Synthese (2014) 191:79–95
DOI 10.1007/s11229-013-0305-1

Evaluating distributed cognition

Adam Green

Received: 8 July 2012 / Accepted: 1 June 2013 / Published online: 12 June 2013

Abstract Human beings are promiscuously social creatures, and contemporary epistemologists are increasingly becoming aware that this shapes the ways in which humans process information. This awareness has tended to restrict itself, however, to testimony amongst isolated dyads. As scientific practice ably illustrates, information-processing can be spread over a vast social network. In this essay, a credit theory of knowledge is adapted to account for the normative features of strongly distributed cognition. A typical credit theory analyzes knowledge as an instance of obtaining success because of or through the ability of the individual knower. The extended credit theory developed here broadens this framework so as to accommodate team-like epistemic achievements. The extended credit theory is then contrasted with some similar proposals given from within a process reliabilist framework. Once one isolates pairs of cases of distributed cognition in which there is a difference between sheer reliability and reliability grounded in ability, one can see that the extended credit theory maps the normative terrain better than the alternatives.

Keywords Distributed cognition · Virtue epistemology · Extended reliabilism

Analytic epistemology has traditionally focused on individuals and the epistemic goods to which individuals can aspire under their own power. As the explosion of interest in testimony has demonstrated, however, there is more to being a human knower than the traditional focus of epistemology might lead one to suspect.[1] In this paper, I will focus on what I will call distributed cognition, and I will argue that a modified credit theory of knowledge can do a better and more complete job of mapping

[1] One might date the shift in attention to Coady (1992) and the responses it elicited.

A. Green (✉)
University of Innsbruck, Sankt Nikolaus Gasse 27, 6020
Innsbruck, Austria
e-mail: greenab3@gmail.com

the normative terrain than competing theories. I will begin by developing the view in interaction with John Greco's agent reliabilism and then I will compare it to potential competitors put forward by Sandy Goldberg and Joe Shieber.

1 Describing distributed cognition

Let us begin with a series of cases that will prove useful in illustrating the phenomena of interest. Though I do not take distributed cognition to be restricted to scientific enquiry, some generic examples of this brand will be helpful.

> Case 0: Bob invents a field of research from scratch without relying on anyone else.
>
> Case 1: Josephine is a maverick researcher. She reads all the relevant literature in her field, but she does not accept anyone else's word on their say so. She double checks anyone's work that she uses in her own research. She does make more progress than she would otherwise be capable of due to what she learns from the successes and failures of others, but any successes Josephine has are the result of information and skills in her possession.
>
> Case 2: Charles does research that depends upon the viability of an established body of research. Though he has personally verified only a very small proportion of the empirical or mathematical work that undergirds the research paradigms within which he operates, he is sensitive to the way in which research is regarded by the relevant subset of the scientific community. He is careful not to base any of his research on assumptions or methods that would not be widely considered acceptable in the relevant scientific circles.
>
> Case 3: Lydia is part of a research group at the cutting edge of a new area of research that intersects several extant fields. Though Lydia is sensitive to the values, standards, and assumptions that characterize her own subfield, she is not similarly sensitive to the state of other fields. The intersection of the different subfields is enough to guarantee the quality of the research, but neither Lydia nor any other single person who can give Lydia testimony is in a position to distinguish this case from similar cases in which an inferior output is the outcome of the collective efforts of the different researchers.[2]

Cases 0–3 illustrate an increasing degree of social distribution of the processing of information. Each case lies further on a continuum of which case 0 is one terminus and the other terminus is well past the position of case 3. In this section, I want to draw attention to the applicability of the evaluative terms of virtue epistemology to these cases, especially cases far from 0. To borrow the terminology of Ernest Sosa (Cf. Sosa 2007, pp. 22–43), one can ask in each case whether the output of the research is accurate, that is, whether it succeeds in its aim. One can ask whether the output is adroit, whether it manifests skill. And one can ask whether the output is apt, which is being accurate because adroit and thus "creditable" to the agent manifesting the skill (Sosa 2007, p. 22). To use an alternative expression from Sosa, one can evaluate the

[2] I do not mean to invoke the technical modal notion of sensitivity here. Rather, I mean something more colloquial such as that changes in the subfield often result in matching changes in the agent.

outputs of each of these cases for whether and to what extent the research is getting things right due to the manifestation of competence.[3] To use John Greco's preferred nomenclature, one can in each case evaluate the scientific output for whether it is a success that is through or because of ability (Cf. Greco 2010, p. 3).

