computer-human system, or an epistemic collective such as a research collaboration.

For this aim, we need at our disposal the tools of a non-individualistic epistemology. Virtue epistemology is particularly suited as a framework to formulate an account of scientific expertise, since it places the emphasis on reliable epistemic competence to produce knowledge without presupposing what the nature of that competence should be (i.e., whether it should be biologically bound, embedded in or scaffolded by external factors, or socio-technologically extended). A non-individualistic construal of virtue epistemology, namely extended virtue epistemology can offer a conception of the scientific expert as a reliable informant whose epistemic competence may be technologically and/or socially extended as demanded by the research question.

The structure of the paper is as follows. I will first present a virtue-epistemological account of scientific expertise where I characterize the former in terms of the notion of informant reliability (1). I will analyze informant reliability in terms of epistemic competence (2) and epistemic responsibility (3). Then I will examine three cases where informant reliability can be considered at a supra-individual level (4). These are cases where the epistemic competence that is manifest in a scientific assertion does not supervene on individual epistemic competence: (i) adversarial collaborations, (ii) technologically extended knowledge-generation, (iii) socially distributed

knowledge-generation. The first case is significant to see that a group of scientists can be a more reliable informant than its individual constituents when the core epistemic competences of the individual scientists do not optimally correlate with informant reliability due to shortcomings in their overall epistemic dispositions. The second case is significant to see that in some situations the reliability of the scientific expert as an informant increases when external cognitive sources are incorporated into the expert's core epistemic competence, thereby extending it. The last case is significant to see that in some contexts a group of scientists can constitute a genuinely collective epistemic agent which manifests a higher level of epistemic competence than any individual scientist possibly can. I conclude the paper by briefly outlining the advantages of reconsidering scientific expertise as a socially relational concept under certain circumstances (5).

1. Expertise as informant reliability

An expert is commonly considered as a special sort of "knower." According to Merriam-Webster, an expert is "one with the special skill or knowledge representing mastery of a particular subject." In the Cambridge dictionary an expert is listed as "a person with a high level of knowledge or skill relating to a particular subject or activity." More technical definitions are also highly similar. Let us take Goldman (2001, p. 92), who defines an expert as "someone who possesses an extensive fund of knowledge (true belief) [in a domain] and a set of skills or methods for apt and successful deployment of this knowledge to new questions in the domain."

However, the possession of knowledge and skill by itself does not tell us much about the social epistemic role or function of the expert, which places normative demands on how and why that knowledge is put to use.⁴ Without reference to this social dimension, we cannot properly distinguish an expert within the much broader category of skillful knowers. In fact, Goldman (2016) himself advances a more recent definition of expertise which acknowledges its social function:

⁴ Quast (2018b) accuses this common approach to defining expertise with reductionism; namely, that it takes expertise to consist in the relevant dispositions an individual possesses. Accordingly, he proposes to balance this approach to expertise with a another that focuses on its function, to wit, fulfilling "a special service role in a restricted domain." This latter notion of expertise has to do with the social role of experts rather than their cognitive properties. While Quast is right at identifying the reductionist character of what I called the common conception of expertise, his attempt at balancing it arguably goes well beyond the bounds of epistemology proper. What he calls the functional sense of expertise is a much more strongly social notion than the one I develop here, and formulated in accordance with other normative considerations. By drawing attention to the social epistemological function of expertise I merely extend the notion of expertise to include the epistemic requirements of proper public affirmations or assertions, and thus remain completely within an epistemological framework.

S is an expert in domain D if and only if S has the capacity to help others (especially laypersons) solve a variety of problems in D or execute an assortment of tasks in D which the latter would not be able to solve or execute on their own. S can provide such help by imparting to the layperson (or other client) his/her distinctive knowledge or skills.

Scientific expertise clearly delimits a smaller region than this broad characterization of capacities, tasks, and problems. Moreover, a significant portion of scientific communication is among experts and thus is much less asymmetrical. But as Goldman's definition also underlines, what is essential to being a scientific expert in distinction to being a knowledgeable and skillful agent is the deployment of these to the epistemic benefit of others. Reference to this social epistemic function of expertise usually requires us to specify additional characteristics that are largely independent of knowledgeability and skill, such as intellectual honesty, sincerity, and transparency, which are more closely related to the *epistemic responsibility* of the expert towards those who consult expert testimony. Knowledgeability and skill may fail to reflect in an expert's testimony due to lack of honesty and transparency – a scientific expert may testify differently from what she knows to be the case (e.g., by falsifying findings), may fail to share the totality of the relevant body of knowledge (e.g., by omitting the mention of a known defeater to the expert's evidence or by selectively reporting findings), or withdraw information (e.g., by abstaining to make negative findings public because they would not bring significant credit from the scientific community). Thus, when an expert testifies unreliably, this does not directly mean they are not knowledgeable or skillful, and the expert's knowledgeability and skill in a domain does not imply that she will duly perform as socially expected of an expert with respect to that domain.

Aristotle's treatment of medical expertise is a good example to illustrate this point. Aristotle characterizes an expertise as a productive and rational capacity. As such, its end is a particular product and involves a true account (*logos*). Thus, an expertise is a productive kind of knowledge (both *epistēmē* and *technē*, see *Meta* Θ2, 1046b3–4; 1046b6–15), which unites the terms knowledgeability and skill in one notion. Like all rational capacities, medical expertise is a two-way capacity; namely, it can give rise to opposite outcomes (i.e., health and its privation, disease). Thus, a medical expert can potentially heal as well as harm an individual. However, the proper application of the medical expertise is producing health and the medical expert (*iatros*) is basically an expert healer because health is the natural end of medicine. If a medical expert uses the medical art to produce disease, she would be acting contrary to the nature of that art (for a more extensive discussion, see Baker, 2021).

I suggest that our concept of a scientific expert should similarly guide us by itself towards the proper epistemic performance of a scientific expert within an information-sharing community: *acting as a reliable informant in a particular scientific domain*. It can only do so if expertise is understood in a way that comprises a responsible epistemic disposition. Scientific experts can potentially use their knowledge and skill in ways that put the others who epistemically rely on them in even worse epistemic situations than they would follow their uneducated opinion, such as when they distort their findings or fabricate data. If we conceive expertise as informant reliability, then the impropriety of such epistemic performances is suggested by the very concept, just as it is the case with a healer that harms.

A vivid example would be a case of scientific fraud that triggered the so-called credibility crisis in Psychological Science. Diederik Stapel was a deeply knowledgeable and skillful researcher in Social Psychology by all standards at his day. He was also a highly eminent academic figure who received numerous distinctions, awards, and plenty of research funding, which are often taken as indirect indicators of excellence in research by the scientific community and the public. In 2011 the news broke out that he consistently fabricated data, which led to the retraction of 59 papers and even of his PhD title.

When we look at the case more closely, we see that Stapel used to design whole experiments with specification of the theoretical background, formulation of the hypotheses, derivation of empirical predictions, and configuration of the experimental interventions and measurements. Then, instead of actually running these experiments he used to fabricate the data in a way that is very difficult to detect the fraud and continue the scientific inquiry procedure as expected with the analysis of the data and inferences about the tested hypotheses and theories (Levelt Committee, 2011).

Is Stapel really an expert? Let us take Goldman's definition. Stapel is "someone who possesses an extensive fund of knowledge (true belief) [in a domain] and a set of skills or methods for apt and successful deployment of this knowledge to new questions in the domain." An expert's core knowledge and skill, in this case that of an experimental scientist's, is typically manifest in the quality and rigor of epistemic performances such as formulation of hypotheses based on the knowledge of extant theories and past observations, experimental design, and statistical analyses of data. What the actual data implies with respect to the tested hypotheses, i.e., whether the predictions were confirmed/falsified or the results were inconclusive, is irrelevant for the assessment of expertise. A good scientific experimenter after all is not someone who is always right but someone who poses valid and epistemically significant empirical questions and develops reliable experimental methods to find out true answers to these

questions. So regardless of the nature of his data, on Goldman's criteria Stapel is an outstanding expert because his epistemic performance in developing social psychological experiments manifested "an extensive fund of knowledge" in the domain and a high level of skill for the "apt and successful deployment of this knowledge to new questions in the domain." However, the actual answers he gave to these questions were even more misleading than those of an imaginary ignoramus who makes up data randomly, because he falsified data skillfully and this made it difficult for others to reveal the fraud. For instance, his results were in the directions predicted by the dominant theories of the field, so they did not arouse suspicion. Moreover, his conclusions were asserted from a position of epistemic authority – so instead of acting as a reliable source of information he actually diminished the others' chances of finding out true answers to those questions. Thus, he did not fail as an expert in the sense of lacking knowledgeability and skill, but he failed in realizing the due epistemic performance of the expert as a member of the epistemic community.

