Erkenn (2016) 81:305–334
DOI 10.1007/s10670-015-9741-2

ORIGINAL ARTICLE

Spreading the Credit: Virtue Reliabilism and Weak Epistemic Anti-Individualism

Spyridon Orestis Palermos[1]

Received: 29 September 2014 / Accepted: 19 May 2015 / Published online: 29 May 2015

Abstract Mainstream epistemologists have recently made a few isolated attempts to demonstrate the particular ways, in which specific types of knowledge are partly social. Two promising cases in point are Lackey's (Learning from words: testimony as a source of knowledge. Oxford University Press, Oxford, 2008) dualism in the epistemology of testimony and Goldberg's (Relying on others: an essay in epistemology. Oxford University Press, Oxford, 2010) process reliabilist treatment of testimonial and coverage-support justification. What seems to be missing from the literature, however, is a general approach to knowledge that could reveal the partly social nature of the latter anytime this may be the case. Indicatively, even though Lackey (Synthese 158(3):345–361, 2007) has recently launched an attack against the Credit Account of Knowledge (CAK) on the basis of testimony, she has not classified her view of testimonial knowledge into any of the alternative, general approaches to knowledge. Similarly, even if Goldberg's attempt to provide a process reliabilist explanation of the social nature of testimonial knowledge is deemed satisfactory, his attempt to do the same in the case of coverage-support justification does not deliver the requisite result. This paper demonstrates that CAK can in fact provide, *pace* Lackey's renunciation of the view, a promising account of the social nature of *both* testimonial and coverage-supported knowledge. Additionally, however, it can display further explanatory power by also revealing the social nature of knowledge produced on the basis of epistemic artifacts. Despite their disparities, all these types of knowledge count as partly social in nature, because in all these cases, according to CAK, the epistemic credit for the individual agent's true belief must *spread* between the individual agent *and* certain parts of her epistemic community.

✉ Spyridon Orestis Palermos
S.O.Palermos@ed.ac.uk

[1] School of Philosophy, Psychology and the Language Sciences, Department of Philosophy, University of Edinburgh, Dugald Stewart Building (Rm 5.04), 3 Charles Street, Edinburgh EH8 9AD, Scotland, UK

Accordingly, CAK is a promising candidate for providing a *unified* approach to several and, perhaps all possible, instances of what we may call 'weak epistemic anti-individualism' within mainstream epistemology: i.e., the claim that the nature of knowledge can occasionally be both social and individual at the same time.

1 Introduction

Despite the traditionally individualistic approach of the field, mainstream epistemology has lately started being increasingly perceptive to the social dimensions of knowledge (Alston 2006; Fuller 2007, 2012; Goldman 1999, 2004, 2010; Palermos and Pritchard 2013). Nevertheless, this shift, or rather widening, of focus, raises the following question: How can we allow our epistemic inquiries to be socially oriented without abandoning the methodological individualism that underlies mainstream epistemology? Or, to put it the other way around, how can we pursue mainstream epistemology while accommodating our socio-epistemic intuitions?

This is a pressing question and, indeed, there have already been some attempts to outline such an approach. Each time, however, the focus has only been on the details of the partly social nature of specific types of knowledge, *in isolation*. Lackey's dualism in the epistemology of testimony (2008) and Goldberg's case for 'coverage-support' (2010)[1] are perhaps the most noteworthy examples, both of which constitute remarkable attempts to spell out the particular ways in which these two types of knowledge are—each in its own way—partly social. What seems to be missing from the literature, however, is a general approach to knowledge that could potentially reveal the partly social nature of knowledge *anytime* this may be the case. In some more detail, even though Lackey (2007) has launched an attack against the Credit Account of Knowledge (CAK)—i.e., roughly put, the idea that if *S* knows that *p*, then *S* deserves credit for believing truly that *p*—on the basis of testimony, she does not classify her dualism in the epistemology of testimony (2008) into any of the alternative, general approaches to knowledge. And even if Goldberg's (2010) attempt to provide an explanation of the social nature of testimonial justification in terms of process reliabilism—according to which knowledge is true belief that is the product of a reliable process—is deemed satisfactory, his treatment of coverage-supported beliefs does not deliver the requisite result of 'socializing' coverage-support justification. In other words, a general account that could bring to the fore possibly all the ways in which knowledge might be partly social is yet to be disclosed.

The aim of this paper is to rectify this. Specifically, it will be argued that we can arrive at such an account by elaborating—*pace* Lackey's renunciation of the view—on a virtue reliabilist version of CAK.[2] As we shall see, despite its orthodox

[1] That is, the interesting observation that, on many occasions, we can come to know that *p*, because if not-*p* were the case we would have heard about it by now (by means of the informational channels provided by our community).

[2] The Credit Account of Knowledge, as it will be here considered, is similar to what Lackey (2007) calls the 'Deserving Credit View of Knowledge'. The only difference is that the former, contrary to the latter, is agnostic as to whether the credit *S* deserves for truly believing that *p* in cases of knowledge accounts for

pedigree, CAK—with its focus on the attribution of epistemic credit—is particularly apt for unraveling the exogenous, social nature of several distinct types of knowledge-acquisition: It can provide a promising account of the social nature of *both* testimonial and coverage-supported knowledge, but it can display additional explanatory power by also revealing the social nature of knowledge produced on the basis of epistemic artifacts (in fact, accommodating this type of knowledge may be thought to be quite a demanding test for virtue reliabilism, though, as we shall see, not an insurmountable one): In all these cases, according to CAK, the cognitive success may indeed be creditable to the cognitive agency of the individual subject whose knowledge status is being assessed, but only to a limited degree; the rest of the credit should, or so the argument will go, be attributed to other specific individuals, or to the epistemic society that the individual subject is embedded in. Accordingly, CAK can provide a unified approach to several and perhaps all cases of what we may call '*weak epistemic anti-individualism*' within mainstream epistemology: i.e., the claim that the nature of knowledge can occasionally be both social and individual at the same time.

One important caveat is in order, however. The point here won't be that knowledge might, in certain cases, be *entirely social*—i.e., the objective is not to argue for the possibility of *strong epistemic anti-individualism* within mainstream epistemology. In order to make this clear we need to examine Hardwig's claim (1985) that the sort of epistemic dependence exhibited in cases like the ones mentioned above leads to either of the following two, seemingly unpalatable conclusions: (1) Either we must reject the individualist idea that in order to know one needs to be *intellectually autonomous* in possessing evidence (such as sound arguments and factual information) for the truth of what one knows; or (2) accept that such knowledge is not possessed by any individual alone, but by the epistemic community as a whole.[3] As we shall see, the present approach is an attempt to avoid the second option while explaining how it is possible to retain the first, but in a way that preserves the methodological individualism that underlies mainstream epistemology.

2 The Credit Account of Knowledge

Before we turn to any of the specific ways in which knowledge can be partly social in nature, it is important to first introduce the general approach to knowledge that we will be here focusing on and against the background of which the most specific

Footnote 2 continued
the additional value of knowledge over mere true belief. In other words, while the Credit Account of Knowledge claims that knowledge is creditable true belief, it says nothing about whether the relevant credit is of positive, negative or merely neutral value.

[3] Hardwig (1985), in other words, hints towards strong epistemic anti-individualism. Likewise, Kusch (2002)—motivated partly by considerations similar to those of Hardwig—has put forward a communitarian epistemology. Hardwig's arguments, however, as we will argue later on, are not sufficient for properly motivating strong epistemic anti-individualism. And, while there is no doubt that Kusch's view is strongly anti-individualistic, his communitarian epistemology is rather alienated from mainstream epistemology such that juxtaposing it with what we here call weak epistemic anti-individualism—in mainstream epistemology—would be considerably misleading.

social aspects of knowledge will be later accounted for. The Credit Account of Knowledge (CAK), as captured by virtue reliabilism,[4] can be defined as the combination of the following two intuitions: (1) The ability intuition on knowledge—i.e., the idea that knowledge is belief that is *true in virtue of the manifestation of cognitive ability*[5]; and (2) the intuition that credit is rightfully attributed only in cases of success through ability. Specifically, according to CAK, knowledge—or, at least, a necessary aspect of it—is creditable true belief (Greco 2007, p. 57), which is creditable to a specific individual *S*, because, not only is it a belief that is true, but it is *true in virtue of* the manifestation of *S*'s cognitive ability.[6]

On this view, cognitive ability is understood as a reliable belief-forming process that has been appropriately integrated into the agent's cognitive character, where the agent's cognitive character mainly consists of the agent's cognitive faculties of the brain/central nervous system (CNS), including her natural perceptual faculties, memory, and the overall doxastic system. In addition, however, it can also consist of "acquired skills of perception and acquired methods of inquiry including those involving highly specialized training or even advanced technology" (Greco 1999, p. 287). Here is a relatively weak formulation of virtue reliabilism we can work with:

> $COGA_{weak}$[7]
> If *S* knows that *p,* then *S*'s true belief that *p* is the product of a reliable belief-forming process, which is appropriately integrated within *S*'s cognitive character such that her cognitive success is to a significant degree creditable to her cognitive agency (Pritchard 2010a, pp. 136–137).

What is the reason virtue reliabilists turn to an account of knowledge that stresses the creditable nature of the cognitive success (i.e., the true belief) as well as its origin in the agent's cognitive ability? Mainly, the knowledge-undermining luck involved in Gettier cases. As Gettier demonstrated, one's justified belief may turn out to be true without thereby counting as an instance of knowledge. In the typical scenario, one's belief, which is the product of a faulty justificatory process, *just happens* to be true for reasons that are extraneous to one's justification: In a lucky

[4] So far the Credit Account of Knowledge has only been spelled out in terms of (specific versions) of virtue reliabilism. In principle, however, there is no reason to think it is incompatible with any of the existing alternatives.

[5] For some defenses of this intuition, see (Greco 1999, 2003, 2007; Plantinga 1993; Pritchard 2006, 2010a, b, 2012; Sosa 2007, 2011, 1988, 1988).

[6] According to Greco (1999, 2003, 2010), "in virtue of" should be here understood in causal explanatory terms.

[7] $COGA_{weak}$ stands for Weak COGnitive Agency. Despite the name, this is a particularly promising formulation of virtue reliabilism that is able to accommodate most (if not all) of the problems facing its alternatives. For extensive defenses, see (Pritchard 2010a, 2012; Palermos 2011a, 2014a, forthcoming. In fact, the only reasons it is supposed to be a 'weak' formulation of virtue reliabilism are technical ones: Firstly, it is a necessary (rather than both a necessary and sufficient) condition on knowledge: Several epistemologists hold that virtue reliabilism is a necessary component, but to have an adequate theory of knowledge, they argue, it must be further supplemented by either the safety or the sensitivity principle (see also fn. 15). For such an example, see (Pritchard 2012). Secondly, for reasons to be discussed in Sect. 3.1, notice that $COGA_{weak}$ requires that one's cognitive success be *significantly*, as opposed to *primarily*, creditable to one's cognitive agency.

turn of events, one's belief, which would otherwise be false (given it is produced in a defective way), turns out to be true. Contrast this with cases of success through the manifestation of ability. "There is a sense of 'luck' on which lucky success is precisely opposed to success through virtue or ability" (Greco 2007, 58). When one's true belief is the product of the manifestation of one's ability then believing the truth cannot have been lucky. Of course, one may still be lucky to believe anything at all (because, say, one could have easily been killed), but believing the *truth* is not lucky itself. Accordingly, and since credit is normally attributed in cases of success through ability, virtue reliabilists hold that when some agent knows, his belief must be *true because of his cognitive ability*, such that the cognitive success be, more or less,[8] creditable to him.