These evaluations can be applied at several levels. One can evaluate the non-mental informational states produced such as the graphs, numerical data, and linguistic descriptions brought about. One can instead focus one's attention on the status of group mental states if one thinks there are such things. One can also evaluate the mental informational states of individual researchers and consumers of research. Much of the work in cognitive science, philosophy of science, or cognitive anthropology that has concerned itself with the social distribution of information has concerned itself with the first two evaluative levels. In his classic study of ship navigation, for example, Edwin Hutchins' focus is on the limitations of individual members of the navigation team and on the cognitive work that is done by information technologies and ways of organizing a team (Hutchins 1995).

In this essay, I will focus on the third level, the level of the individual consumer of distributed cognition, and I will do so for two reasons. First, focusing on this level is useful for showing how distributed cognition informs traditional epistemic concerns. It would otherwise be open to an epistemologist to claim, for instance, that the normative dimension of public representations and that of belief do not intersect in any interesting manner. It is helpful, then, to show that distributed cognition changes what one should say about something epistemologists were already talking about.

Second and perhaps of more interest to non-epistemologists, it is not guaranteed that excellences acquired at a group level will be transferrable to individuals. Thus, it is a substantive and interesting question how it could be that individuals could derive an epistemic benefit from distributed cognition. For instance, one of the themes of Christian List and Philipp Pettit's recent work is that rational agents can produce irrational groups and that restricting certain epistemic goods at the individual level can have positive epistemic results at the group level (Cf. List and Pettit 2011). Likewise, Ronald Giere is careful to qualify his work on distributed cognitive processes by claiming that individual agents must still be considered the epistemic locus of states like knowing (cf. Giere 2012, 2006, pp. 96–116). Thus, since the relation between goods at the group level and the individual level cannot be guaranteed, it is useful to have a picture of how one might come to know something as an individual cognizer in virtue of distributed cognition.

Turning our attention to the individual level, then, a little reflection suffices to show that groups can provide cognitive benefits to an individual member of a group that are hard to account for using traditional epistemology. Let us back our way into understanding how that is so by first attending to how an individual might fail to appropriate excellence in information-processing that obtains at the group-level. Consider an elaboration of case 3. Suppose that the work that Lydia and her colleagues are doing is skillful qua information-processing and that the cooperative research derives the result that p. Suppose further that, were one person to possess all the information and skills

[3] These expressions are realist, but I take it that one could adapt the language so as to make similar points within the framework of anti-realism in the philosophy of science.

brought to bear by Lydia and her colleagues in the broader scientific community in the manner of case 1 or 2, then that person would be in a good position to know that p by the lights of traditional epistemologies. Nonetheless, Lydia might easily not be in a position to know that p by the lights of a traditional epistemology, nor to receive testimony from anyone that counts as knowing that p by traditional lights.

One obvious reason that Lydia could fail to know that p that applies anywhere on the continuum of distributed cognition is that she might fail to believe that p. Establishing that p through empirical investigation does not entail that one will believe what one has established. Another reason that applies throughout the continuum is that Lydia might have an undefeated defeater that diminishes the epistemic standing of her belief to something short of knowledge. It could be, for example, that misleading empirical evidence that does not figure directly in her research undercuts her belief or that strongly held phenomenological, moral, or religious considerations conflict with the belief that p in a way that precludes her belief that p amounting to knowledge.

Despite there being general considerations of this type that allow the epistemic evaluations of research and researcher to diverge, there are special reasons involving the distribution of cognition that are worth attending to as well. For instance, there may be defeaters to which Lydia has access that are themselves defeated by information possessed by other members of the extended research community. It could be that the structure of the cooperative whole of the extended research community is such that Lydia is walled off, cognitively, from what would defeat her defeater in the sense that she has conscious access to her defeater but she does not have conscious access to the defeater of the defeater. Thus, it could be that the output of the research could be protected from the defeater that bothers Lydia due to the possession of a defeater defeater by others, while Lydia, nonetheless, is prevented from having full-blown knowledge that p.