The due epistemic performance of an expert is to act as a reliable informant in a domain. To explicate this, I will briefly introduce Craig's notion of a good informant in connection to Sosa's social epistemic account of the value of knowledge in the following.

Craig's notion of a good informant gives us a fruitful perspective to approach scientific expertise, since it underlines the social roots and function of knowledge. Craig offers a genealogical account of knowledge which traces it back to an imaginary State of Nature. According to Craig, the concept of knowledge is anchored in the ancestor concept of a *good informant*, which reflects the fundamental epistemic needs in any human society; namely, acquiring true beliefs sufficiently often, relying on others for information, and distinguishing good information from bad. A good informant is somebody who is reliable enough as a source of information, open for communication, and whose reliability and openness for communication can be recognized by others (see also Fricker, 2012). This account also finds resonance in Sosa, who says that "the honorific term "knowledgeable" is to be applied only to those who are reliable sources of information, surely an important category for a language-using, social species" (1991, p.27) and, more specifically, that we value justification ultimately because of its function for an information-sharing species; namely, that it indicates that somebody is a "dependable source of information over a certain field in certain circumstances" (1991, p. 275; also 2015, p.66 and Greco, 2012).

Invoking our social epistemic needs to account for why we value knowledge or justification is a common strategy, and it should even more commonsensical to account

for the value of expertise in terms of its social epistemic function. Expertise is a special case of knowledge where the social function and origins are so prominent that they do not require a genealogical argument. We are clearly interested in and value expertise because we need to rely on some members of the society as the best and dedicated sources of information over a certain field. Christian Quast (2018a) expresses this function of expertise within a framework of resource allocation for epistemic needs by saying that it serves to “improve the social deployment of available agential resources apt for an accurate attainment of cliently relevant ends.” This is largely resonant with Hardwig’s (1988) argument for the rationality of relying on expert testimony instead of our own judgment in domains where we do not have sufficient epistemic competence. When we conceive expertise in terms of informant reliability, the social epistemic function of expertise can be explicated with minimal commitment to extra, non-epistemic considerations. An expert produces accurate information on a domain with a low enough misinformation rate, so that we can rationally and responsibly rely on them as an information source in making epistemic judgments in that domain. Moreover, informant reliability captures both the senses of knowledgeable ability and skill and the social epistemic dimension of responsibility under one simple term. Thus, conceiving a scientific expert as a reliable informant rather than a skillful knower gives us a much better target of analysis.

Sosa’s reliable-competence virtue epistemology (2015) is particularly fit as a general framework for this reconstruction, because it gives primacy to agential virtuous competences that can be considered both as reliable knowledge-constitutive competences and as responsibilist intellectual virtues. But we need to distinguish more clearly the virtues that pertain to social epistemic acts from those that pertain merely to individual doxastic states.

Sosa categorizes knowledge as apt affirmation, an affirmation that is epistemically successful because of being made competently, where affirmations can be private or public (2015, p.66f.). Sosa’s account does thus not explicitly distinguish a reliable informant from any other knower. In Sosa’s framework a knower is already a reliable informant – someone who affirms aptly, whether to oneself or to others. Assertion for Sosa is teleologically constituted by virtue of being a sub-class of judgment, where judgment is constitutively affirmation with the aims of truth and aptness (2015, n.4, ch. 8). But we can make a useful distinction here based on Sosa’s telic analysis. The expert is more than a knower, whose aim would be fulfilled with the private attainment of knowledge. The expert’s aim, in accordance with the expert’s function in the epistemic community, is to make knowledge public. We can say that the proper aim of the expert’s epistemic performance is to inform aptly (i.e., to inform

accurately through competence), which goes beyond the proper aim of the knower, which is to believe correctly or indeed aptly (i.e., to obtain accurate beliefs through competence).

For Sosa, “aptness, success that manifests competence, is a norm of performance” (2015, p. 171). Informing (asserting/ testifying) is an epistemic performance, and apt informing is normatively more demanding than apt believing. Thus, while the epistemic performance of the knower is chiefly judged on the basis of the properties of private or mental affirmations (i.e., beliefs), the epistemic performance of the informant should be judged on the basis of the properties of public affirmations or *assertions*. An expert qua reliable informant fails as such if she believes aptly but does not assert aptly when she does so, which constitutively involves an intention to inform aptly. The epistemic performance of the expert qua reliable informant is thus guided by the intention to inform aptly, in resonance with the expert’s social function.⁵ To elaborate, Sosa distinguishes affirmations (whether private or public) into mere alethic affirmations and judgments. While a mere alethic affirmation is done in the endeavor to affirm correctly, a judgment is affirmation in the endeavor to affirm with apt correctness (2015, p.66). We can thus characterize a proper expert assertion by reformulating Sosa’s definition of judgment:

A proper expert assertion is a judgment that is made in the endeavor to inform with apt correctness, when testifying on a subject matter.

Sosa’s account of epistemic competence is sensitive not only to the aim but also to the context and situation of the epistemic performance. Where the threshold will be set for an epistemic performance to be considered sufficiently reliable depends also on the particular epistemic (and possibly also pragmatic) risks that attach to failure. We can differentiate between the respective evaluation-criteria of private vs public affirmations also in this regard (see also Sosa, 2015, p. 173). Expert assertion implies epistemic authority, thus epistemic dependence of many on few. For this reason, the threshold of sufficient reliability for an expert assertion to be considered virtuous is higher than that for private belief. This means that the standards for epistemic vigilance and negligence are stricter for expert assertion to manifest competence, besides the accompanying requirements of providing honest and complete information.

⁵ This telic approach is highly resonant with Craig’s distinction between a good informant and a reliable source of information (although Sosa does not explicitly make a terminological distinction). For Craig (1990, p. 36), a mere source of information does not have the intention to inform properly, which a good informant does: “What I have in mind is the special flavor of situations in which human beings treat each other as subjects with a common purpose, rather than as objects from which services, in this case true belief, can be extracted.”

Informing reliably, thus informing aptly in the endeavor to inform aptly, requires the scientific expert to draw on the best available evidence. Producing or assessing the best available evidence, however, may sometimes require the expert to rely on an epistemic competence that is not completely seated in her. More specifically, in order to increase her reliability as an informant, an expert may have to assert propositions for which she does not have autonomous justification. But this might counter-intuitively imply that the expert asserts what she does not know, because according to the received view expert knowledge requires autonomous justification. Locke (1894, I.iii) says famously that “in the sciences, every one has so much as he really knows and comprehends. What he believes only, and takes upon trust, are but shreds.” The very distinction of expert knowledge from lay knowledge is that the expert is someone who has autonomous justification for the propositions she affirms in her capacity as an expert (see e.g. Schmitt, 1987) So, following this premise, we would have to say that the epistemic responsibility of the expert qua informant might sometimes conflict with her epistemic responsibility qua knower.

Let us take the case of testimonial justification. The layperson’s dependence on experts might indicate lack of nontestimonial reasons altogether, and this dependence might be justified. Experts also rely on testimonial justification, since it is practically impossible and inadmissible to discharge epistemic dependence on testimony altogether. But we expect of the expert that a substantial portion of their justification when asserting in the domain of expertise is non-testimonial. According to Schmitt (1988), experts have an epistemic responsibility qua experts “to discharge testimonial reasons” when “such a discharge would not entail a significant loss of justification” and “to add nontestimonial reasons in the social network of testimonial processes.” An expert qua expert should not rely on undischarged testimony in informing others within the domain of expertise.