Moreover, and before we can leave CAK temporarily to the side, it will be helpful to settle one more question: According to virtue reliabilism, in order for a belief-forming process to count as a cognitive ability it must be part of the agent's cognitive character. Accordingly, an important question to ask is what could it be required in order for a process to be so integrated? As far as common-sense intuitions are concerned, Greco (1999, 2010) has noted that the relevant belief-forming process must be neither strange nor fleeting (i.e., it must be a normal, dispositional cognitive process).[9] Despite such broad intuitions, however, Greco has noted in later work (2010) that in order for a process to be appropriately integrated within one's cognitive character it must interact cooperatively with it. Specifically he writes: "cognitive integration is a function of cooperation and interaction, or cooperative interaction with other aspects of the cognitive system" (2010, p. 152).

3 The Dual Nature of Testimonial and Coverage-Supported Knowledge

While keeping the above in mind, it is now interesting to examine the existing attempts to construe certain types of knowledge as both social and individual. Specifically, the focus will be on Lackey's dualism in the epistemology of testimony and Goldberg's process reliabilist treatment of testimonial and coverage-support justification. All of these accounts constitute remarkably insightful attempts to accentuate the epistemic distribution of labor that takes place in these *specific* types of knowledge. Nevertheless, they fall short of providing a general approach to knowledge; one that can reveal the sense in which the nature of knowledge is partly social in several, and hopefully, all instances whereby this may be the case. In particular, even though Lackey rejects CAK on the basis of considerations having to

[8] As we shall see later on, there is disagreement on the extent to which one's true belief should be creditable to one. For example, Lackey (2007)—an opponent of the view—argues that Greco's Agent Reliabilism (1999, 2010) must be understood as requiring that one's cognitive success be *primarily* creditable to one's self. In contrast, Pritchard's $\text{COGA}_{\text{weak}}$ demands that it only be *significantly* creditable to one's cognitive agency.

[9] The reason why a belief-forming process must be integrated to the agent's cognitive character has to do with epistemic responsibility/subjective justification. For a detailed discussion of subjective justification/epistemic responsibility in terms of cognitive integration as well as in terms of the above (common-sense) intuitions, see Palermos 2014a.

do with testimonial knowledge, she does not classify her view under any of the existing alternative, general approaches to knowledge, and even if one accepts Goldberg's process reliabilist explanation of the social nature of testimonial knowledge, his attempt to 'socialize' the justification of coverage-supported beliefs is not likely to be equally convincing. The aims of this section are (1) to critically examine the insights provided by the aforementioned attempts to reveal the particular ways in which these specific aspects of knowledge are social in nature and (2) to demonstrate that, despite Lackey's renunciation of the view, CAK—with its focus on the attribution of epistemic credit—can in fact prove particularly helpful in accommodating and spelling out the social aspects of both testimonial and coverage-supported justification.

If this section is successful in demonstrating the above, the effect will be that CAK will start figuring as a strong candidate for providing a general approach to knowledge—one that can account for all known, and perhaps all possible, ways in which knowledge may be partly social in nature. Subsequently, the following Sect. 4 will seek to corroborate this expectation of theoretical progress by demonstrating how CAK can also reveal the social nature of knowledge produced on the basis of epistemic artifacts.

3.1 Testimonial Knowledge

Testimonial knowledge has always been a central topic in epistemology. The reason is simple: Very many of our everyday beliefs appear to have testimonial origins. Accordingly, any adequate account of knowledge should be able to accommodate this powerful source of knowledge.

Traditionally, the two opposing sides within the debate concerning testimonial knowledge are those of reductionism, which assigns the entire epistemic burden to the hearer, and non-reductionism, which shifts the epistemic burden to the speaker.[10] In some more detail, reductionists ascribe to the 'positive reasons' thesis, according to which, justification or warrant is attached to testimonial beliefs only by the presence of appropriate positive reasons that depend on sense perception, memory and inductive inference. Testimonial justification or warrant is therefore ultimately reducible to the justification/warrant of these basic epistemic sources.[11] On the contrary, according to non-reductionism, testimony is as epistemically basic as sense perception, memory and inductive inference. Therefore, acquiring testimonial knowledge does not require the possession of any positive reasons on

[10] Goldberg and Henderson (2006), however, argue that it is possible to stick to the letter of non-reductionism, while still assigning *some* of the epistemic burden to the hearer—in the form of a subconscious, monitoring-for-reliability condition.

[11] Hume (2011) is often regarded (quite possibly unjustly) as the best-known supporter of reductionism regarding the epistemology of testimony. Amongst contemporary supporters of the view is Faulkner (2000, pp. 587–588) who claims that "it is doxastically irresponsible to accept testimony without some background belief in the testimony's credibility or truth", and that "an audience is justified in forming a testimonial belief if and only if he is justified in accepting the speaker's testimony." Similarly, Fricker (1994, pp. 149–150) claims that "the hearer should be discriminating in her attitude to the speaker, in that she should be continually evaluating him for trustworthiness throughout their exchange, in the light of the evidence, or cues, available to her".

the part of the hearer but only *the absence* of any negative ones: So long as there are no relevant undefeated defeaters,[12] hearers can acquire testimonially based knowledge merely on the basis of a speaker's report.[13]

In any case, however, regardless of whether one is a reductionist or a non-reductionist, there is no doubt that the speaker does play at least *some* role in the testimonial exchange, and this is a claim we need to examine in far closer detail for the following two reasons. Firstly, it renders testimonial knowledge the perfect candidate for a type of knowledge that is partly social in nature. Secondly, it is this particular, *social* aspect of testimonial knowledge that Lackey has exploited (2007) in order to *argue against* CAK. Accordingly, before we can continue we need to first put this worry to rest.

To appreciate Lackey's objection, consider the following example:

> *Jenny*[14]
> Jenny gets off the train in an unfamiliar city and asks the first person that she meets for directions. The person that she asks is indeed knowledgeable about the area, and gives her directions. Jenny believes what she is told and goes on her way to her intended destination.

Obviously, unless we want to deny a great amount of knowledge that we suppose we have, we must admit that Jenny can gain knowledge in this way. On the basis of this thought experiment, however, Lackey (2007) has argued that, apparently, given the way Jenny gains knowledge, her cognitive character does not have much to do with her true belief. Instead, it is the informant's cognitive character that figures most importantly in the explanation of how Jenny believes the truth.

While this remark may initially sound harmless, Lackey argues that, if true, it is actually very problematic for proponents of CAK. This is because, if they want to account for testimonial knowledge in a way that will accommodate the Jenny case, they must loosen their demands by requiring that the agent's cognitive success of believing the truth only be *significantly* (as opposed to primarily) creditable to the agent's cognitive agency. By doing so, however, virtue reliabilists run the risk of rendering their view impotent with respect to the problem it was initially called to resolve. Remember that virtue reliabilists turn to an account of knowledge that

[12] Following Lackey (2008), there are two relevant types of defeaters that could affect one's acquisition of testimonial knowledge. First, there are psychological defeaters, which are beliefs or doubts that are had by the hearer and which indicate that the hearer's beliefs are either false or unreliably formed. Notice that psychological defeaters may not be objectively correct. Second, there are normative defeaters, which are doubts or beliefs that the hearer ought to have, and which indicate that the hearer's beliefs are either false or unreliably formed. In other words, normative defeaters are beliefs or doubts that the hearer should have (despite whether or not the hearer does actually have them), given the presence of certain available evidence.

[13] Non-reductionism is traditionally associated with the work of Reid (1786). Some contemporary proponents of the view are Burge (1986, p. 47): "a person is entitled to accept as true something that is presented as true and is intelligible to him, unless there are stronger reasons not to do so"; Weiner (2003, p. 57): "we are justified in accepting anything that we are told unless there is positive evidence against doing so"; and Audi (1998, p. 142): "gaining testimonially grounded knowledge normally requires only having no reason for doubt about the credibility of the attester."

[14] (Pritchard 2009, p. 68). It is adapted from Lackey's 'Morris case' (see Lackey 2007, p. 352).

stresses the creditable nature of the cognitive success precisely in order to do away with the knowledge-undermining luck involved in Gettier cases. If, however, in order to accommodate testimonial knowledge, virtue reliabilists demand only a weak—as opposed to strong—degree of creditability, they run the risk of loosing their grip on the problem posed by Gettier: Even though the agent's belief is only luckily true in Gettier cases, he still deserves credit for employing his cognitive ability in order to believe something—no matter whether it is true or false—as opposed to nothing at all. Accordingly, it seems that on a weak version of CAK, victims of Gettier cases incorrectly count as knowing.

This may sound worrisome for $\text{COGA}_{\text{weak}}$, which does indeed demand that the cognitive success be only significantly creditable to the agent's cognitive agency. But if we pay attention to the details of what we have said so far, we will see that it is actually not worrying at all. First, remember that $\text{COGA}_{\text{weak}}$ is only a necessary condition on knowledge that may as well be supplemented with an anti-luck condition in order to account for Gettier cases.[15] But even if we don't want to supplement $\text{COGA}_{\text{weak}}$ with any such modal principle, but we want, instead, to treat it as a full account of knowledge, we can still avoid the problem laid out above. This is because, as Greco (2010) has pointed out, virtue reliabilism ($\text{COGA}_{\text{weak}}$ included) does not merely require that one believe in virtue of cognitive ability *and* that one's belief happen to be true. Instead, virtue reliabilism, properly understood, requires that one's belief be *true in virtue of cognitive ability*—a crucial qualification that allows virtue reliabilists to deal with Gettier cases without necessarily requiring that the cognitive success be *primarily* creditable to one's cognitive agency. To wit, so long as we demand that one believes the *truth* in virtue of one's cognitive character—as opposed to merely demanding that one's belief be true *independently* of how one arrived at it—we can do away with knowledge undermining luck (in the sense operating in Gettier cases), without worrying ourselves with whether the cognitive success may only be *significantly,* as opposed to *primarily,* creditable to one. Accordingly, in principle, virtue reliabilists can also account for testimonial knowledge, no matter that, in such cases, the credit for the hearer's cognitive success is not predominantly down to her.[16] Therefore, virtue reliabilism, in general,

[15] Consider for example *Anti-Luck Virtue Epistemology: S knows that p if and only if S's safe belief that p is the product of her relevant cognitive abilities (such that her safe cognitive success is to a significant degree creditable to her cognitive agency)* (Pritchard 2012). Again, in (Pritchard 2010b, p. 76) we read: "knowledge is safe belief that arises out of the reliable cognitive traits that make up one's cognitive character, such that one's cognitive success is to a significant degree creditable to one's cognitive character". The *sensitivity principle* is usually formulated as follows: If *S* knows that *p,* then *S*'s true belief that *p,* is such that, had *p* been false, *S* would not have believed *p.* The classic defenses of the sensitivity principle can be found in Dretske (1970) and Nozick (1981). The *safety principle* is usually understood thusly: if *S* knows that *p*, then *S*'s true belief that *p,* is such that *S's* belief that *p* could not have easily been false. For recent defenses of the safety principle see Sosa (1999, 2000) and Pritchard (2002, 2008). For a very good discussion concerning the relation between the ability and the anti-luck intuition on knowledge see Pritchard (2012).