Notice, however, that Lydia could easily find herself cognitively walled off from positive grounds for p in a way exactly parallel to the case of the defeater defeater. In fact, one would expect this to be the case often when what is at issue is strongly distributed cognition, such as in case 3. As we shall see in the next section, it is this scenario that puts pressure on more traditional, individualistic accounts. The nature of the social network can enhance the individual's reliability in a way not countenanced by the traditional focus on the individual.

Suppose, for instance, that Lydia is conducting research at a time when global warming was an emerging theory for a wide variety of phenomena impinging on the specializations of many scientists whose work would otherwise not interact. I do not propose to be doing real history of science through this example. Rather, it is useful to imagine a phenomenon such as global warming that has an evidential footprint in many different kinds of empirical data. The time that one should bear in mind is one prior to the theory being widely recognized as a consensus opinion of the scientific community. Suppose, furthermore, that the narrow specialization of Lydia and her team relates all and only to tracking fluctuations in solar energy hitting the earth's atmosphere. It could have easily been the case that Lydia's results, produced at a time when global warming was only an emerging theory, could be taken up into the information processing of a distributed scientific community so as to produce grounds that would be sufficient for knowledge if cognitively possessed by one individual. This

could have been the case while it was yet true that the grounds to which Lydia had access at the time were insufficient for her to know that global warming was true by the lights of traditional epistemologies. I take it to be obvious that this is a perfectly generalizable scenario. It can easily be the case that one is cognitively walled off from some of the features of a social network that enhance the reliability of one's beliefs despite the belief being describable in the terms of virtue epistemology.[4]

In the next section, I want to explore the ability of virtue epistemology to accommodate the normative features of distributed cognition with special attention to cases of strongly distributed cognition as exemplified by the Lydia case. We will focus here on Greco's development of the view.[5]

2 Greco and testimony: the need for an extended credit view

For Greco, knowledge is an instance of something familiar. To know is to succeed in the cognitive domain because of or through ability. It is for the truth of one's belief to be an intellectual achievement, and, equivalently, to know is for the having of a true belief to redound to one's credit, for it to be a creditable success. This account of knowledge subsumes epistemic normativity into the familiar genus that governs skillful endeavours of all kinds from baseball to calculus to cooking (Cf. Greco 2003, 2010, pp. 3–4, 7).

Greco's account can take advantage of an interesting axiological fact about other instances of this familiar kind (Greco 2010, pp. 97–98). One values succeeding through ability more than one values success simpliciter. A good soufflé that is the result of skill is generally more valuable to the one who cooks it than one that comes out right through a fortuitous accident, not because it tastes better per se but because we would rather produce a soufflé through skill than by accident. Winning a tennis match through skill is more valuable to us than winning through biased refereeing. Arguably, doing the right thing on purpose is more valuable to us than stumbling into the action that happens to be required morally. This general axiological pattern coheres well with the putative value of knowledge. We value knowledge more than mere true or Gettierized belief, despite there being no clear difference in instrumental value between these different states.

In each of the cases on the continuum of cognitive distribution we examined previously, one can describe what happens in terms of success in information processing

[4] One does not have to endorse the existence of collective minds or the extended mind to accept that there is such distributed cognition. It is enough to think that social networks, whatever they are, exist and that individuals interface with them. One can think of the relationship as being one of mind to cognitive niche. Thanks to an anonymous referee for drawing my attention to the need to clarify this.

[5] I will focus primarily on interacting with Greco's position. This is useful because Greco's view can be contrasted sharply with the view I want to develop, though I take it that Greco's thinking on this aspect of his theory is continuing to evolve. Sosa, by contrast, is a bit less clear when it comes to what role other people do or do not play in cases like that of testimony. He claims that it is enough if an agent receives partial credit for a testimonial belief and that credit can be distributed through a group (Sosa 2007, pp. 92–97). He is not very clear, however, on what the theoretical substance of this move is. As contrasted with Sosa, one could think of this project as trying to draw out and make explicit what Sosa's suggestion would commit one to if it were doing substantive theoretical work.

that is or at least should ideally be the result of ability. In each case, we end up with a success that seems valuable in a way that surpasses the value of lucky success. Thus, it is worth investigating how Greco's views might accommodate the normative properties that can be supported by distributed cognition. It is instructive, in this regard, to look at what Greco has to say about testimony.