However, Schmitt’s formulation of the expert’s epistemic responsibility is not easily applicable to the cases where scientific knowledge production demands substantial reliance on other experts due to an interdisciplinary division of cognitive labor. Hardwig (1988) underlines a potential conflict between the responsibilities of gathering the strongest justification and discharging testimonial justification when there is division of cognitive labor:

Research at the research front is possible in most disciplines only on the basis of division of cognitive labor. Thus, the process of gathering the strongest non-testimonial evidence (first-order justification) depends on employing qua experts «undischarged testimony.»

In research collaborations we see the strongest case of epistemic dependence due to the division of cognitive labor within a single research project. But epistemic dependence is a necessary feature of scientific inquiry in many other research contexts as well, such as when a researcher relies on technological artifacts to realize a part of the process of inquiry but has only a superficial understanding of how and why they are (or sometimes are not) reliable. Considering the complexity and inter/trans-disciplinarity of many research questions that are investigated today, division of cognitive labor is also characteristic of the broader scientific community. The experts may often have to rely on external epistemic sources such as other scientific literatures or the testimony of other experts in order to be cognizant of the best and the total evidence available to the scientific community. Thus, asserting on the basis of the highest level of scientific justification often implies at least partial dependence on other agents or even whole epistemic collectives. Consequently, the scientific is not obligated to be epistemically autonomous.

I argue for the case of non-testimonial but partially dependent or socially distributed expert knowledge on the basis of the concept of extended epistemic competence. To lay the ground, let me first talk about epistemic competence in the context of scientific expertise in general. I will do so in reference to the notions of epistemic competence and epistemic responsibility.

2. Epistemic competence

From the perspective of Sosa's competence-based virtue epistemology, an epistemic competence is a disposition to succeed in forming correct beliefs reliably enough, when one tries (Sosa, 2015, p.43). Sosa's account is based on a two-level conception of epistemic competence and thereby knowledge. Knowledge, according to Sosa, is an *apt* epistemic performance; that is, it is successful and its success manifests a suitable kind of epistemic competence relative to one's aim and situation. Knowledge simpliciter (i.e., animal knowledge (2007) or functional knowledge (2015)) is an epistemic performance that aims at correctness and fulfils this aim competently. Higher level, reflective or judgmental knowledge is an epistemic performance that aims not only at correctness but also at aptness. This further aim requires that the agent is reflectively aware whether the epistemic performance is sufficiently competent, where the threshold for sufficiency reflects the stakes in one's situation (see 2015, ch. 8). If this further aim itself is fulfilled in a way that manifests competence, we speak of reflective knowledge.

Reflective knowledge requires thus not only that one succeeds in acquiring correct beliefs about the subject matter of interest competently, but also that one can make a correct and competent evaluation of one's epistemic competence to do so, the

conditions of success and failure of exercising one's epistemic competence, and thereby the epistemic risks involved in one's attempt to acquire true beliefs competently. Reflective knowledge thus requires a second-order or reflective competence to evaluate one's competence to know and in the appropriateness of one's situation for the exercise of this competence. It is thereby "fully "or "aptly apt" (2015); that is, it is an apt epistemic performance which is aptly judged as being apt by the epistemic agent, hence its aptness manifests not only first-order but also second-order epistemic competence or meta-competence.

Scientific knowledge by default requires reflective evaluation of the methods and processes whereby it is obtained. It has a very low tolerance for epistemic risk and thus a very high threshold for sufficient competence. In the scientific context a fully apt epistemic performance typically produces first and second order scientific justification. First-order scientific justification is typically evidence E towards the truth of a scientific claim p and second-order scientific justification is higher-order evidence E' for the evidential status of E . Second-order justification concerns a broad range of reliability assessments regarding the data, methods, instruments, inferences, or background theories. For instance, good calibration of the astronomer's telescope gives the astronomer second-order evidence for the accuracy of the measurements made with it. Convergence of empirical results from several distinct experiments using dissimilar methods is second-order evidence for the reliability of the first-order evidence obtained in each experiment. Unlike first-order evidence, second-order evidence may be piecemeal and cumulative, but there are often a set of primary reliability-assessment tasks (depending on the study and the field) that have to be completed before conferring any evidential status to a body of findings. The whole body of such assessments constitute second-order scientific justification. Accordingly, the relevant epistemic (meta-)competences in the scientific context are those which serve the production and evaluation of first and second-order evidence. The scientific expert should possess both relevant object-level competences for producing and/or evaluating scientific evidence and meta-competences for evaluating the reliability of object-level competences, instruments, and methods.⁶

This is also precisely the difference between an expert and an enthusiast. An enthusiast might be very knowledgeable in a particular domain, for instance she might

⁶ Reflective knowledge is meta-knowledge, but it is not necessarily binary, as in either knowing that I know or having completely unreflective and (necessarily purely external) justification. On the contrary, knowledge of knowledge in the absolute sense can only be an ideal limiting case. We can more sensibly talk about varying degrees and kinds of meta-competences for evaluating the reliability of object-level epistemic competence.

be deeply familiar with many theories in Psychology and with the state of the literature on a wide range of issues, and well informed about the best research methods. But she would not know why exactly a particular research method is good for a particular research question (e.g., on psychometrical or statistical grounds). The epistemic competence of the enthusiast would be confined to first-order scientific justification and fall short of the epistemic meta-competence that would enable cognitive access to second-order justification, which pertains to the expert qua expert. The enthusiast would thereby lack the reflective knowledge the expert has.

Sosa's analysis of epistemic competence is based on reliable cognitive skills such as perception, memory, and rational inference. In the scientific context, besides such more or less inbuilt general skills we can speak of acquired, specialized cognitive skills such as statistical inference, experimental design or utilization of particular scientific instruments. These and similar skills can be said to constitute a scientific expert's *core* epistemic competence. One's core epistemic competence with respect to a subject matter, let us say whether p , renders one capable to find out whether p if one would properly try to do so under the right conditions. Possessing the core competence to find out whether p does not guarantee that one will exercise it with the intention to find out whether p , will do so under the right conditions, and will assert truthfully and completely what one has found out. Beyond such core epistemic competence, thus, we can speak of a *complete* epistemic competence which implies that one is *properly disposed to publicly answer the question whether p competently and correctly*. This involves further, a competent assessment of the actual fit between the expert's skills and the particular research question at hand and a competent analysis of the context of scientific inquiry (e.g., an experimental setting) for potential sources of error and the epistemic risks associated with these. I elaborate on these in the following. Finally, it also implies an epistemically responsible disposition qua informant, which I will elaborate in section [3](#).

Here it is worthwhile to invoke Sosa's SSS analysis of competences (2010; 2015). According to Sosa's SSS account, a competence in general consists of an innermost skill (i.e., *seat*) that is exercised in proper *shape* and *situation*. To follow his example (2017), for a driver's innermost driving skill (seated in the driver's nervous system and body) to give rise to a complete driving competence manifested in successful driving, the driver must also be awake and sober (i.e., in proper *shape*) and the road and weather conditions must be so that they allow the flow of traffic (i.e., proper *situation*). In the scientific context, proper shape may imply several things. It implies first and foremost that the expert has adequate meta-cognitive control over their epistemic performance so as to be sensitive towards risks of failure in the proper application of cognitive skills. Equally importantly, being in proper shape implies that the expert manifests a generally

knowledge-conducive disposition; that is, they are motivated to draw valid inferences, apply the most suitable and reliable methods, gather correct empirical results and so on ultimately to find out a true answer to whether p . The expert would clearly not be in proper shape to find out whether p if they are motivated to gather data which supports that p irrespective of the truth value of p . Proper shape also implies that the expert can to some significant extent assess whether they are *well situated*; namely, they are sufficiently knowledgeable about the particularities of the research setting and whether they allow or preclude a reliable investigation of whether p . Nonetheless, the right inner competence (*seat plus shape*) may still fail to result in a complete competence, as quite commonplace in the history of science, since there can always be a previously unknown confounding factor in the research setting or the phenomenon that is studied might manifest a hitherto unknown property, which would lead to failed predictions or erroneous results.

As scientific inquiry is a species of reflective knowledge, complete epistemic competence in this context must incorporate a meta-competence for assessing own epistemic performance relative to one's aim and situation. Since scientific inquiry often has novel objects and operates in unfamiliar knowledge situations, a sufficiently reliable meta-competence involves context-sensitive evaluation and adjustment of own epistemic performance, which requires active monitoring of own epistemic performance, reflective assessment of own skills, sensitivity and vigilance towards various sources of error, and competent choice and calibration of methods and instruments. Depending on the case the set of competences and meta-competences might substantially differ. For instance, the competence required for applying a survey questionnaire might fall short of the competence required for assessing its cross-cultural validity.