[16] In cases of testimonial knowledge, for example, it may be argued (as we shall see below) that even though the hearer's cognitive success is not primarily creditable to her (since at least a significant part of the credit must also go to the speaker for delivering a reliable report), the hearer's belief is still true in virtue of the hearer's cognitive abilities: It is on the basis of these abilities that the hearer spots an appropriate informant, rationally accepts the offered true report, and thereby ends up with the truth.

and $COGA_{weak}$, in particular, are well positioned to avoid Lackey's worries with respect to the dilemma posed by Gettier cases and testimonial knowledge.

And in fact, as we are about to see, Lackey (2008) has recently offered a detailed account of testimonial knowledge that—despite her previous attack to CAK—is quite in line with $COGA_{weak}$. Of course, this is an interesting observation on its own, but unpacking here Lackey's account in terms of $COGA_{weak}$ can also reveal how CAK is well positioned to spell out the distribution of epistemic burden that takes place in cases of testimonial knowledge—a first, yet significant step towards demonstrating that CAK can provide a unified approach to possibly every instance of *weak epistemic anti-individualism.*

To start with, the first thing we need to note is that, contrary to reductionism that puts all of the epistemic responsibility on the hearer and non-reductionism that assigns at least the lion's share—if not the entirety—of the epistemic work to the speaker, the main idea that motivates Lackey's *dualist account* is that it "takes two to tango": "[A]n adequate view of testimonial justification or warrant needs to recognize that the justification or warrant of a hearer's belief has dual sources being grounded in both the reliability of the speaker and the rationality of the hearer's reason for belief" (Lackey 2008, p. 177). Accordingly, Lackey formulates her account this way:

> *Dualism in the Epistemology of Testimony*
> For every speaker, A, and hearer, B, B knows (believes with justification/warrant) that *p* on the basis of A's testimony only if:
> (D1) B believes that *p* on the basis of the content of A's testimony;
> (D2) A's testimony is reliable or otherwise truth conducive;
> (D3) B is a reliable or properly working recipient of testimony;
> (D4) The environment in which B receives A's testimony is suitable for the reception of reliable testimony;
> (D5) B has no undefeated (psychological or normative) defeaters for A's testimony;
> (D6) B has appropriate positive reasons for accepting A's testimony. (Lackey 2008, pp. 177–8)

Now, before explaining how $COGA_{weak}$ can accommodate Lackey's dualism on testimonial knowledge, we should make clear that the former can (easily) handle the Jenny case: Although the cognitive success is not *primarily* creditable to Jenny—but instead a significant part of the credit goes to the stranger—Jenny, by being responsive to epistemically relevant factors (for example, to epistemically reliable contexts and contextual figures, reliable classes of report and reliable speakers)[17] has the right sort of abilities and employs them in the right sort of way so as to deserve significant credit for accepting the stranger's report. After all, she wouldn't believe someone who is obviously drunk or evasive and she wouldn't accept the

[17] For more details on the types of inductively based positive reasons that could allow normal subjects to identify reliable (or unreliable) testimonial reports, see Lackey (2008, 182–183).

suggestion to go past the city hall while she is obviously in a village. Therefore, according to $COGA_{weak}$, Jenny can gain knowledge in this way.

Should we then be confident that $COGA_{weak}$ is in agreement with Lackey's dualist account? And more importantly to the purpose of this paper, can it accommodate the distribution of the epistemic burden that Lackey's account brings to light? Putting aside conditions D1 and D4, which are only meant to ensure that *B* comes to believe *p* on the basis of *A*'s testimonial report (D1) and that there is nothing tricky in the environment the testimonial exchange takes place (D4), it is important to see whether $COGA_{weak}$ can encompass conditions D2, D3, D5 and D6.

Consider, first, conditions D3, D5 and D6, according to which the hearer must: (D3) be a reliable or properly working recipient of testimony; (D5) have no undefeated defeaters for the speaker's testimony; and (D6) have appropriate reasons for the speaker's testimony. The incentive for discussing these three conditions together is that they are jointly meant to ensure the rational or at least, not irrational, acceptance of the speaker's report. Moreover, notice that the absence of any undefeated defeaters against and the possession of positive reasons for a testimonial report—all the while the agent is a proper recipient of testimony in that had he had any negative or positive reasons in mind he would respond appropriately—[18]can be thought of as the two sides of the same coin. Specifically, both conditions require that one be aware of, able and expected to detect any such reasons, should they become evidentially available. The only difference is that in order to acquire testimonial knowledge, no undefeated defeaters must remain in the end, while positive reasons must have been acquired. Importantly, however, both conditions require a fairly active stance on the part of the hearer in the sense that she must be on a continuous lookout—even if subconsciously so—for satisfying them.[19] Satisfaction of D3, D5 and D6 therefore ensures that acquiring a true belief on the basis of a speaker's report will at least be significantly creditable to the hearer's cognitive agency, such that, according to $COGA_{weak}$, her testimonially derived true belief will amount to knowledge.

This leaves us with condition D2, which is the only condition that pertains to the speaker: The speaker's testimony must be reliable or otherwise truth-conducive. First, we should here concentrate on the epistemic burden distribution entailed by D2. As previously noted, Lackey's dualism, contrary to reductionism and non-reductionism that only focus on the hearer or the speaker, distributes the epistemic burden across both parties of the testimonial exchange. So how can $COGA_{weak}$ account for the dual origins of the epistemic justification/warrant?

Remember that according to $COGA_{weak}$, knowledge can be attributed to *S* only if the cognitive success of believing the truth can be significantly credited to *S*'s

[18] For more details on D3, see (Lackey 2008, ch. 7) and (Palermos 2011b).

[19] For an excellent account of how an agent can satisfy conditions D5 and D6 in a *critical* yet unreflective (i.e. subconscious) "snooze" mode, see Fricker (2007) (especially ch. 3). To the contrary, Shieber (2012) argues that evidence from social psychology suggests that we should be skeptical about our abilities to be sensitive either to the trustworthiness or to the deceptiveness of our interlocutors. Still, however, satisfaction of D5 and D6 does not seem to hinge merely on the perceived trustworthiness of the speaker, but also on the reliability of the offered report itself, which depends or several other epistemic factors, such as understanding, plausibility, consistence, coherence and so on.

cognitive agency. Crucially, however, $COGA_{weak}$ denies that the cognitive success must be *wholly* attributed to the hearer's cognitive agency, thereby allowing, in cases of testimonial knowledge, for the rest of the credit to be, at least in part, attributed to the speaker's epistemic effort. We can go back to the Jenny case to see how this works in practice: It is not that Jenny's cognitive character has nothing to do with her believing the truth; it is just that the informant's cognitive character is also important. In some more detail, a significant part of the credit can indeed be attributed to Jenny's cognitive agency for employing the right sort of belief-forming processes so as to detect an appropriate speaker and rationally accept her words on the basis of positive reasons and the lack of any negative ones with respect to the offered report. At the same time, however, the rest of the credit can be, at least in part, attributed to the speaker's cognitive agency for delivering a reliable report on the basis of which Jenny forms her true belief. $COGA_{weak}$, therefore, can accommodate the very essence of Lackey's dualism in the epistemology of testimony—i.e., the distribution of the epistemic credit across both the speaker and the hearer: The hearer deserves a significant amount of the epistemic credit for detecting an appropriate informant and rationally accepting her report, whereas the relevant speaker can be credited with delivering a reliable or otherwise truth-conducive statement that shapes the content of the hearer's true belief.[20]

Overall then, even though Lackey has argued that testimonial knowledge appears to pose a problem for CAK as captured by virtue reliabilism, we see that a virtue reliabilist condition on knowledge, viz., $COGA_{weak}$, can accommodate Lackey's own account of testimonial knowledge—admittedly one of the most promising and detailed accounts on offer.[21] For the purposes of the present discussion, however, what is distinctive of Lackey's account is that it explicitly points out the dual sources of testimonial justification. That is, testimonial justification is not fully reducible to the hearer's reasons for detecting an appropriate speaker and rationally accepting her report. Instead, it also supervenes on the reliability of the speaker's report. Remarkably, $COGA_{weak}$, which captures at least a necessary aspect of knowledge, allows the hearer to acquire knowledge, because the cognitive success can be significantly creditable to her cognitive agency. At the same time, however, it allows for the rest of the credit to be attributed to the speaker for offering a reliable report. Accordingly, $COGA_{weak}$ has the resources to do justice to Lackey's insight regarding the distribution of epistemic burden that takes place in cases of testimonial knowledge. And since knowledge is here understood as *creditable true*

[20] Sosa (2007) too has attempted to spell out the social nature of testimonial knowledge in CAK terms, by arguing that "testimonial knowledge can […] take the form of a belief whose correctness is attributable to a complex social competence only partially seated in that individual believer" (97). Nevertheless, Sosa further notes, such belief may still count as knowledge, "despite how little of the credit for the belief's correctness may belong to the believer individually" (97). In contrast to Sosa's assessment of testimonial knowledge, the point here is that, in cases of testimonial knowledge, the hearer does not deserve little but *significant* part of the epistemic credit, because it is the hearer that is responsible for believing truly, by having detected an appropriate informant and having rationally accepted the offered report. Though, admittedly, a significant part of the credit must again go to the speaker for having delivered a reliable report that shapes the content of the hearer's belief, and in the absence of which the hearer would have to remain agnostic about the relevant matter. See also fn. 28.

[21] For more details on how $COGA_{weak}$ can accommodate Lackey's dualist account see Palermos 2011b.

believing and credit must be attributed to both parties of a testimonial exchange, testimony appears to provide a first instance of weak epistemic anti-individualism within mainstream epistemology: i.e., the claim that knowledge can occasionally be both social and individual in nature.

3.2 Epistemic Coverage Support

In his recent book, Goldberg (2010) appears to share Lackey's insight with respect to the epistemic burden distribution that occurs in cases of testimonial knowledge. In fact, in order to accentuate the speaker's involvement in the production of a reliable testimonial belief, he goes so far as to claim that the belief-forming process that produces the hearer's justified true belief is a single belief-forming process that supervenes on both the hearer and the speaker's cognitive sub-processes: "far from being merely local features of the subject's environment, the testimony itself, along with the cognitive processes implicated in the production of that testimony, are more appropriately regarded as *part of the testimonial belief-forming process itself*. Call this the 'extendedness hypothesis'" (2010, p. 79).[22]

In some more detail, and to familiarize with Goldberg's terms, his argument follows a Goldman-style Process (Historical) Reliabilism (1979), according to which, in order for a belief *p* to be reliable, it is not enough that the final phase of the process that leads to *p* be reliable; it is also necessary that the entire history of the process be reliable. For instance, in cases of belief-dependent, belief-forming processes like memory—where the input to the process of recollection is a belief that has been encoded at some point in the past—the reliability of the final, recollected belief does not only depend on the reliability of the process of recollection. Instead, it also rests on the reliability of the initial process of storing the original belief.