Critics of Greco's view and credit theories of knowledge in general have often brought up the example of testimony as a supposed counterexample. What might be the most widely discussed case comes from Jennifer Lackey.

> Having just arrived at the train station in Chicago, Morris wishes to obtain directions to the Sears Tower. He looks around, approaches the first adult passer-by that he sees, and asks how to get to his desired destination. The passer-by, who happens to be a Chicago resident who knows the city extraordinarily well, provides Morris with impeccable directions to the Sears Tower by telling him that it is located two blocks east of the train station. Morris unhesitatingly forms the corresponding true belief. (Lackey 2007, p. 352)

Lackey presses the question of who deserves credit for Morris' success. The obvious answer, according to Lackey and other opponents of the credit theory, is that the testifier deserves the credit and that Morris does not. It is the testifier's abilities that explain why Morris forms a true belief about the location of the Sears tower. What is of interest in this context, however, is the nature of Greco's response to the objection.

Greco describes his virtue theoretic approach to testimony as one on which "testimonial knowledge requires that the believer is a reliable receiver of testimony". He explains that this "will plausibly involve reliable capacities for discriminating reliable sources of testimony from unreliable sources" (Greco 2010, p. 81). As a result, Greco claims that the Morris case is under-described. One needs to know whether Morris's success is "grounded in his ability to discriminate good from bad testimony" (ibid). Moreover, if Morris's success is conditional on his ability to discriminate between sources of testimony, then it is plausible that his abilities are salient enough for his true belief to count as a creditable success. The skills of the testifier create a friendly environment for achievement, but so long as Morris employs skill that explains his success, he gets credit.[6]

Can Greco's response show us a way in which his theory can account for the normative features of distributed cognition? Not unmodified. Let us give some formal structure to a scenario in which Lydia comes to believe, at a time when global warming was emerging as a theory, that global warming is occurring and in which Lydia's research relates all and only to measuring the solar energy that reaches the atmosphere, an insufficient basis by itself for knowing that global warming is occurring. Suppose that Lydia is sensitive to the standards, values, methods, live hypotheses, and background assumptions of her subfield. Put another way, changes in her subfield tend to result in similar changes in Lydia's standards, values, methods, etc. as well as her beliefs about her field and its subject matter. Call the set of things associated with Lydia's subfield X. X mutually and imperfectly influences the standards, values,

[6] See also Greco (2007).

methods, live hypotheses, and assumptions of several other fields. Call these fields Y and Z. These other fields mutually influence each other and link up to still other subfields in a web. The case, then, would be as follows.

> Lydia's research group produces a result q, which in conjunction with Lydia's explicit or implicit awareness of X produces the belief in her that global warming is the case. The conjunction of Lydia's awareness of both q and X is a reliable basis for believing truly propositions concerning global warming. This reliability is dependent on the mutually reinforcing web of connections between X, Y, and Z. If X were insulated from Y and Z, it would not be a reliable basis for so believing. Lydia is cognitively walled off from this fact about X, and she is cognitively walled off from the ways in which X, Y, and Z are mutually reinforcing. She is only responsive to X, and she, consequently, does not have the ability to discriminate between cases in which the conjunction of q and X is reliable and cases in which it is not reliable as regards the proposition that global warming is occurring.

One can visually represent these relationships in the following manner, where the small circle connected to the big circle is Lydia, the other small circles are the fellow researchers in her lab, the large circle is Lydia's own field of research, and the other symbols are fields of research to which Lydia's own field connects but which are opaque to Lydia.