3. Epistemic responsibility

The complete epistemic competence that is expected from a scientific expert has one more dimension, which we can analyze under the title of epistemic responsibility.

According to the received view, epistemic responsibility attaches to individual doxastic attitudes such as belief, disbelief, or suspension of judgment. For a person to be epistemically responsible for a doxastic attitude is for that person to deserve epistemic credit or blame for forming or maintaining that attitude. Credit or blame in turn are determined by what a person is epistemically obliged to do or avoid to achieve epistemic success, such as justified belief or knowledge (see also Feldman, 2002, p. 376). Since scientific inquiry mostly aims to answer novel questions, the scientific expert's epistemic obligations typically involve diachronic ones that attach to epistemic acts

influencing future doxastic attitudes, such as observing and experimenting.⁷ The scientific expert is thus typically obligated to generate reliable evidence on whether p and do so in a way that is sensitive to possible defeaters. On this basis we can say that the epistemic responsibility of the expert *qua knower* consists in putting her relevant epistemic skills into full use in seeking, evaluating, synthesizing evidence on whether p , vigilance for sources of error and bias, critical monitoring of own epistemic performance, and critical evaluation of the reliability of external sources of information (cf. Sosa, 2015, p.82-83).

The notion of a complete epistemic competence which I have just explicated in reference to Sosa's SSS analysis already implies the epistemic responsibility of an expert *qua knower*. As an account of the proper knowledge-conducive epistemic disposition relative to a situation and a research question, it involves those epistemic responsibilities as pertain to a *reflective knower*. A reflective knower exercises the relevant first-order competences and meta-competences that are responsible for an agent to exercise and irresponsible to neglect in order to attain knowledge (see also Sosa, 2015 p. 82-83). What remains, however, is the expert's epistemic responsibility *qua informant*; namely, those further epistemic obligations that attach to the expert's social epistemic actions. These are epistemic actions that influence others' doxastic attitudes, such as asserting and testifying on whether some states of affairs obtain. Quast (2020, p. 421) rightly argues, in a way that is very rare in the literature, that the expert's responsibilities towards others should be regarded as intrinsic to the relevant set of dispositions that justify expertise ascription: "the character of expertise ascription is dynamic, so that expertise can be retracted even though the subject remains extraordinarily competent, that is, an authority. In other words, the structure of expertise ascription is default and query. Once expertise has been ascribed, it is vulnerable to a later defeat."

Building on the concept of a complete epistemic competence, a *reliable informant* is an agent who exercises the relevant first-order competences and meta-competences to inform others in a way that manifests epistemic responsibility towards others. They are not only properly disposed to know, but also *properly disposed to inform*. The reflective knower might fail their epistemic responsibility towards others by not drawing on the relevant competences and meta-competences in informing others (e.g., drawing on a priori beliefs instead) or refusing to share knowledge in whole or in part (e.g., selectively reporting empirical results or whole studies). Thus, in addition to what we have said regarding expert competence in the preceding section, we can say that the

⁷ Millar (2019) similarly states that epistemic obligations can pertain to the performance or omittance of certain belief-influencing actions that are required for epistemic success.

inner competence (*seat plus shape*) of a scientific expert comprises also the motivation to inform truly and fully regarding whether p .⁸

In addition to the expert's epistemic responsibility as a knower (or information gatherer), her epistemic responsibility as a reliable informant consists then in putting the *relevant epistemic competence into full use and using it as the main source* in making an assertion on whether p ,⁹ and *not to falsify or withdraw information* when asserting in the area of competence. To put it more concisely, the expert is obligated to inform *competently, honestly and completely*.

This general account does not necessarily commit to epistemic individualism, because it leaves open how the "relevant epistemic competence" will be individuated. It can refer to an individual or an extended epistemic competence, or can possibly be understood even as pertaining to a socially distributed cognitive system (see section 4). There are certain limits to this flexibility of course, which are dictated by the status of being an expert. For instance, an expert cannot use another individual's epistemic competence as the main source in testifying or making an assertion as an expert.

We can accordingly elaborate on this general description. We can alternatively say that the epistemic responsibility of the expert as informant consists in using the *relevant extended epistemic competence* as the main source in making assertions on whether p whenever the relevant extended epistemic competence would yield stronger first-order justification (i.e., evidence) for p than the *relevant individual epistemic competence* in isolation. This formulation is largely in line with Hardwig (1988) who maintains that "epistemic interdependence is often epistemologically better than epistemic independence" because "reasons for believing p based partly on testimony can be epistemologically superior to even the best strong nontestimonial evidence for p that anyone, even a researcher in her field of expertise, can have" (For a similar point, see Goldberg, 2021).

I have now sketched what the concept of a reliable informant comprises in terms of the complete epistemic competence of an expert, which involves their core epistemic competence and epistemic responsibility towards others as an informant. Next, I investigate how this picture would look like when being a reliable informant might

⁸ The complete competence to inform (*seat, shape and situation*) may comprise further that the expert is operating in a sufficiently low-noise information environment or that the audience does not by default disregard factual information or dismiss expert opinion.

⁹ Scientific assertions can take a variety of forms: from promising discoveries not yet well-established to uncertainty about a topic or evidence of lack of finding. Moreover, we ask expert opinion to learn whether there is evidence and how strong the evidence is (especially among experts, or in policy making or scientific advisory boards). Thus, I refrain from saying asserting/testifying "that p ." Conditions of competence and responsibility can equally be applied when the expert asserts or testifies that p is *not* the case or there is inconclusive evidence for p .

require one to epistemically depend on external sources and thus partly give up one's epistemic autonomy.

4. Extended Competence

In individual scientific inquiry an expertise in the relevant field typically can enable scientists to competently 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 E' so as to form a judgment on the evidential quality of E . This situation, however, is an idealized case which does not represent a sizeable portion of contemporary science.

Sometimes the epistemic competence that is manifest in an expert's apt assertion may not supervene on the expert's competence as a biological individual. It can be the case that the core or the complete epistemic competence that is primarily relevant for the scientific justification of an assertion incorporates in a substantial sense the epistemic properties/actions of other experts or external cognitive sources such as technological artifacts. Then we can speak of an extended competence, which is only partly seated in the individual agent.¹⁰ By an extended competence I mean an epistemic competence that involves epistemic dependence in a substantive sense. In its stronger sense, extended competence means that the *core* epistemic competence is technologically extended or socially distributed in line with the theses of *extended cognition* (see e.g., Clark, 1997; Clark & Chalmers, 1998; also Hurley, 1998; Menary, 2006; Rowlands 1999; 2009; Wilson, 2004)¹¹ and *distributed cognition* (see e.g., Clark 2017; Hutchins, 1995). Extended competence in this (stronger) sense implies that the relevant cognitive processes are realized by cognitive systems that are composed through the ongoing and reciprocal interaction between a biological and technological system or between multiple cognitive agents and technological instruments.¹² A weaker sense of

¹⁰ In elaborating the specific reference of the *seat* of an extended competence we can resort to the thesis of extended cognition and say that it is a supra-individual cognitive system or to a more modest thesis like embedded cognition and say that it is a complex causal process involving cognitive as well as environmental factors (for a discussion, see Bernecker, 2014). The present argument does not assume a one-to-one correspondence between competence and cognition, so it is broad enough to include extended cognition as well as embedded cognition.

¹¹ I use the term extended cognition rather broadly here. There are distinct theses put forward under this or closely associated terms, such as active externalism, vehicle externalism, or environmentalism, but the present argument does not call for an elaboration on the differences.

¹² In a relatively weaker sense in which core competences can be extended, we can speak of cases where technological instruments are less deeply integrated within the cognitive processes but nonetheless are substantial cognitive supports. Such cases are arguably more in line with the thesis of embedded cognition, where the dependence relations may also be rather asymmetric. However, we do not have to conceive embedded and extended cognition in binary terms. Heersmink (2015) argues that "Seeing

extended competence can be formulated in reference to cases where the complete epistemic competence is extended and not the core epistemic competences or cognitive skills. I use the term extended competence rather broadly to cover both of these.