Moreover, early on in his book, Goldberg makes a point about distinguishing between what he calls local and global reliability (2010, p. 12): On one hand, a process is globally reliable (G-reliable) if it produces, or would produce, a preponderance of true over false beliefs, when employed in circumstances similar to the ones it is standardly used. On the other hand, a process is locally reliable if it produces, or would produce, a preponderance of true over false beliefs in circumstances that are relevantly similar to the ones it is being currently used.

Accordingly, as Goldberg further notes on the basis of the above distinction, a true belief will count as known if is both globally and locally reliable, but in order for it to count as *doxastically justified*, it only needs to be globally reliable (*bid.,* p. 159). This is the reason why—and since he is not committed to virtue- but only to process-reliabilism as an account of doxastic justification—Goldberg attempts to unravel the partly social nature of testimonial beliefs by specifically 'socializing' testimonial (G-) reliability. Testimony is a 'quasi-belief dependent', belief-forming process, whose overall reliability is a function of the reliability of its input which, in

[22] Notice that Goldberg does not make his claim on the basis of the extended cognition hypothesis (Clark and Chalmers 1998; Clark 2008). Rather, Goldberg holds that testimony is a belief-forming process that *epistemically* "extends" to the cognitive capacities of the speaker. For Goldberg's disavowal of the hypothesis of extended cognition, see (2010, ch. 5).

turn, depends on the reliability of the cognitive processes that go on in the speaker's head. Accordingly, "doxastic justification [i.e., G-reliability] of a testimonial belief depends [also] on the reliability of the relevant cognitive processes in the hearer's informant" (*ibid.,* p. 81).[23]

There is no doubt that Goldberg's account of testimonial knowledge is very interesting indeed, and despite any possible shortcomings,[24] it does remarkably well in exposing the social dimensions of testimonial knowledge in process reliabilist terms. There is, however, no need to further dwell on its details here, as the aim of this section is to focus on another very interesting epistemological phenomenon, which, even though it is somewhat related to testimony (or, rather, the lack thereof), it is not, as Goldberg himself notes, an instance of his "extendedness hypothesis". I am referring to the *'coverage-reliability' of one's community*—an epistemic phenomenon that is nowadays increasingly effective due to the widespread use of information media such as the printed press, TV, radio and, most importantly, the World Wide Web.

To illustrate Goldberg's point, consider that you know that there is no World-War taking place at the moment, that none of your colleagues was fired in the past few days (especially since it was only yesterday that you saw Jim the gossipmonger), that there are no protests taking place at the city center right now (as you've had the radio on for the last 20 min and there has been no relevant

[23] For a process reliabilist account of testimonial knowledge that is very similar to Goldberg's (2010), see Shieber (2013). Michaelian (2014) is also very sympathetic to Goldberg's process reliabilist account, though he attempts to radicalize Goldberg's extendedness hypothesis, in order to also apply in cases where the reliability of one's beliefs does not only rest on the reliability of other agents but on the reliability of artifacts as well. This may count as an improvement on Goldberg's view, but it will also inherit the problems that Goldberg's process reliabilism faces with respect to coverage-support (to which we are about to turn). Interestingly, in his (2009), Goldberg puts forward a predecessor to his (2010) argument not on the basis of process reliabilism, but on the basis of virtue reliabilism. Even there, however, Goldberg avoids associating virtue reliabilism with the attribution of epistemic credit, thereby distancing himself from CAK. Finally, following, Goldberg's analysis of the cognitive processes that testimonial justification relies on, Green (2012) has attempted to provide an idiosyncratic credit account of knowledge that is specifically designed to account for testimonial knowledge: "CREDIT FOR US: If x knows that p, then the abilities that contribute to the formation and sustenance of x's belief that p deserve primary credit (or something close to it) for x knowing p *whether those abilities are contributed solely by x or also by other agents*" (125). However appealing, one worry is that this is an *ad hoc* account that is motivated by and crafted to accommodate solely considerations pertaining to testimonial knowledge. Other than that, however, the main problem with Green's account, and why we here need to opt for the traditional approach to CAK instead, is that, by failing to tightly tie the agent's cognitive success (i.e., believing the truth of the matter) to the agent's ability, it allows for testimonial knowledge whereby the hearer has done nothing to ensure the reliability of the speaker's report: No one would count me as knowledgeable if Peter Higgs suddenly gave me an anonymous call to report the existence of the Higgs Boson (about which I don't have the foggiest idea), even though, in such a case, credit for my cognitive success could in principle be attributed to both of us.

[24] One possible worry about Goldberg's account is that it may lead to counterintuitive results with respect to epistemic responsibility. Specifically, it is intuitive to think that there is a close correlation between doxastic justification and epistemic responsibility, such that the absence of the former entails the absence of the latter. In cases where the reliability of the speaker's—but not the hearer's—testimony-related processes are defective, however, Goldberg's account rules that the hearer's belief is doxastically *un*justified all the while there being nothing epistemically culpable about the hearer himself. This appears to be a significant worry that needs to be clearly addressed before the 'extendedness' hypothesis can get off the ground.

report), that Messi has not signed a contract with Real Madrid (nothing remotely related to this was mentioned last night during the sport news) and that Madonna, this morning, was still alive (your daughter, who avidly follows her on twitter, looked perfectly fine during breakfast). One of the underlying reasons for all these instances of knowledge, Goldberg claims, is that *if any of those beliefs were true, you would have heard about it by now*. Call the italicized conditional the 'true-to-testimony conditional' (TTTC).

In some more detail, coverage-supported beliefs are a "species of inferential belief, where one of the premises involved is none other than (something like) the truth-to-testimony conditional itself" (Goldberg 2010, p. 174). Specifically, a subject's coverage supported belief that *p* is justified by her current belief that she has no memory of having been informed that not-*p*, *together with* her belief in the relevant instance of the truth-to-testimony conditional. Of course, there will also be relevant inferences that will not hold. Accordingly, Goldberg (2010, pp. 158–164) provides the following set of five conditions that he deems jointly sufficient for a subject's coverage-supported beliefs to count as knowledge:

i) *Source existence condition*: there must be some subgroup of members of the hearer's community—we will call this group "the source"—who are disposed to report about the relevant sort of matters.
ii) *Reliable coverage condition*: the relied-upon source must be reliable in uncovering and subsequently publicizing truths about the domain in which the subject is exhibiting coverage-reliance.
iii) *Sufficient interval condition*: there must be some sort of coordination between the time-related expectations of [the hearer], on the one hand, and the abilities of [the source] to make any relevant discoveries, on the other.
iv) *Silence Condition*: in point of fact, *H* has not encountered any relevant report to date.
v) *Receptivity Condition*: *H* must be such that she *would* come across whatever relevant reports were offered by the source(s) on whom she was relying, were one to be made.

Now, letting the details aside, the above conditions should be relevantly uncontroversial; coverage-supported knowledge does exist, and it most likely behaves in ways very close to what Goldberg's innovative account suggests.

Goldberg's analysis, however, appears to face some serious problems when he attempts to account for the social nature of coverage support justification. As mentioned above, Goldberg avoids accounting for coverage-supported beliefs in terms of his 'extendedness hypothesis', primarily because, in such cases, the agent's environment does not contribute any input to the agent's belief-forming process. Granted, just like testimony, coverage-supported beliefs, on Goldberg's construal, may still count as the product of *belief-dependent*, inferential processes: "The relevant inference would be from the subject's currently formed belief that she has no memory of having been informed that not-*p*, together with her belief in the relevant instance of truth-to-testimony conditional, to the conclusion that *p*" (*ibid.,* pp. 174–175). Since, however, there is no input from the social environment (if

anything, the social environment affects the agent's belief-forming process *via* the *lack* of any input to it), coverage-supported beliefs cannot be construed in terms of the extendedness-hypothesis. But then, if the reliability of coverage-supported beliefs does not depend on the reliability of any input received from the agent's social environment (as in the case of testimonial beliefs), how can Goldberg account for the social nature of such doxastically justified beliefs?

Goldberg summarizes his strategy by noting the following: "The doxastic justification of coverage-supported belief will depend on the doxastic justification of the subject's belief in the relevant truth-to-testimony conditional; and this, in turn, will depend on whether the subject is sensitive to the conditions under which she has relevant reliable coverage" (*bid.,* p. 175). Specifically, Goldberg draws our attention to the "process by which the expectations of the coverage-relying individual are calibrated so as to be brought in line with the prevailing social practices and institutions in her community (*ibid.,* 178).

The problem with the above claims, however, is that, contrary to Goldberg's aim to 'socialize' the doxastic justification of coverage-supported beliefs, they actually accentuate the importance of the individual's efforts to pitch her expectations about the relevant TTTC at the right level. In fact, Goldberg's further remarks confirm this: "While the social institutions and practices I have been discussing constitute the background conditions on a subject's coverage-supported belief, *it is the subject's sensitivity* to the existence and nature of these institutions and practices, *and her sense* of what they portend in terms of the coverage that she is receiving, that determine the G-reliability of her coverage-supported beliefs (*ibid.,* pp. 179–80; my italics). In other words, the social practices and institutions of one's community do constitute the enabling (background) conditions that may allow one to form expectations of coverage-reliability and thereby coverage-supported beliefs. But even though such conditions are necessary for the subject to be able to form *any* coverage-supported beliefs whatsoever, the very G-reliability (i.e., doxastic justification) of her coverage-supported beliefs depends on *her ability* to form the right coverage expectations and thereby on herself, alone.

Clearly, however, this is not the requisite result of 'socializing' the reliability of coverage-supported beliefs. One would, therefore, expect Goldberg to have an ace up his sleeve that could help him turn the tables somehow. Curiously, however, the only move he makes is to offer the following remarks:

> Two different coverage-relying subjects, as alike skin-in as any two distinct individuals can be, might nevertheless differ in the G-reliability of their respective coverage-supported beliefs, as one subject lives in a community in which these institutions and practices provide her with highly reliable coverage on the issue at hand, whereas the other lives in a community where the coverage is less highly reliable (and where there are more issues of interest to her that are not covered). Whatever difference there is in the G-reliability of their respective beliefs supervenes on more than what is going on in their respective heads: it also supervenes on the social practices and institutions that surround them.