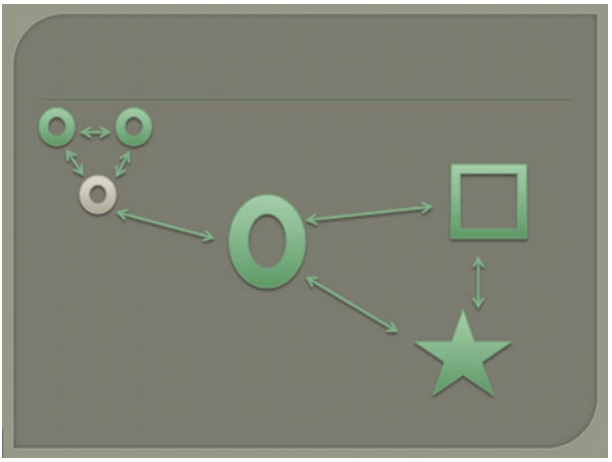

Notice that employing Greco's line on the Morris case to the case of Lydia produces the following result. Lydia does not know that global warming is occurring. After all, she cannot discriminate between when X is a reliable source for propositions of the sort in question and when X is not reliable. X is performing a role analogous to that of a testifier. X is a conduit of information processing performed in places other than Lydia's mind. If Morris only can be said to know if he has the ability to discriminate between reliable and unreliable sources of testimony, then by parity of reasoning it would appear that Lydia does not have knowledge in this case.

Much like the Morris case, the abilities and skills that make information processing effective in this case are primarily those of others. In this case, they are the result of many others. The results of Lydia's group taken alone would be radically insufficient to establish the truth of what Lydia believes, and she is only one of a group producing those results. The bulk of the skill and ability that determines whether Lydia's belief tracks the truth lies in the large set of agents that together bring about a web of information processing to which Lydia has only very limited conscious access. Lydia's case differs from Morris' in the following respect. It is much more plausible that human beings could have an ability to sort individual testifiers who tell us about subjects on which they are competent taken by themselves, though even this thesis is challenged on empirical grounds (cf. Shieber 2012). It is much less plausible that human beings have the sort of abilities that would be required for Lydia to do what Greco requires of Morris.

For Greco, however, knowledge is supposed to be success from ability, where the success is true belief and the ability is intellectual ability. Lydia's belief appears to be a cognitive success that comes about because of or through ability. It is not an accident that Lydia arrives at a true belief. It is not a matter of luck in any straightforward sense. Rather, Lydia's true belief is reflective of skill and competence. In fact, it is reflective of a great deal of expertise on her own part that allows her to conduct the research she does and to be sensitive to the state of her subfield, and it is reflective of a great deal of expertise by the other researchers that together make up the scientific communities involved in X, Y, and Z. In the Morris case, it is unclear whether or not Morris is a skillful cognitive agent, not so with Lydia.

Just as Lydia's belief is not lucky in any ordinary sense of the term, it is obviously not Gettierized. There is no relevant causal interloper. After all, if Lydia had contained within her mind all the information processing used by those she depends on, there would be no question as to whether she has knowledge. Nonetheless, if Lydia does not have knowledge in this case, Greco is committed to denying that Lydia's true belief is a success through cognitive ability (since, for him, that's just what knowledge is).

One can press the case further by considering a parallel, intra-mental case. Suppose that the mind is modular to at least the following extent. Cognitive processing is not wholly conscious, and conscious processing is often dependent on an awareness of the output of subconscious processing modules, which, in turn, may rely mutually on other subconscious processing modules. This is not an unreasonable fancy given what we know about the mind, but this background allows one to construct an intra-mental case parallel to Lydia's extra-mental one.

> John's conscious processing produces a result q, which in conjunction with his awareness of the output of a subconscious processor X produces the belief that p. The conjunction of q and John's awareness of the output of X is a reliable basis for believing truly that p and other propositions in the neighborhood of p. This reliability is dependent on the mutually reinforcing web of subconscious processors Y and Z within which X is embedded. If X were insulated from Y and Z, the conjunction of q and the output of X would not be a reliable basis for believing p. John is cognitively walled off from this fact about X, and he is cognitively walled off from the ways in which X, Y, and Z are mutually reinforcing. He is only responsive to the outputs of X, and he, consequently,

does not have the ability to discriminate between cases in which the conjunction of q and X ends up being reliable and cases in which it is not reliable.

Denying that this kind of modular cognitive architecture could produce knowledge for John would be a rather unhappy result. In fact, it is plausible that this cognitive architecture is implemented in human beings. If I observe someone else getting touched on the arm, I register this fact with a mirror neuron system that uses some of the same neural circuitry that registers when my own arm is touched. I am prevented from having an experience as of my arm being touched when I see another person being touched only because feedback from my own arm sends signals to the brain to indicate that no touch is had by me (Ramachandran 2009, 4:42–6:20). Thus, the component systems that track who all is being touched in one's environment would not be reliable if it were not for the fact that they are integrated in the right way, but whether or not one's subconscious processing modules are correctly integrated is not something to which one tends to have access.