In the following cases 4.1 through 4.3, I examine situations where even if individual scientists with relevant field expertise can competently evaluate (i), their meta-competence is to evaluate (ii) is either insufficient or undermined. In all three cases, apt epistemic performance (i.e., getting it right aptly on whether p) requires a more complex epistemic competence, an extended competence, and this is achieved in different ways. These cases are presented as different (non-exhaustive) forms of extended epistemic competence in science; namely, adversarial collaborations, technologically extended knowledge-generation and socially extended or distributed knowledge-generation.

The first case is one where the meta-competence of the scientific expert to evaluate the total available higher-order evidence is undermined without necessarily implying epistemic irresponsibility on the expert's part. The undermining factor is the existence of peer disagreement on the extent and nature of the total evidence regarding whether p due to divergent higher-order evidence each expert has in accordance with certain theoretical and/or methodological assumptions the other experts do not share. In a successful adversarial collaboration, a group of experts manifest a higher level of epistemic competence as a group than any of its members thanks to the counterbalancing of individual bias towards certain theories or methods, hence they can be more reliable as a collective informant than any of the individual experts that are its members.

In the second and third cases the scientific expert's meta-competence to evaluate the total higher-order evidence for p is insufficient, hence an adequate assessment of whether p requires the expert to integrate external cognitive sources into their epistemic performance. In the case of technologically extended epistemic competence, the core epistemic competence of the expert is extended through integration of technological artifacts functioning as external cognitive sources. In the case of socially extended or distributed epistemic competence, the relevant epistemic competence is a complex social competence which is partially seated in multiple agents. Thus, the relevant expertise is manifested by a supra-individual or collective epistemic agent. The case of

situated systems as either embedded or extended is not a particularly fruitful way of conceptualizing such systems, as some may be more embedded or extended than others. The nature of their embeddedness or extendedness depends on the degree of integration and, consequently, there is a grey area between systems that are embedded or extended" (see also Sutton et al. 2010; Sterenly, 2010).

distributed knowledge-generation is partially challenging for any account of scientific expertise. For this reason, I will dwell on this last case in particular detail.

4.1 Adversarial collaborations

Some research questions may be highly contested; that is, the scientific literature on them may contain substantial empirical support for all of the opposing claims and no party may possess evidence that is not defeated by the findings of another party or otherwise undermined on theoretical or methodological grounds. Such contested research questions may be extremely difficult to definitively answer, because often the parties endorse incompatible theoretical assumptions, lack common criteria for falsifying points of view and disagree on key methodological issues. The idea of adversarial collaboration has been put forward (Tetlock, 2006; Mellers, Hertwig, & Kahneman, 2001) to organize empirical testing of such contested questions by bringing together scientific experts who have opposing views on the answer to a research question to collectively agree on a research strategy.

Adversarial collaborations are intended to resolve a special case of epistemic peer disagreement; namely, where two or more scientific experts are competent peers with respect to a scientific proposition p but hold opposing or mutually incompatible doxastic attitudes towards p . Typically each contesting party disagrees with the other(s) over the evidential status of the findings the other party put forward in support of p , let us say E_1 and E_2 . This evidential disagreement is due in turn to their theoretical and/or methodological disagreement. More particularly, the parties disagree over the evidential status of E_1 and E_2 because they disagree over the evidential status of each other's higher-order evidence E_1' and E_2' for the evidential support relations $E_1 - p$ and $E_2 - p$.

What concerns us in the case of adversarial collaborations is the fact that a scientific expert's overall epistemic disposition or character may involve certain theoretical and/or methodological commitments which are not necessarily epistemically irresponsible or blameworthy, but nonetheless introduce a certain bias to observations. Scientists often show some degree of dedication to certain research methods, theories or to whole research programs in the absence of overwhelming evidence in favor or even in the presence of some counterevidence or known weaknesses (Kuhn, 1996; also Cowan et al., 2020). This is not necessarily irrational and can sometimes be vital to scientific progress, because it might prevent the premature abandonment of a theory with significant potential or enable the further development and calibration of a method (Lakatos, 1978). However, incompatible commitments to methodologies or theories (as background assumptions) might make it highly difficult for rival scientists to agree on which observations have evidential value and which do

not (Longino, 2020). In cases of such peer disagreement, individual commitments might prevent experts from putting their epistemic competence into use in a way that would maximize their reliability as informants.

In adversarial collaborations the aim is to have contesting parties agree on a research procedure that they all can regard as a suitable test of the research question, and to have them accept and publicly affirm the outcomes of that research procedure in a collective expert assertion (typically in a scientific publication) – no matter what the findings may imply regarding their own positions. When an adversarial collaboration is successfully implemented, the group of experts (or the spokesperson for the group) would be a more reliable informant than any of the individual members, because the resulting group assertion will manifest a complete epistemic competence that is more reliable than those of the individual experts, however high a level of core epistemic competence each may possess vis-a-vis the same research questions. This is because the influence of individual bias would be counter-balanced or at least mitigated.¹³

As an example of a successful adversarial collaboration, we can consider the test of two contenting theories of social evaluation (Koch et al., 2020). According to the stereotype content model, people spontaneously differentiate social groups based on the dimensions of warmth and competence. The agency-beliefs-communion model, on the other hand, proposes the two fundamental stereotype dimensions of agency/socioeconomic success (somehow similar to competence) and conservative-progressive beliefs, while it contests that people spontaneously use warmth as a spontaneous stereotype dimension. The core of the discrepancy is that the agency-beliefs-communion model contests the fundamentality of the warmth (also termed communion) dimension, while the stereotype content model contests the fundamentality of the beliefs dimension. The adversarial collaboration resulted in a consensus finding that the previous studies supporting the agency-belief-communion model were biased by their use of mean-level analysis against warmth/communion compared to agency and beliefs. By using an individual-level analysis, the adversarial collaboration showed all three dimensions as spontaneously used by people in stereotyping groups, thereby demonstrating a way to reconcile the two models. On the other hand, it demonstrated that the stereotype content model had disregarded the fundamental belief dimension, which shows a theoretical limitation.

¹³ It should be noted that I am speaking only of epistemic bias due to theoretical and/or methodological commitments and not of other possible kinds of bias due to, for instance, ideology, moral or aesthetic concerns.

4.2 Extended knowledge-generation

The individual scientist in isolation may also be lacking the set of *core* competences and meta-competences that are required for the investigation of a complex research question. A significant portion of the scientific knowledge produced today relies for this reason substantially on technologically complex procedures. Such knowledge manifests *technologically extended competences* when technological instruments are not mere external tools but become constitutive parts of the cognitive processes underlying research. According to the thesis of extended cognition, this would happen when an artifact is coupled with skin-bound cognitive operations by virtue ongoing two-way interactions (see e.g., Clark & Chalmers 1998). In the context of scientific research there are numerous examples of coupled agent-artifact systems that jointly realize processes of observation, data analysis, modelling, measurement, simulation and so on. Let us briefly consider these from the perspective of epistemic competence and scientific justification.

In many cases of technological instrument use, the expert possesses the relevant meta-competence for evaluating the reliability of the instruments, hence can produce or at least understand the higher-order evidence (E') for the evidential quality of the observations made by virtue of those instruments (E). The expert thus has autonomous reflective knowledge of the scientific propositions they assert; for instance, when experts use software for speed and ease of calculation and can individually cross-check their computations by other means or even with pen and paper. Here I do not think that we should refer to extended cognition in describing the epistemic competence in question, because even though there might be ongoing feedback between the instrument and the expert, the expert's meta-competence for assessing its reliability eliminates any substantial reliance on the instrument.