Above, I say 'curiously', because, given Goldberg's previous remarks (1) that a process is G-reliable if it is reliable in the environment it is *standardly* employed and (2) that the G-reliability of coverage-supported beliefs is a matter of the subject's sensitivity to the nature and existence of the relevant institutions and social practices, one would have expected Goldberg to have presented a story that is substantially different from the one provided above. Specifically, for any two different coverage-relying subjects, who are as alike skin-in as any two distinct individuals can be, but who live in communities with different social practices and institutions, *if they both are to be G-reliable* in their respective coverage-supported beliefs *at all,* they must differ internally in one crucial respect: They must have different expectations about their communities and thereby relevant TTTCs.[25] Accordingly, for any given domain of beliefs, and for any two individuals who are both G-reliable in their coverage-supported beliefs in general, any belief *p* (under the given domain) must be *equally* coverage-supported G-reliable—that is, G-reliable *simpliciter* (see also fn. 25)—for both of them, because they are both in a position to hold some appropriately pitched TTTCs with respect to the relevant domain. Otherwise, if the two individuals differ with respect to their doxastic justification with respect to *p*, it will be because one of them does not even hold a coverage-supported belief in *p*, because her community does not even provide the necessary background conditions in order for her to form (on the basis of some appropriately-pitched TTTC) such a doxastically justified (i.e., G-reliable) coverage-supported belief in *p*. In any case, however, the G-reliability of their respective coverage-supported beliefs—if any—supervenes only on their internal goings-on. Specifically, their respective communities provide *only* the necessary structure for the two individuals to be in a position to merely form coverage-supported beliefs in the first place—G- or not-G-reliable alike. But whether any of these beliefs will end up being G-reliable indeed is a whole other issue, which depends solely on whether the individuals have managed to calibrate their expectations to the relevant TTTCs appropriately.

Accordingly, even though Goldberg's analysis demonstrates that the existence of coverage-supported beliefs would be impossible in the absence of one's social environment, it fails to reveal the epistemically social nature of such beliefs on the basis of process reliabilism, as it falls short of making the case for the social nature of the doxastic justification (i.e., G-reliability) of coverage-supported beliefs. Nevertheless, there might be an alternative, easier way to account for the epistemically social nature of coverage-supported beliefs on the basis of CAK.

Specifically, drawing on the above discussion, it is obvious that, in cases of coverage support, believing the truth (i.e., believing that *p* when *p* is the case) is due

[25] Another way to put the same point is to note that, in the above quoted passage, Goldberg's talk of the G-reliability of coverage-supported beliefs as being a matter of degree is not consistent with the rest of his analysis. For one thing, TTTCs can be either true or false; accordingly, coverage-supported beliefs can be either G-reliable or not-G-reliable—and not more or less G-reliable. Accordingly, with respect to any two very similar individuals that inhabit communities with different social practices and for any given domain of beliefs, the two individuals will either both be coverage-supported G-reliable in their respective, relevant beliefs *at all*—but they will possess different relevant TTTCs—or they will both have the same TTTCs, but at least one of them will be coverage supported G-*un*reliable in her respective, relevant beliefs.

to the agent's cognitive abilities: It is because of her abilities to calibrate with the relevant informational channels, recall that she has not encountered a report of not-*p* in the past, and draw the relevant inference on some appropriately pitched TTTC, that she comes to believe the truth with respect to *p*. Undoubtedly then, the agent's cognitive success will at least be significantly creditable to her cognitive agency. Therefore, according to $COGA_{weak}$, a subject *can* gain knowledge on the basis of coverage support. At the same time, however, given that in order for the individual to be in a position to draw inferences on the basis of relevant TTTCs, such that coverage supported beliefs can even be part of her doxastic repertoire, the relevant relied-upon source must be in place, part of the credit for the agent's cognitive success must also go to the relevant aspects of her community.

Of course, it may be objected here that, by analogy of reasoning, in cases of knowledge on the basis of, say, vision, we should give credit to *light* (and so on and so forth to every environmental factor that makes the rest of our beliefs possible). There is a crucial difference, however, between this sort of natural causal contribution and the effect of one's community in cases of coverage support; a difference that has to do with the fact that credit attributions have been traditionally associated with *intentional* agents. Accordingly, even though several extra-organismic factors may contribute *causally* in almost every case of belief-formation, credit for the resulting true beliefs should only be directed to those factors that, somehow, contribute *intentionally*. And since in the case of coverage support, the relied-upon source clearly intends to report information on matters that it is being relied upon, part of the credit can plausibly, and quite intuitively, be directed to it.

In other words, given (1) it is clear that one's epistemic community does contribute causally and *intentionally* to the acquisition of coverage-supported true beliefs—even if only as a background, enabling condition and not by ensuring that the subject gets to the truth of the matter—and given (2) we also have a strong intuition that, in such cases, the individual subject does not deserve the entirety of the credit for her true belief, the remaining part of the credit for the ensuing cognitive success should be directed to the relevant relied-upon source.

Accordingly, if the epistemic phenomenon that Goldberg has unearthed obtains—and indeed it appears that it does—we can here finish the job that Goldberg's attachment to process reliabilism appears to prevent him from running to completion. Specifically, if instead of process reliabilism we here follow CAK—according to which knowledge is understood as creditable true belief—and given that the credit for any cognitive success resulting from coverage support must be attributed to both the individual and the community she is relying upon, we can successfully claim that coverage-supported true beliefs constitute another important fragment of our everyday knowledge, which is both social and individual in nature.

4 The Dual Nature of Knowledge on the Basis of Epistemic Artifacts

When unpacked in terms of CAK, testimonial and coverage-supported true beliefs demonstrate the social nature of individualistic knowledge, as in both cases the subject's success is not only creditable to her cognitive agency but also to other

individual epistemic agents who have either offered the reliable testimonial reports that the subject rationally accepts as true, or have contributed to the source that allows the subject to have expectations of coverage and thereby coverage-supported beliefs. In this section, the aim is to demonstrate how CAK can go even further than these two testimony-related cases, by accommodating yet another type of cases that falls under the general heading of weak epistemic anti-individualism.

The cases I have in mind are cases whereby the agent comes to know something on the basis of the operation of some epistemic artifact. Think, for example, of perceiving a chair through a tactile visual substitution system (TVSS),[26] detecting the existence of a specific molecule by using a microscope, or coming to know the position of a satellite on the basis of a telescope. Before moving on to uncovering the social nature of this type of knowledge, however, we must focus on its details first, because, on closer inspection, it could turn out to be a particularly problematic type of knowledge for CAK—especially as motivated by virtue reliabilism.

To see why, consider that, according to the underlying ability intuition on knowledge, belief must be true in virtue of cognitive ability.[27] In the cases we here have in mind, however, the agent's true belief arises out of the interaction of his organismic cognitive abilities with the epistemic artifact: An interaction, during which the artifact plays an *integral* and particularly crucial role with respect to detecting the truth. To make clear how this may be so, think of a well-trained agent, whose telescope has been recently broken. Even though the agent may still be able to form many relevant beliefs (if, say, he still wants to take his chances), none of them will be non-accidentally true, no matter how much he tries. By contrast, think of another agent in possession of a properly working telescope. In this case, not only will the agent be able to form pertinent beliefs, but, unlike the first agent, his beliefs will also come out systematically true. Therefore, it seems that, in such cases, the epistemic artifact is a crucial factor in one's cognitive success that explains not why the agent merely possesses relevant beliefs but why his beliefs come out *true*.

To be clear, this is not to claim that one's internal processes are not a significant factor in how one *forms* one's beliefs. The factor, however, that seems to explain how one *believes the truth of the matter* is the artifact. Accordingly, in order to hold fast to the ability intuition on knowledge, virtue reliabilists must find a way to account for the artifact as being part of the agent's cognitive character.[28] No doubt,

[26] See (Bach-y-Rita and Kercel 2003). Briefly, tactile visual substitution systems consist of a mini video camera that collects visual input, which is then converted into tactile stimulation on the back, tongue or forehead of (usually blind) subjects. On the basis of their practical understanding of sensorimotor contingencies (Noë 2004), subjects can then interact with the device by moving around, which allows them to perceive shapes and objects and orient themselves in space.

[27] Remember that, according to Greco (1999, 2003, 2010), "in virtue of" is supposed to be understood in causal explanatory terms.

[28] To preempt a possible worry here, it is important to note that, in cases of testimony, one cannot really advance a similar argument to the effect that, according to CAK, the speaker must be part of the hearer's cognitive character. In contrast to cases of employing an artifact, in cases of testimony, one's internal capacities are a crucial factor in detecting the truth: It is one's internal capacities that explain how the hearer detects an appropriate informant and rationally accepts the offered report, thereby ending up with the truth of the matter . Of course, if we imagine a parallel to the broken artifact case above, say a case where the hearer is in a testimonially unfriendly environment, such that almost everyone around is a

this sounds like a counterintuitive or even intractable demand that can put virtue reliabilists in a particularly awkward position. Nevertheless, there may actually be a promising way to meet this challenge by invoking the hypothesis of extended cognition.

According to the hypothesis of extended cognition, "the actual local operations that realize certain forms of human cognizing include inextricable tangles of feedback, feedforward and feed-around loops: loops that promiscuously criss-cross the boundaries of brain, body and world (Clark 2007, sec. 2). Cognitive processing, in other words, can, and under the appropriate conditions, literally extends to the agent's surrounding environment. Think about solving a mathematical problem by using pen and paper. According to the hypothesis of extended cognition, the involved artifacts are proper parts of the ongoing cognitive processing.

In fact, it has been previously argued that virtue reliabilism in general and $COGA_{weak}$ in particular are apt for an interpretation along the lines suggested by the hypothesis of extended cognition (Pritchard 2010a; Palermos 2011a, 2014a). First of all, as the reasoning goes, there is nothing in the formulation of $COGA_{weak}$ that restricts knowledge-conducive cognitive abilities to processes within the agent's head. To the contrary, the idea of a cognitive character that may consist of "acquired methods of inquiry including those involving highly specialized training or even advanced technology"(Greco 1999, p. 287) seems to be compatible with or even prefigure the hypothesis of extended cognition (Clark and Chalmers 1998; Clark 2007, 2008; but also see Palermos 2014b).[29]

If we focus on the details of the two theories, however, we can make a much stronger claim. Specifically, both theories put forward the same condition in order for a process to count as part of the agent's cognitive system/character (and, thereby, by the lights of virtue reliabilism, as knowledge-conducive): Just as Greco claims that cognitive integration of a belief-forming process (be it organismically internal or external) is a matter of cooperative interaction with other parts of the cognitive system, so cognitive scientists have recently argued that in order for an artifact to count as a proper part of an agent's cognitive system all that is required is that the two of them be non-linearly related on the basis of ongoing mutual interactions

Footnote 28 continued
compulsive liar, it will indeed be very difficult for the hearer to gain any true beliefs. But if she somehow comes to eventually believe the truth on the basis of a speaker, this does not mean that her cognitive success is not down to her internal cognitive capacities; all the more so for having managed to detect the only person that could provide her with reliable information and for having rationally accepted their words, despite the epistemically hostile setting she finds herself in. In contrast to cases of employing epistemic artifacts, therefore, in cases of testimony, virtue reliabilists do not need to claim that the speaker must be part of the hearer's cognitive character, because, in cases of testimony, it is solely the hearer's internal capacities that explain how the hearer detects the truth.