I think we should speak of extended competences in science especially (or arguably only) when scientific instruments form constitutive and often irreplaceable parts of the relevant cognitive processes. More particularly, when an instrument generates information (not merely processes it) and it so complex that understanding its functioning and assessing its reliability requires its own specialization. For instance, an astronomer often would not have the expertise to know in detail how a space telescope functions so that she can individually calibrate it, or an expert in microbiology who relies on machine learning for classification problems could depend on other experts in computer science or statistics for assessing the reliability of the process of training models.¹⁴ In such cases experts substantially rely on the instrument, in the sense that the

¹⁴ As with any novel technology of scientific instrumentation, there might also be cases of technologically extended scientific cognition where the scientific community as a whole does not yet possess sufficient evidence of the reliability of the instruments, due to not yet having explored the whole range of boundary conditions under which an

instrument constitutes part of the core epistemic competence necessary for producing first-order scientific evidence. Just as we often simply trust our memory (an organic part of our skin-bound cognitive system) without possessing higher-order evidence on its reliability (except maybe for lack of a clear inconsistency or anomaly), experts simply trust the coupled scientific instrument as a source of information in their ongoing interaction (see “trust and glue conditions” in Clark, 2010). As they cannot autonomously assess the reliability of the instruments they rely on to a sufficient extent, some portion of the higher-order evidence must be produced by other experts.

This situation can be said to undermine expertise. If knowledge is a norm of assertion, then expert knowledge is a norm of expert assertion. Expert scientific knowledge is difficult to attribute in cases of epistemic dependence, especially when we speak of extended knowledge-generation, because we commonly expect experts to be epistemically autonomous. Extended processes of scientific justification pose a challenge to the idea of epistemic autonomy, according to which epistemic subjects can be said to know if they are solely or primarily responsible in the production of their knowledge.¹⁵ The challenge is to provide an account of how the individual scientific expert can be said to make apt assertions non-autonomously, which demands that we update our view of expert knowledge to accommodate epistemic dependence.

Sosa’s account already makes room 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

instrument yields reliable measurements, lacking other means to calibrate the instrument with, or not yet having fully discovered its whole spectrum of utilization. For instance, in the case of machine learning some scientists utilize so-called opaque or black box models to make successful predictions, which may nonetheless suffer from what is called link uncertainty: “the amount, kind, and quality of scientific and empirical evidence supporting the link connecting the model to the target-phenomenon” (Sullivan, 2002). High link uncertainty (or insufficient higher-order evidence), for Sullivan, prevents the scientific understanding and explanation of target phenomena. So we might say that while scientific instruments may still prove useful (e.g., in making predictions) in the absence of higher-order scientific evidence, for higher goals like scientific knowledge and understanding sufficient higher-order evidence is a must. High link uncertainty may arguably restrict our ability to ascribe scientific expertise to scientists in some cases. But it is important to note here that it is not necessarily the technological extension of cognition, nor even a lack of insight into the inner workings of the complex algorithms in question that prevents scientific knowledge and understanding of the target phenomena in these cases, but a lack of sufficient higher-order evidence available to the scientific community. Resonantly, Sullivan (2022) argues that “it is not the complexity or black box nature of a model that limits how much understanding the model provides. Instead, it is a lack of scientific and empirical evidence supporting the link that connects a model to the target phenomenon that primarily prohibits understanding.” This issue is arguably an example of a widely discussed issue in the philosophy of measurement, and one of the key insights found in this literature is that evidence for the reliability of novel technologies of scientific instrumentation may accumulate in leaps or even in an iterative manner (see Chang, 2004). Of course, not only progress in available empirical evidence but also theoretical progress plays a role in the second-order justification of research. See e.g. the debate surrounding the behavior of deep learning models with adversarial examples, Buckner, 2020.

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.

From the perspective of a weak epistemic anti-individualism, in certain cases knowledge can be creditable to external factors as well as to the individual. 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. Pritchard's formulation of positive epistemic dependence (2015, p. 307) and his weak cognitive ability condition on knowledge (2010, p. 13-137) together gives us a virtue-theoretical conception of knowledge that commits to such a weak 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.

$COGA_{WEAK}$: 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 light of these we can formulate epistemically dependent expert knowledge as follows:

Extended expert scientific knowledge: A scientific expert A can come to know that p through the exercise of an extended competence partially seated in A which produces scientific evidence E only if (i) E is the kind of evidence that would justify affirming that p (ii) A has cognitive access to the evidential connection between E and p , (iii) and the total higher-order evidence E' for the evidential status of E is available to A (minimally, that E is not defeated).

The cognitive access condition (ii) satisfies the weak cognitive ability condition on knowledge. By a cognitive access to E, I mean a competence to affirm that there is an evidential connection between E and p . The availability of the total higher-order evidence (iii) on the other hand furnishes us the enabling factor that is (wholly or partly) external to A's agency.

4.3 Distributed knowledge-generation

We can also have much more intricate webs of epistemic dependence. A complex epistemic competence may involve multiple complex instruments and multiple experts, and division of epistemic labor between these, such as an interdisciplinary collaboration. In interdisciplinary research collaborations, the relevant epistemic

competences and meta-competences that are manifest in the epistemic success of a scientific assertion may supervene on vast, multi-agent cognitive systems. Such systems of scientific knowledge production have been described in the literature as distributed cognitive systems (see, e.g., Giere & Moffatt, 2003). Here we can speak of a distributed epistemic competence that is partially seated in all individual experts who make a substantial agential contribution to the process of knowledge-generation. Since this is the most challenging case for an account of scientific expertise, I will dwell rather more extensively on it.

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 scientific 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 larger cognitive system comprising multiple agents. As an integral part of the research plan, a research collaboration also engages in various activities for scrutinizing the reliability of the distributed research process, whereby it produces higher-order evidence for the evidential value of first-order scientific evidence. This comprises a wide range from the calibration of instruments to comparison of independent calculations and nested review committees. In a large research collaboration, this process of higher-order evidence generation is typically also distributed; that is, the required reliability-assessments are made via a distributed social process, which we can call the *distributed process of criticism*, where different collaboration members realize different parts of the whole reliability-assessment task and provide diverse kinds of higher-order evidence in accordance with their expertise. For this process to be reliable, the collaboration must actively monitor sources of error and have the necessary social and technological means at its disposal to detect and fix errors when they are present.

The individual scientist is only a contributor to the production of that information. More particularly, the individual scientist will at best have a superficial understanding of the total evidence E for the scientific claim C , and their expertise will enable them to produce and assess only a *part* of the total available higher-order evidence E' for E . No individual expert will have expert level understanding of the total available higher-order evidence E' for E . Thus, the production and assessment of the total first and second-order scientific justification (E and E') will require a complex,

socio-technologically extended competence only partially seated in each individual scientist.

For this reason, it is more appropriate to conceive large research collaborations as a single (supra-individual) informant, which asserts through spokespersons or in impersonal forms such as consortium authorship. Karin Knorr-Cetina (1999) accordingly feels the need to depict the High Energy Physics experiments she observed during her field research stay at CERN in terms of a collective epistemic subject:

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).

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, 2004; Goldman, 2014; Lackey, 2021; List and Pettit, 2011; Tuomela, 2004). In particular, the phenomenon of distributed cognition compels us to develop a non-individualist concept of epistemic competence. This point is also at the heart of several virtue-epistemological perspectives of collective scientific knowledge. 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, *forthcoming*; Kallestrup, 2020; Palermos, 2020).

How can we say that a collective or supra-individual informant asserts with sufficient epistemic justification, especially in regard of the fragmented and distributed nature of the scientific evidence thereby generated? We can formulate a preliminary account of justified group assertion as applies to an interdisciplinary research collaboration that implements a distributed research process as follows. Firstly, a group assertion by a research collaboration is justified only if both the distributed research process for generating the body of first-order evidence E and the process of criticism for generating the body of higher-order evidence E' are indeed reliable. This is broadly in analogy to the conditions of epistemic competence and epistemic responsibility which I analyzed at the individual level. The reliability of the distributed research process is the manifestation of the collaboration's *epistemic competence* as a whole to investigate the research question at hand and similar others. This is the collective analogue of individual epistemic competence. The reliability of (distributed) process of criticism is

manifestation of the collaboration's *meta-competence* to assess the reliability of the particular research process at hand and similar others. This gives us a collective analogue of individual epistemic responsibility.¹⁶ The distributed research process and the distributed process of criticism together constitute the *epistemic performance* of the research collaboration.

Secondly, there is the question of a *group* assertion itself, which typically takes the form of a publication or public announcement stating a scientific claim C by or on behalf of the research collaboration. As a general requirement, the group assertion that C should be properly related to member views on C and their bases. Lackey's (2021) Group Epistemic Agent account of justified group belief gives us a quite suitable framework to concretize this general requirement. According to this, a group justifiably believes that *p* only if a significant percentage of the operative members of the group justifiably believe that *p*. The operative members for Lackey are those who are in a decision-making position. In our case, we can say that these are the members who make significant agential contribution to the research process or at least to its criticism, whereby they produce some part of the evidence E or at least of higher-order evidence E'. This is important as it is what essentially distinguishes non-testimonial from testimonial scientific knowledge.¹⁷ Next, these operative members must affirm the scientific claim C with sufficient justification in order for the collaboration to justifiably assert that C. Sufficient justification in the case of research collaborations can be formulated by drawing on the account of extended expert scientific knowledge I formulated in the preceding.

An operant collaboration member's affirmation that C would be justified if (i) they have a basic cognitive access to the evidential connection between the resulting body of empirical findings E and the scientific proposition C under investigation, (ii) and the outcome of the (distributed) process of criticism is available to them (minimally, whether E is defeated).

¹⁶ Orestis Palermos (2020) similarly argues that epistemic responsibility emerges as a collective property in distributed cognitive systems through self-regulation.

¹⁷ From a virtue epistemological perspective, if an individual researcher makes a significant agential contribution to an epistemic process, 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, 2020). This is not to say that external experts cannot qualify for non-testimonial knowledge. 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 using 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.

A scientist with sufficient expertise in one of the core constituent fields of the interdisciplinary research would typically have the competence needed for a basic grasp of the evidence and its relationship to the scientific claim. Such basic cognitive access to A can be missing in the case when an expert offers merely technical services without contributing to the substantive aspects of a research project, who would than fail to be an operant member of the collaboration.)

In light of the preceding considerations, we can formulate an account of justified group assertion as applies to research collaborations that implement a distributed research process as follows:

Justified group assertion: A research collaboration G can justifiedly assert that C on the basis of implementing a distributed research process X of which evidence E for C is the outcome only if (i) X is a reliable process for establishing the kind of evidence that would justify affirming that C, (ii) G also implements a reliable distributed process of criticism Y, whose outcome is higher-order evidence E' for the evidential status of E, (iii) the operant members of G have cognitive access to the to the evidential connection between E and C, and (iv) the outcome of Y is available to the operant members of G (minimally, that E is not defeated).

Research collaborations are formed when a given research questions requires a complex expertise that no individual agent can manifest by themselves. When justified group assertion obtains in accordance with these conditions, a research collaboration as a collective body will be a much more reliable informant than any individual expert on the research question at hand, as its investigation will manifest the adequate complex expertise that is to be found only at the level of the whole collaboration. Epistemic responsibility thus demands that the investigation of certain research questions are entrusted to research collaborations instead, and that they are identified as “the expert.”

5. Expertise as a socially relational concept

I tried to explicate and defend the advantages of reconstructing the concept of expertise on the basis of informant reliability. In summary, conceiving an expert as a reliable informant enables us to situate the core notion of expert epistemic competence within the social epistemic context that defines its aim and value. From this perspective, an expert is a reliable source of information in a particular domain, which implies that when testifying on matters in that domain they assert competently, honestly and completely. Competent expert assertion involves the epistemic responsibility to draw on nothing but the highest degree of epistemic competence relevant to the given

context. Thus, being a reliable informant may in some situations require that the expert draws on an epistemic competence that goes beyond their innermost epistemic skills and abilities. In other words, the expert may be obligated to assert on the basis of an extended competence which involves epistemic dependence on external sources. Moreover, depending on the kind of epistemic competence required by a given research question, a supra-individual entity such as a research collaboration may be in a much better position to act as an informant, typically through consortium authorship or spokespersons who assert on behalf of a collective body. This does not imply that the knowledge norm of assertion is undermined, though. If we are to speak of any scientific knowledge at the frontiers of many fields today, which would arguably be commonsensical to do, we have to revise the traditional individualist notion of knowledge which requires epistemic autonomy. The kind of epistemic competences that a substantial portion of researchers today possess can yield scientific knowledge only in combination with the competences of others and with complex scientific instruments. Thus, when we judge expert knowledge, we should do so not on the basis of the knowledge that one already has but considering the knowledge one can generate as part of an epistemic network.

From the perspective of epistemic interdependence, an expert would be conceived as a reliable gatherer and provider of information who functions as a node within a network of epistemic processes. This renders the experts' interpersonal epistemic responsibilities particularly salient. These involve manifesting certain *socially relational* properties that are not typically prominent in individual-centered knowledge situations, to which I come last.

I use "relational" in a broad sense here, comprising interpersonal properties as well as certain socially relational dispositions. From a reliabilist standpoint I would particularly underline interpersonal properties that indicate reliability, such as one's good track-record as an epistemic source, or properties that facilitate reliability assessment such as transparency about data, methods, or computer codes. Indicators of reliability will be expected to refer to certain relational dispositions. Most of these are meta-competences, that is competences required for second-order justification – quite similar to those for evaluating one's own epistemic skills and performance, but *relational* in the sense that they are for the justification of epistemically interdependent research processes. Examples can be competence in seeking and synthesizing externally provided information, engaging in interdisciplinary collaboration, or critically evaluating the scientific work of others within the same discipline as well as across disciplines.

Relational meta-competences are crucial in many contemporary contexts of scientific inquiry. We see the importance and value of cultivating relational meta-competences especially in fields with a high degree of division of cognitive labor or those that feature forms of socio-technologically extended scientific inquiry. However, cultivating and exercising relational competences are often not sufficiently rewarded in the current credit system of science which uses individual-level criteria, and this has arguably negative consequences for science. One term that comes close to my notion of relational meta-competences in the meta-science literature is indirect effects (Tiokhin et al., 2021), which is meant to denote the causal effect of a scientist's epistemic performance on a measurable scientific outcome that goes through (or mediated by the epistemic performances of) other scientists. Scientists produce such indirect effects by, for instance, engaging in rigorous peer review, serving as good teachers and mentors, facilitating interdisciplinary communication and coordination, undertaking replication research, or detecting fraud and statistical errors in the scientific literature (p. 3).¹⁸ Ignoring indirect effects in the credit economy of science is argued to increase competition and reduce cooperation between scientific experts and to hinder specialization in skills and roles that are useful mainly in the context of large-scale multidisciplinary collaborations (p. 20).

To refer back to the example of Diederik Stapel's fraud case, it is worth underlying that it was mainly due to the "poor functioning of scientific scrutiny and criticism" (Verfaellie & McGwin, 2011) that he could get away with fabricating data for so long. In closer scrutiny his data were too good to be true; namely, the effect sizes were too large, and the amount of noise was untypically low. If statistical error detection was sufficiently widespread, the ongoing fraud could have been detected much earlier or even deterred. Thus, the individualist conception of expertise not only fails to capture the nature of actual scientific practice but also harms scientific progress by discouraging the exercise and cultivation of relational meta-competences.

A socially relational conception of expertise has the advantage that it defines and characterizes an expert in reference to the specific place they occupy in a network of epistemic dependence. Placing the focus on individual-driven or collective-driven scientific inquiry is ultimately up to the scientific community in a particular field, but if properly crediting the cultivation and exercise of relational meta-competences would be

¹⁸ Some of the items on this list are clearly epistemic performances in that they contribute to higher-order scientific justification (hence to knowledge-generation), such as conducting replications, reviewing, and detecting errors, while others are performances that are rather auxiliary. The term relational meta-competence applies more clearly to the dispositions underlying the former group of performances, while the latter group can be said to manifest auxiliary virtues that enable (rather than substantially contribute to) scientific knowledge generation. For the distinction see also Sosa (2015, ch.2)

set as an aim by a scientific community, then replacing the traditionally individualist notion of expertise with a socially relational notion will be of significant import. Proliferation of research collaborations and increasing specialization is already a visible trend in numerous fields, and a socially relational conception of expertise more faithfully reflects the form scientific inquiry is increasingly taking.