[29] It is worth noting that the idea of cognitive extension has also been invoked (Vaesen 2011) within the literature in order to argue against CAK. Nevertheless, as Vaesen notes himself, his argument relies on a weak notion of cognitive extension that philosophers of mind would more appropriately categorize under the heading of 'embedded cognition' (see also fn. 31). For responses to Vaesen's arguments see (Kelp 2013) and Greco (2012), and for a rejoinder, see (Vaesen 2013). Moreover, but on a slightly different note, Green (2014) has attempted to wed CAK to the hypothesis of distributed cognition, according to which cognition may not only extend beyond an individual's organism but it may even be distributed between several individuals at the same time, in order to account for team-like epistemic achievements.

(Sutton 2008; Chemero 2009; Froese et al. 2013; Palermos 2014b; Theiner et al. 2010; Tollefsen and Dale 2012; Wegner et al. 1985).[30] Specifically, the claim is that in order to have an extended cognitive system—as opposed to a cognitive system that is merely *embedded* in the sense of being dependent on, but not constituted by, certain environmental aspects (cf. Adams and Aizawa 2001, 2008; Rupert 2004, 2009)—all we need is that the contributing parts (i.e., the relevant cognitive agents and their artifacts) interact continuously and reciprocally with each other.[31]

We see, then, that both virtue reliabilism and the hypothesis of extended cognition put forward the same criterion (viz., cooperative interaction with the rest of the agent's cognitive system) in order for a process to count as integrated into an agent's cognitive system, and thereby as knowledge-conducive. Accordingly, there is no principled theoretical bar disallowing extended belief-forming processes from counting as knowledge-conducive: An agent may extend his cognitive character by incorporating epistemic artifacts to it.[32]

So, for example, in this way, we can explain how a subject might come to know the position of a satellite on the basis of a telescope, while holding fast to the idea that knowledge is belief that is true in virtue of cognitive ability. Even though the belief-forming process in virtue of which the subject believes the truth is for the most part external to his organismic cognitive agency, it still counts as one of his cognitive abilities as it has been appropriately integrated into his cognitive character.[33] Moreover, the subject satisfies $\text{COGA}_{\text{weak}}$, since his overall cognitive success is significantly creditable to his cognitive agency (i.e., his organismic cognitive apparatus): It is the subject's organismic cognitive faculties that are first and foremost responsible for the recruitment and sustaining of the extended belief-

[30] Moreover, it has been elsewhere (Palermos 2011a, 2014a) argued that both theories put also forward the same broad, common-sense functionalist intuitions on what is required from a process to count as a cognitive ability. Briefly, both views state that the process must be (a) normal and reliable, (b) one of the agent's habits/dispositions and (c) integrated into the rest of the agent's cognitive character/system.

[31] To use the standard terminology from philosophy of mind and cognitive science, it is very important to distinguish between the hypothesis of embedded cognition (HEMC) (Rupert 2004, 2009) and the hypothesis of extended cognition (HEC). Invoking what has come to be known as the 'causal-constitution' fallacy (Adams and Aizawa 2001, 2008), according to which proponents of HEC mistake a *causal* for a *constitutive* dependence between the agent and her artifacts, proponents of HEMC suggest that we should rather settle for the less provocative view that the mind is merely embedded rather than extended to its environment. As noted above, however, several proponents of HEC have replied that the presence of non-linear relations between the agent and her artifacts provides a clear criterion for distinguishing between HEC and HEMC as well as putting the 'causal-constitution' fallacy to rest. For an overview of the debate and a detailed approach to how we can distinguish between HEC and HEMC as well as avoid several other worries with respect to HEC (including the 'cognitive bloat' worry and the 'causal-constitution' fallacy), see 2014b.

[32] See also (Alfano 2014, forthcoming) for very interesting discussions on further potential connections between virtue epistemology and the extended cognition and extended mind hypotheses.

[33] Making observations through a telescope can clearly qualify as a case of cognitive extension as it is a dynamical process that involves ongoing reciprocal interactions between the agent and the artifact. Moving the telescope around, while adjusting the lenses, generates certain effects (e.g., shapes on the lens of the telescope), whose feedback *drives* the ongoing cognitive loops along. Eventually, as the process unfolds, the coupled system of *the agent and his telescope* is able to identify—that is, see—the target satellite.

forming process (i.e., telescopic observation) in virtue of which the truth with respect to the satellite's position is eventually arrived at.

In cases like this, therefore, even though it is the external component that accounts for the truth-status of the agent's belief, the agent's cognitive agency—i.e., his organismic cognitive faculties—is still significantly creditable for having appropriately integrated the relevant external component into his cognitive system, and so the agent can, by the lights of $COGA_{weak}$, count as knowledgeable.

Interestingly, however, at this point, the following question may arise: Whereto should the rest of the credit be attributed? This is a fair worry, for as it was argued above, in such cases, the prevailing factor in the causal explanation of the agent's overall cognitive success is the *integrated,* extended belief-forming process that consists of both of one's cognitive agency and the epistemic artifact, operating in tandem. Accordingly, since a significant part of the credit has been attributed to the agent's cognitive agency, should we attribute the rest of the credit to the external aspects of the relevant reliable belief-forming process? That is, should we attribute credit to telescopes, microscopes, computers and so on? It seems that the answer to these questions should be negative.

The reason has to do with a consideration we have already alluded to in the discussion of credit attribution in cases of coverage-support: Even though, as Greco (2003) claims, credit attributions are very much akin to causal explanations, attributions of responsibility, praise, or merely neutral action i.e., attributions of positive, negative or merely neutral credit, respectively) have been traditionally associated with *intentional* agents. Therefore, to attribute credit to artifacts would be a categorical mistake. Notice, however, that artifacts can be defined as objects that have been *intentionally* made or produced for a certain purpose. Accordingly, the remaining credit should be attributed not to the artifacts themselves but to the individuals that intentionally brought the relevant extended belief-forming processes about. Notice further, however, that, frequently, it won't be possible to attribute the rest of the credit to only one single individual, because, in most cases, a (potentially very large) number of individuals has contributed to the development of such reliable belief-forming processes, by having, for example, provided even more belief-forming (sub-) processes or relevant data produced on the basis of the latter. Accordingly, many times, the remaining credit, i.e., the credit that is associated with the external portion of the epistemic agent's extended cognitive ability, will have to *spread* among a large part of the agent's epistemic community.[34]

To make the above more intuitive, let us pause to consider the interesting example of the FIA Formula One Championship. The F1 season consists of a series of races, the results of which are combined to determine *two* Annual World Championships; one for the drivers and one for the constructors. Accordingly, the analogy to be drawn is that the drivers play the role of the cognitive agents and the cars that of the epistemic artifact.[35] Now, what is remarkably to the point is the fact that, according to FIA's rules, the credit for winning cannot be solely attributed to

[34] To be clear, "epistemic community", as it is here intended, refers only to a sum of individual epistemic agents and not to any entity that is over and above that sum.

[35] Driving the car then, plays the role of the overall extended belief-forming process.

the drivers; hence the two championships. Moreover, quite reasonably, the points for the second championship do not go to the cars themselves, but to the constructor's team that *built* the cars. In other words, the credit for winning does not get to be attributed solely to the cognitive agent that drives the car, but, also, to the *team* that brought his racing artifact (i.e., the car) about.

Now, if we think about it, very similar considerations may apply, *mutatis mutandis,* to the case of telescopic observation. In order to come to know the position of the satellite, the astronomer has to rely on an extended belief-forming process that was brought about on the basis of knowledge of long generations of mathematicians, physicists, opticians, machinists, astronomers, computer scientists (depending on the kind of telescope we are talking about) and, in general, a series of experts whose length could go on for a while. Had the astronomer not been part of this epistemic community, and therefore lacked the necessary reliable belief-forming process, he would be incapable of gaining knowledge of the target proposition.

Overall, then, the cognitive success of coming to believe the true position of the satellite is to a significant degree creditable to the particular astronomer—it is *he* who came to know the target proposition by employing the necessary belief-forming process—but the rest of the credit must be attributed to the individuals and in general the epistemic community that brought about the necessary belief-forming process.

If that's true, however, consider how a similar description of the process of gaining propositional knowledge could apply within the fields of mechanics, physics, biology, chemistry, neuroscience and in general any discipline that involves the operation of epistemic artifacts. In order to come to know the truth of some proposition *p,* many times, epistemic agents have to employ reliable belief-forming processes produced by long generations of mathematicians, engineers, experimentalists, scientists, philosophers and many other experts.[36]

In all these cases, therefore, just as in cases of testimonial knowledge and knowledge by coverage support, the epistemic credit for the agent's cognitive success must spread between the individual agent and the relevant parts of her epistemic community. Before closing this section, however, it is perhaps worth pointing out what the difference is between the distribution of epistemic credit that takes place in cases of knowledge on the basis of epistemic artifacts and the spread of credit that occurs in the rest of the cases that the previous sections were dedicated on.

The main contrast has to do with the role that the social factors play in the cognitive ability that the agent employs in order to detect the truth. According to CAK and virtue reliabilism, the agent must not only believe the truth but he must arrive at it on the basis of his cognitive ability. In cases of knowledge by testimony,

[36] For a detailed argument on how we can reveal the social nature of several aspects of scientific knowledge by combining virtue reliabilism with active externalism, see Palermos (forthcoming). For an account of the epistemology of scientific artifacts from a philosopher's of science point of view see (van Fraasen 2001, 2008). See also (Toon 2014).

the agent ends up with a true belief on the basis of her own cognitive capacities that allow her to detect an appropriate informant and then rationally accept the offered report—provided the hearer deems the report reliable or at least not unreliable. The speaker therefore plays only a background role in how the agent forms a true belief, by offering a report that is in fact reliable, and which subsequently forms the content of the hearer's belief. Similarly, in cases of coverage support, it is the agent's internal cognitive abilities that allow her to believe the truth of the matter by appropriately calibrating with the relevant informational channels, recalling that she has not encountered a report of not-*p* in the past, and drawing the relevant inference on some appropriately pitched TTTC. Accordingly, the relevant parts of her epistemic community play again only a background role in how the agent forms a true belief, by making it possible for her to draw inferences on the basis of appropriately pitched TTTCs.

By contrast, in cases of knowledge on the basis of epistemic artifacts, the relevant aspects of the agent's community do not merely play a background role. Admittedly, they affect the agent's cognitive success only *indirectly*, since the agent does not interact with the other individuals themselves, but only with the artifact these individuals produced. Nevertheless, as noted above, in such cases, the artifact plays an *integral* role in how the agent detects the truth of the matter. Put another way, in such cases, the agent does not interact with the relevant external components in a linear, one-way dependence in order, for instance, to merely access information or enable herself to draw further inferences. Instead, the agent detects the truth by interacting with the artifact in a mutual, non-linear way. What this means, according to HEC, is that, in contrast to testifiers or the channels that support coverage support, in such cases, the external component does not play just a causal but a *constitutive* role in very process that is responsible for the agent's cognitive success. Accordingly, when the agent knows on the basis of epistemic artifacts that she has appropriately integrated to her cognitive character, the epistemic community plays a much deeper role—even if indirectly so—and thereby deserves a bigger part of the epistemic credit in comparison to cases of knowledge by testimony or coverage-support.