References

- Baker, S. (2021). Aristotle on the Nature and Politics of Medicine. *Apeiron*, 54(4), 411-449. <https://doi.org/10.1515/apeiron-2019-0052>
- Bernecker, S. (2014). How to Understand the Extended Mind. *Philosophical Issues*, 24(1), 1-23. <https://doi.org/10.1111/phis.12023>
- Bird, A. (2014). When Is There a Group that Knows? In J. Lackey (Ed.), *Essays in Collective Epistemology* (pp. 42-63). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199665792.003.0003>
- Bird, A. (2010). Social knowing. *Philosophical Perspectives*, 24, 23-56.
- Buckner, C. (2020). Understanding adversarial examples requires a theory of artefacts for deep learning. *Nature Machine Intelligence*, 2(12), 731-736.
- Cambridge University Press. Expert (n.d.). In *Cambridge dictionary*. Retrieved February 11, 2022, from <https://dictionary.cambridge.org/dictionary/english/expert>.
- Carter, J. Adam (forthcoming). Collective (Telic) Virtue Epistemology. In Mark Alfano, Jeroen de Ridder & Colin Klein (Eds.), *Social Virtue Epistemology*. London: Routledge.
- Chang, Hasok (2004). *Inventing Temperature: Measurement and Scientific Progress*. Oxford University Press.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*, Cambridge, MA: MIT Press.
- Clark, A. (2010). Memento's Revenge: The Extended Mind, Extended. In R. Menary (Ed.), *The Extended Mind* (pp. 43-66). The MIT Press. <https://doi.org/10.7551/mitpress/9780262014038.003.0003>
- Clark, A. (2017). Embodied, Situated, and Distributed Cognition. In *A Companion to Cognitive Science* (pp. 506-517). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781405164535.ch39>
- Clark, A. and Chalmers, D. (1998). The extended mind. *Analysis*, 58: 7-19.
- Craig, E. (1990). *Knowledge and the state of nature: An essay in conceptual synthesis*. Clarendon Press; Oxford University Press.
- Collins, H. (2014). *Are we all scientific experts now?* John Wiley & Sons.
- Collins, H., & Evans, R. (2008). *Rethinking expertise*. In *Rethinking Expertise*. University of Chicago Press.
- Cowan, N., Belletier, C., Doherty, J. M., Jaroslawska, A. J., Rhodes, S., Forsberg, A., Naveh-Benjamin, M., Barrouillet, P., Camos, V., & Logie, R. H. (2020). How Do Scientific Views Change? Notes From an Extended Adversarial Collaboration.

- Perspectives on Psychological Science*, 15(4), 1011–1025.
<https://doi.org/10.1177/1745691620906415>
- Eyal, G. (2019). *The crisis of expertise*. John Wiley & Sons.
- Feldman, R. (2002). Epistemological Duties. In *The Oxford Handbook of Epistemology*. Oxford University Press. <https://doi.org/10.1093/0195130057.003.0013>
- Giere, R. N. and Moffatt, B. (2003). Distributed Cognition: Where the Cognitive and the Social Merge. *Social Studies of Science*, 33(2), 301–310.
<https://doi.org/10.1177/03063127030332017>
- Gilbert, M. (2004). Collective Epistemology. *Episteme*, 1(2), 95–107.
- Goldberg, S. C. (2021). What epistemologists of testimony should learn from philosophers of science? *Synthese*. <https://doi.org/10.1007/s11229-021-03342-x>
- 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>
- Heersmink, R. (2015). Dimensions of integration in embedded and extended cognitive systems. *Phenomenology and the Cognitive Sciences*, 14(3), 577–598.
- Hurley, S. (1998). *Consciousness in action*, Cambridge, MA: Harvard University Press.
- Hutchins, E. (1995). *Cognition in the wild*, Cambridge, MA: MIT Press.
- Schmitt, F. F. (1987). Justification, sociality, and autonomy. *Synthese*, 73(1), 43–85.
<https://doi.org/10.1007/BF00485442>
- Schmitt, F. (1988). On the road to social epistemic interdependence. *Social Epistemology*, 2(4), 297–307. <https://doi.org/10.1080/02691728808578497>
- Sosa, E. (2017). *Epistemology*. Princeton, N.J.: Princeton University Press.
- 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.
- Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge University Press.
- Sterelny, K. (2010). Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465–481. doi:10.1007/s11097-010-9174-y

- Sutton, J., Harris, C. B., Keil, P. G., & Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, 9(4), 521-560.
- Kallestrup, J. (2020). Group virtue epistemology. *Synthese*, 197(12), 5233–5251.
<https://doi.org/10.1007/s11229-016-1225-7>
- Koch, A., Imhoff, R., Unkelbach, C., Nicolas, G., Fiske, S., Terache, J., Carrier, A., Yzerbyt, V. (2020). Groups' warmth is a personal matter: Understanding consensus on stereotype dimensions reconciles adversarial models of social evaluation. *Journal of Experimental Social Psychology*, 89, 103995.
 doi:10.1016/j.jesp.2020.103995
- Kuhn, T. (1996). *The structure of scientific revolutions* (3rd ed.). Chicago: University of Chicago Press.
- Lackey, J. (2021). *The epistemology of groups*. Oxford University Press, USA.
- Lakatos, I. (1978). *The Methodology of Scientific Research Programmes* (J. Worrall & G. Currie, Eds.). Cambridge University Press.
- Longino, H. E. (2020). *Science as Social Knowledge*. Princeton University Press.
- Makin, S. (2006). *Aristotle: Metaphysics Book Θ*. Oxford: Clarendon Press.
- Mellers, B., Hertwig, R., & Kahneman, D. (2001). Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science*, 12, 269-275.
- Menary, R. (2006). Attacking the bounds of cognition. *Philosophical Psychology*, 19(3), 329-344.
- Merriam-Webster. (n.d.). Expert. In *Merriam-Webster.com dictionary*. Retrieved February 11, 2022, from <https://www.merriam-webster.com/dictionary/expert>.
- Millar, B. (2021). Shared Epistemic Responsibility. *Episteme*, 18(4), 493–506.
<https://doi.org/10.1017/epi.2019.21>
- Palermos, S. O. (2011). Belief-forming processes, extended. *Review of philosophy and psychology*, 2(4), 741-765.
- Palermos, O. (2016). 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. (2020). Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism. *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00258-9>
- Palermos, O., & Pritchard D. (2013). Extended Knowledge and Social Epistemology. *Social Epistemology Review and Reply Collective*, 2 (8), 105-120.

- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175(S1), 133–151. <https://doi.org/10.1007/s11229-010-9738-y>
- Pritchard, D. (2015). Epistemic Dependence. *Philosophical Perspectives*, 29(1), 305–324. <https://doi.org/10.1111/phpe.12067>
- Quast, C. (2018a). Expertise: A Practical Explication. *Topoi*, 37(1), 11–27.
- Quast, C. (2018b). Towards a Balanced Account of Expertise. *Social Epistemology*, 32(6), 397–419.
- Quast, C. (2020). On How Expertise Ascriptions Work. *Analyse & Kritik*, 42(2), 399–429.
- Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*, Cambridge: Cambridge University Press.
- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22(1), 1–19.
- Tetlock, P.E. (2006). *Adversarial collaboration: Least feasible when most needed? Least needed when most feasible?* Presentation to Board of Directors of Russell Sage Foundation, New York City.
- Levelt Committee. (2011 October 31). *Interim Report Regarding the Breach of Scientific Integrity by Prof. D. A. Stapel*. Tilburg University, the Netherlands. Archived from the original on 27 June 2016. Retrieved on 15 February 2022 from: https://web.archive.org/web/20160627142859/https://www.tilburguniversity.edu/upload/547aa461-6cd1-48cd-801b-61c434a73f79_interim-report.pdf
- List, C., & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Locke, J. (1894). *An Essay Concerning Human Understanding*. Clarendon Press, Oxford.
- Tiokhin, L., Panchanathan, K., Smaldino, P. E., & Lakens, D. (2021, October 28). Shifting the level of selection in science. <https://doi.org/10.31222/osf.io/juwck>
- Tuomela, R. (2004). Group Knowledge Analyzed. *Episteme*, 1(2), 109–127.
- Verfaellie, M., & McGwin, J. (2011, December). The case of Diederik Stapel. *Psychological Science Agenda*. <http://www.apa.org/science/about/psa/2011/12/diederik-stapel>
- Wilson, R. A. (2004). *Boundaries of the mind: The individual in the fragile sciences – cognition*, Cambridge: Cambridge University Press.