So with that said, and to return to our overall argument, the upshot of this section is this: In all cases where the agent comes to believe truly on the basis of a belief-forming process that extends to some epistemic artifact, the knower *is* the individual, because the cognitive success is to a significant degree creditable to her cognitive agency for appropriately employing the relevant reliable belief-forming process that she has integrated into her cognitive character. Crucially, however, the rest, and perhaps the greatest part of the credit for the agent's cognitive success will have to be attributed to the individuals, and, in general, to the epistemic community that brought the relevant extended belief-forming process about. Therefore, and since according to the CAK, as captured by virtue reliabilism, knowledge is creditable true belief, knowledge acquired on the basis of epistemic artifacts is yet another type of knowledge that is both social and individual in nature, in a rather substantial sense.

5 Conclusion: Epistemic Dependence and Weak Epistemic Anti-individualism

So far, it has been argued that individual knowledge, understood in terms of creditable true believing, often appears to be also social in nature. In particular, it was demonstrated that CAK cannot only account for the partly social nature of testimonial and coverage-supported knowledge, but it can also reveal the social nature of knowledge arrived at via the operation of epistemic artifacts. Admittedly, there are probably several other types of knowledge that are both social and individual in nature, but the above three cases represent a particularly significant and diverse set of a social subject's overall means for acquiring knowledge. Accordingly, and contrary to any of the previous attempts to unravel the particular ways in which specific types of knowledge are partly social in nature, CAK looks very promising as a candidate for providing a *unified* account of all the possible ways in which weak epistemic anti-individualism may be instantiated.

Specifically, despite their variegation, it was argued that what all the above cases have in common is that they all satisfy the following claim: The subject's cognitive success is *significantly* creditable to her cognitive agency and thus, according to $\text{COGA}_{\text{weak}}$, the individual can be knowledgeable. What is also true in all these cases, however, is that the rest of the credit must go to one or more individual members of the individual subject's epistemic community. This is so, either because those individuals offer reliable reports, form epistemic channels on which the subject can rely on for her coverage-supported beliefs, or have produced some reliable belief-forming process that the subject can integrate within her cognitive character so as to reliably form true beliefs. Therefore, to repeat the claim, if knowledge is to be understood in terms of creditable true believing and if, in all the above diverse cases, significant part of the credit must be attributed both to the individual subject and the epistemic community of which the subject is a proper part, then individual knowledge, in all these cases, turns out to be also social in nature, thereby clearly qualifying as an instance of weak epistemic anti-individualism.[37]

[37] It should be noted that there is no problem with the act of epistemically grouping all these cases in such a way, even if, in addition to all the obvious differences of mere physical implementation between the above belief-forming processes, they also differ in other *epistemically* relevant ways. For example, one epistemically relevant way to differentiate between the above processes is to accentuate the role that one's society plays with respect to the relevant mechanisms of forming one's belief: Specifically, given the above analysis, in cases of coverage-supported beliefs, one's society plays merely the role of an enabling condition for one's belief-forming process to exist; in cases of testimony, one's society provides one with the information one ends up believing; and in cases of knowledge on the basis of epistemic artifacts, one's society is the source of one's belief-forming process itself. Similarly, another epistemically relevant way to differentiate between these cases is in terms of how one's evidence is associated with one's epistemic community: In coverage support, the agent relies on the community for *not* providing relevant evidence; in cases of testimony the agent relies on the society for providing him with relevant evidence; and in cases of instrument-mediated belief, the agent relies on the community having built some instrument, indicating that, occasionally, there must be relevant evidence 'out there' to be accessed via the target instrument. Despite all these possible ways to (epistemically) differentiate between these cases, however, the fact remains: In all of them, the agents' cognitive successes are partly

But does this partly social nature of knowledge suggest, as some epistemologists have argued, that we should stop considering the individual as the proper object of our epistemological inquiries? To answer this worrying question, let us briefly go through what Hardwig thinks concerning the following case:

A knows that *m*
B knows that *n*
C knows (1) that *A* knows that *m*, and (2) that if *m*, then *o*
D knows (1) that *B* knows that *n*, (2) that *C* knows that *o*, and (3) that if *n* and *o*, then *p*.

Having the above in mind, Hardwig writes: "Suppose that this is the only way to know that *p* and, moreover, that no one who "knows" that *p* knows that *m*, *n* and *o* except by knowing that others know them" (1985, p. 348). If that's the case, then we must agree that either "one can know without possessing the supporting evidence [for the truth of the relevant proposition] or accept the idea that there is knowledge that is known by the community, not by any individual knower" (Hardwig 1985, p. 349). In some more detail, according to Hardwig, cases such as the above demonstrate that knowledge and intellectual autonomy—in the sense of being self-sufficient in possessing the necessary evidence (such as sound arguments and factual information) for the truth of one's belief—should come apart. Otherwise, if the link between knowledge and intellectual autonomy is to be preserved then knowledge, in cases such as the above, is not possessed by any individual alone, but by the relevant community as a whole, because it is that community alone that can be intellectually autonomous in possessing the relevant evidence.

Of course, the latter option is obviously problematic for mainstream epistemology and the methodological individualism that underlies it. Hardwig, however, assumes that the former option of abandoning intellectual autonomy as a prerequisite for knowledge is equally unpalatable for individualism and mainstream epistemology. But how so?

One possible answer is that Hardwig appears to implicitly rely on a popular internalist understanding of knowledge and justification, according to which *S*'s reasons for his true beliefs must be *accessible* to him by reflection alone.[38] Consequently, on the basis of this, Hardwig also comes to believe that the epistemic agent must be intellectually autonomous and that, crucially, a rejection of this requirement signifies a departure from mainstream epistemology and the underlying methodological individualism.

This doesn't have to be so, however. For one, it is possible that such an interpretation of epistemic internalism is too strong and that not all formulations of epistemic internalism entail that one must be intellectually autonomous in

Footnote 37 continued
creditable to their epistemic communities, allowing them, according to CAK, to all qualify as clear cases of weak epistemic anti-individualism.

[38] This is the standard formulation of internalism, known as accessibilism. Roughly stated: Whenever one knows that *p,* then one can become aware by reflection of one's knowledge basis for *p*. For more details, see (BonJour 1985; Chisholm 1977; Steup 1999).

possessing evidence for the truth of the relevant proposition. In other words, it is an open question whether internalism can account for the fact that one can have no evidence with respect to *p* but still know that *p,* simply by knowing that someone else knows that *p*.[39] But even if this is impossible, we can still bring into question the epistemic internalism that Hardwig seems to implicitly rely upon.

It is a happy incident, then, that virtue reliabilism in general and $COGA_{weak}$ in particular are externalist, though mainstream approaches to knowledge all the same, which have been proposed as alternatives to the classical internalist account. This is a happy incident, because $COGA_{weak}$ can easily accommodate Hardwig's first and arguably rather intuitive conclusion that the individual *is* a proper epistemic agent, *even though not autonomously so*. To see how this might be, consider that, according to epistemic externalism, in order for one's true beliefs to count as knowledge, they need not be backed up by reasons that one should in principle have introspective access to. And by so denying the demand for introspective access to one's justification for one's beliefs, epistemic externalism also makes the demand for intellectually autonomy appear under-motivated.

$COGA_{weak}$, however, takes these considerations a step further. In particular, by allowing knowledge to be acquired merely on the basis of reliable belief-forming processes, such that the cognitive success can only be *significantly* creditable to one's cognitive agency, it actually anticipates—if not ascertains—the denial of intellectual autonomy. Either the presence of the relevant input to, the enabling conditions for, or even the very existence of the belief-forming process itself may heavily depend on one's epistemic community (or at least parts of it), and this is a fact that can be explicitly accommodated: $COGA_{weak}$ allows (the rest of the) credit to be attributed to those exogenous (or rather extra-organismic) epistemic factors as well. Nevertheless, at the same time, it recognizes the individual as the proper epistemic subject and stresses his/her importance by demanding that the cognitive success be significantly creditable to his/her cognitive agency.

Overall, then, $COGA_{weak}$ provides a way out of Hardwig's dilemma: The phenomenon of epistemic dependence does not mean that we need to abandon epistemic individualism by admitting that there can be knowledge that is not known by any individual alone, but by the relevant community as a whole. Instead, $COGA_{weak}$ allows mainstream epistemologists to embrace Hardwig's first conclusion that one can know a proposition *p*—even if not autonomously so—all the while remaining well in line with methodological individualism.

It should be made clear, however, that the above is not to claim that all externalist epistemology points away intellectual autonomy and *strong individualism* (as we may call the view that knowledge should be fully down to the individual). Having the reasons of one's justification out of one's reach of awareness is certainly not the same as partly having those reasons out of one's bodily boundaries. For one thing,

[39] Of course this is going to be no easy task, and its feasibility will largely depend on whether one is a reductionist or anti-reductionist about testimonial knowledge. While it is beyond the scope of the present paper to expand on this issue, on one hand, the combination of epistemic internalism and reductionism about testimonial justification is rather problematic, because internalism has a hard time accounting for inductive knowledge (Greco 1999). On the other hand, *prima facie*, epistemic internalism and anti-reductionism about testimonial justification is a rather unfitting match. See also fn. 40.

there can certainly be externalist conditions on knowledge, which are individualistic in nature. Take for example Lackey's (2007) interpretation of Greco's (externalist) Agent Reliabilism, according to which the cognitive success must be *primarily* creditable to *S*'s cognitive character. Such an account can clearly qualify as a form of *weak individualism*; and, if instead it demanded that one's cognitive success be *solely* creditable to one's cognitive agency, it would even qualify as a case of *strong individualism*, while being an externalist approach to knowledge all the same. Therefore, the internalism/externalism distinction is by no means the same as the individualism/anti-individualism distinction. And even though internalist conditions on knowledge are likely to be always tied to strong individualism,[40] externalist conditions appear to come in degrees, with the potential to occupy any available area on the continuum that the individualism/anti-individualism distinction defines.

$COGA_{weak}$, however—with its lenient demands on the creditability of the cognitive success to one's cognitive agency—is able to account for the *whole* spectrum of the continuum defined by the individualism/anti-individualism distinction. So far, we have been liberated from the demand of intellectual autonomy and the concomitant 'duopoly' of strong and weak epistemic individualisms, by pointing out how, on the basis of $COGA_{weak}$, knowledge can be creditable to both some individual and the society of which he/she is a part, thereby arguing for what we have here called *weak epistemic anti-individualism*. Of course, this leaves open the even more liberal possibility of *strong epistemic anti-individualism,* according to which certain instances of knowledge may be *entirely* social in that they may be creditable only to a group of people *as a whole*—a rather interesting dialectical possibility, which is diametrically opposite to strong individualism, and about which more needs to be said in the future.

For now, however, it suffices that we have here taken the first steps towards an outline of how mainstream epistemology can provide a unified account of perhaps all the possible cases whereby our seemingly individualistic knowledge turns out to be social in nature as well: According to CAK, on a multitude of diverse occasions, the cognitive success of having a true belief turns out to be creditable not just to the relevant individual, epistemic agent but to his/her epistemic community as well. At the same time, however, in all these cases, the individual agent's central role is also clearly acknowledged, thereby allowing for mainstream epistemology's methodological individualism to be applied in such weakly anti-individualistic cases equally well.

Acknowledgments I am very thankful to Adam Carter and Duncan Pritchard for feedback in previous drafts of this paper. I am also grateful to two anonymous referees for *Erkenntnis* for their very insightful comments. Research into the area of this paper was supported by the AHRC-funded 'Extended Knowledge' Project, based at the Eidyn research centre, University of Edinburgh.

[40] Perhaps, it would be more accurate to say here that internalist conditions on knowledge are likely to be tied to strong individualism, *are they not to be transformed beyond recognition.* For an extended discussion of how both access and mentalist (Conee and Feldman 2001) internalism can be given an anti-individualistic reading, see (Carter and Palermos 2014).

References

Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology, 14*(1), 43–64.

Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Oxford: Blackwell Publishing Ltd.

Alfano, M. (2014). What are the bearers of virtues?. In H. Sarkissian & J. Wright (Eds.), *Advances in moral psychology*. Continuum.

Alfano, M. (forthcoming). 'The embedded and extended character hypotheses'. In J. Kiverstein (Ed.), *Philosophy of the social mind*. Routledge.

Alston, W. P. (2006). *Beyond "justification": Dimensions of epistemic evaluation*. Ithaca, NY, Bristol: Cornell University Press.

Audi, R. (1998). *Epistemology: A contemporary introduction to the theory of knowledge*. London; New York: Routledge.

Bach-y-Rita, P., & Kercel, S. W. (2003). Sensory substitution and the human–machine interface. *Trends in Cognitive Sciences, 7*(12), 541–546.

BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge, MA: Harvard University Press.

Burge, T. (1986). Individualism and psychology. *The Philosophical Review, 95*(1), 3. doi:10.2307/2185131.

Carter, J. A., & Palermos, S. O. (2014). Active externalism and epistemic internalism. *Erkenntnis*. doi: 10.1007/s10670-014-9670-5.

Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.

Chisholm, R. M. (1977). *Theory of knowledge*. Englewood Cliffs NJ: Prentice-Hall.

Clark, A. (2007). Curing cognitive hiccups: A defense of the extended mind. Retrieved from http://www.era.lib.ed.ac.uk/handle/1842/1719.

Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford, New York: Oxford University Press.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19. doi:10.1093/analys/58.1.7.

Conee, E., & Feldman, R. (2001). Internalism defended. In H. Kornblith (Ed.), *Epistemology: Internalism and externalism*. Cambridge: MIT Press.

Dretske, F. I. (1970). Epistemic operators. *Journal of Philosophy, 67*(24), 1007–1023.

Faulkner, P. (2000). The social character of testimonial knowledge. *The Journal of Philosophy, 97*(11), 581–601.

Fricker, E. (1994). Against gullibility. In B. K. Matilal & A. Chakrabarti (Eds.), *Knowing from words* (pp. 125–161). Springer Netherlands. Retrieved from http://link.springer.com/chapter/10.1007/978-94-017-2018-2_8.

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford: Oxford University Press.

Froese, T., Gershenson, C., & Rosenblueth, D. A. (2013). The dynamically extended mind—a minimal modeling case study. arXiv:1305.1958 [nlin]. Retrieved from http://arxiv.org/abs/1305.1958.

Fuller, S. (2007). *The knowledge book key concepts in philosophy, science, and culture*. Stocksfield: Acumen.

Fuller, S. (2012). Social epistemology: A quarter-century itinerary. *Social Epistemology, 26*(3–4), 267–283. doi:10.1080/02691728.2012.714415.

Goldberg, S. (2010). *Relying on others: An essay in epistemology*. Oxford: Oxford University Press.

Goldberg, S., & Henderson, D. (2006). Monitoring and anti-reductionism in the epistemology of testimony. *Philosophy and Phenomenological Research, 72*(3), 600–617.

Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: D. Reidel.

Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.

Goldman, A. I. (2004). Group knowledge versus group rationality: Two approaches to social epistemology. *Episteme, 1*(01), 11–22. doi:10.3366/epi.2004.1.1.11.

Goldman, A. (2010). Why social epistemology is real epistemology. In A. Haddock, A. Millar & D. Pritchard (Eds.), *Social epistemology* (pp. 1–29). Oxford: Oxford University Press.

Greco, J. (1999). Agent reliabilism. *Noûs, 33*, 273–296. doi:10.1111/0029-4624.33.s13.13.

Greco, J. (2003). Knowledge as credit for true belief. *Intellectual virtue: Perspectives from ethics and epistemology* (pp. 111–134). Oxford: Oxford University Press.

Greco, J. (2007). The nature of ability and the purpose of knowledge. *Philosophical Issues, 17*(1), 57–69. doi:10.1111/j.1533-6077.2007.00122.x.

Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
Greco, J. (2012). A (different) virtue epistemology. *Philosophy and Phenomenological Research, 85*(1), 1–26.
Green, A. (2012). Extending the credit theory of knowledge. *Philosophical Explorations, 15*(2), 121–132.
Green, A. (2014). Evaluating distributed cognition. *Synthese, 191*(1), 79–95.
Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy, 82*(7), 335–349.
Hume, D. (2011). *Enquiry concerning human understanding*. [S.l.]: Simon & Brown.
Kelp, C. (2013). Extended cognition and robust virtue epistemology. *Erkenntnis, 78*(2), 245–252.
Kusch, M. (2002). *Knowledge by agreement: The programme of communitarian epistemology*. Oxford: Oxford University Press.
Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese, 158*(3), 345–361.
Lackey, J. (2008). *Learning from words: Testimony as a source of knowledge*. Oxford, New York: Oxford University Press.
Michaelian, K. (2014). JFGI: From distributed cognition to distributed reliabilism. *Philosophical Issues, 24*(1), 314–346.
Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.
Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
Palermos, S. O. (2011a). Belief-forming processes, extended. *Review of Philosophy and Psychology, 2*(4), 741–765.
Palermos, S. O. (2011b). Dualism in the epistemology of testimony and the ability intuition. *Philosophia, 39*(3), 597–613.
Palermos, S. O. (2014a). Knowledge and cognitive integration. *Synthese, 191*(8), 1931–1951.
Palermos, S. O. (2014b). Loops, constitution, and cognitive extension. *Cognitive Systems Research, 27*, 25–41.
Palermos, S. O. (*forthcoming*). Active externalism, virtue reliabilism and scientific knowledge. *Synthese*. doi: 10.1007/s11229-015-0695-3.
Palermos, O., & Pritchard, D. (2013). Extended knowledge and social epistemology. *Social Epistemology Review and Reply Collective, 2*(8), 105–120.
Plantinga, A. (1993). *Warrant and proper function*. New York: Oxford University Press.
Pritchard, D. (2002). Resurrecting the Moorean response to the sceptic. *International Journal of Philosophical Studies, 10*(3), 283–307. doi:10.1080/09672550210152122.
Pritchard, D. (2006). *What is this thing called knowledge?*. London, New York: Routledge.
Pritchard, D. (2008). Sensitivity, safety, and anti-luck epistemology. In J. Greco (Ed.), *The Oxford handbook of skepticism*. Oxford: Oxford University Press.
Pritchard, D. (2009). *Knowledge*. Palgrave Macmillan. Retrieved from http://www.abebooks.co.uk/Knowledge-Palgrave-Philosophy-Today-Professor-Duncan/10077064975/bd.
Pritchard, D. (2010a). Cognitive ability and the extended cognition thesis. *Synthese, 175*(1), 133–151. doi:10.1007/s11229-010-9738-y.
Pritchard, D. (2010b). Knowledge and understanding. In A. Haddock, A. Millar & D. Pritchard (Eds.), *The nature and value of knowledge*. Oxford: Oxford University Press.
Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy, 109*, 247–279.
Reid, T. (1786). *Essays on the intellectual powers of man*. Dublin: Printed for L. White. Retrieved from http://archive.org/details/essaysonintellec02reiduoft.
Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy, 101*(8), 389–428.
Rupert, R. D. (2009). *Cognitive systems and the extended mind* (1st ed.). New York: Oxford University Press.
Shieber, J. (2012). Against credibility. *Australasian Journal of Philosophy, 90*(1), 1–18.
Shieber, J. (2013). Toward a truly social epistemology: Babbage, the division of mental labor, and the possibility of socially distributed warrant. *Philosophy and Phenomenological Research, 86*(2), 266–294.
Sosa, E. (1988). Beyond scepticism, to the best of our knowledge. *Mind, XCVII*(386), 153–188. doi:10.1093/mind/XCVII.386.153.
Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge: Cambridge University Press.
Sosa, E. (1999). How to defeat opposition to Moore. *Noûs, 33*, 141–153. doi:10.1111/0029-4624.33.s13.7.

Sosa, E. (2000). Skepticism and contextualism. *Philosophical Issues, 10*(1), 1–18. doi:10.1111/j.1758-2237.2000.tb00002.x.

Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge*. Oxford: Clarendon Press.

Sosa, E. (2011). *Reflective knowledge: Apt belief and reflective knowledge* (Vol. II). Oxford: Oxford University Press.

Steup, M. (1999). A defense of internalism. In L. Pojman (Ed.), *The theory of knowledge: Classical and contemporary readings*, 2nd edition. Belmont, CA: Wadsworth Publishing.

Sutton, J. (2008). Between individual and collective memory: Interaction, coordination, distribution. *Social Research*, *75*(1), 23–48.

Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research, 11*(4), 378–395. doi:10.1016/j.cogsys.2010.07.002.

Tollefsen, D., & Dale, R. (2012). Naturalizing joint action: A process-based approach. *Philosophical Psychology, 25*(3), 385–407. doi:10.1080/09515089.2011.579418.

Toon, A. (2014). Empiricism for cyborgs. *Philosophical Issues, 24*(1), 409–425.

Vaesen, K. (2011). Knowledge without credit, exhibit 4: Extended cognition. *Synthese, 181*(3), 515–529.

Vaesen, K. (2013). Critical discussion: Virtue epistemology and extended cognition: A reply to Kelp and Greco. *Erkenntnis, 78*(4), 963–970.

Van Fraassen, B. C. (2001). Constructive empiricism now. *Philosophical Studies, 106*(1/2), 151–170.

Van Fraassen, B. C. (2008). *Scientific representation: Paradoxes of perspective*. Oxford: Oxford University Press.

Wegner, D. M., Giuliano, T., & Hertel, P. T. (1985). Cognitive interdependence in close relationships. In D. W. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253–276). Springer New York. Retrieved from http://link.springer.com/chapter/10.1007/978-1-4612-5044-9_12.

Weiner, M. (2003). Accepting testimony. *The Philosophical Quarterly, 53*(211), 256–264. doi:10.1111/1467-9213.00310.