One might object that the intra-mental case of John is not a proper analogue of Lydia's case. In particular, one might object that John has a track record of success available to him that Lydia does not have, and that this track record makes it possible for John to satisfy the discrimination requirement that Lydia fails.[7] I deny both of the objection's claims. Lydia could have a track record of success related to relying on her subfield while being cognitively walled off from various properties of that subfield that contribute to that track record. Presumably, Lydia and John both take themselves to have acquired many true beliefs from their respective sources. In this respect, the two cases match each other. Neither track record, however, would allow the agent in question to be able to discern reliably whether the hidden structure that explains that track record of success is different in the actual case. After all, in the Morris case, the discrimination requirement is satisfied only if Morris can tell that the testimony he is receiving is reliable. Having a track record with testimony in general that shows it to be a generally reliable source is not what Greco has in mind. Even so, in the intra-mental case, John could be aware of the excellent track record of some introspective doxastic source without being in a good position to tell when some of his subdoxastic processors are operating unreliably.

I find it plausible, then, that Lydia can acquire knowledge by the lights of virtue epistemology in a case of strongly distributed cognition where she does not have the ability to discriminate reliably whether her source is really reliable. Lydia's true belief can be a success from ability in spite of the ability in question not satisfying the discrimination requirement. Lydia forming a true belief in this case can manifest a great deal of information processing skill on her part, and it appears strictly analogous to intra-mental cases where it would be crazy to deny either knowledge or ability, as when one perceives that another person was touched and not oneself.

One might protest, at this point, that, by definition, Lydia can't have knowledge by the lights of virtue epistemology because the discrimination requirement just is part of the view.[8] That's too quick, however. There is a core to the virtue epistemolo-

[7] I thank an anonymous referee for drawing my attention to this objection.

[8] I thank an anonymous referee for bringing this objection to my attention.

gist's position. The core has to do with grounding epistemic normativity in ability or competence. It is plausible that possessing epistemic ability or competence will carry with it some discriminatory capacities. There is a conceptual gap, however, between affirming that much and affirming that one must be able to discriminate whether one's source is actually reliable. This point should be familiar from other domains of ability. In a loud arena, a wide receiver in a football game should be able to discriminate reliably whether a ball thrown by the quarterback can be caught in bounds. That is sufficient for him to contribute to the play by attempting to catch the ball. He does not, in addition to this, need to be able to discriminate reliably whether the play has been nullified due to action taking place on some other part of the field. The receiver needs only to be able to treat his local context as one that demands the application of the relevant skillset. To all appearances, Lydia's contribution to the distributed information processing from which she benefits does not require her to have discriminatory capacities that extend beyond her subfield, even though actual reliability in the case hinges critically on features of the broader network.

One might think the proper response at this point would be for Greco to retract the discrimination requirement from his account of testimony. That cannot be the sole modification to his view, however. If he were to drop the discrimination requirement, it would be unclear, then, what answer Greco could give to the Morris case and cases like it. If one lifts the discrimination requirement, it is unclear why Morris' success is creditable to him.

Consider the following modification of Greco's view, which I have suggested elsewhere.[9]

> 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 a high degree of credit for x knowing p *whether those abilities are contributed solely by x or also by other agents*.

Credit For Us maintains that knowledge is a matter of success through ability. It is an achievement, something for which credit is owed. Where this principle differs from Greco's view[10] is that it makes room for the skills and abilities of others to factor constitutively in the domain of evaluation when one evaluates whether someone knows or not. The thought is that knowing is often a team-like activity. In the limit, an achievement can be credited to one individual and the skills and abilities that she brings to bear, but it is often the case that credit deserves to be more widely distributed. In the athletic domain, for instance, it is often the case that whether one person succeeds from ability is a matter of whether the group of which that person is a part succeeds from ability. A role player on a team can be properly credited with the good of winning from ability despite bringing to bear only a fraction of the ability necessary for an achievement. Moreover, the role player has reason to value the success of the team being through ability and not because of luck or biased refereeing.

[9] Green (2012).

[10] In conversation, Greco relates that he will be dropping the discrimination requirement in a forthcoming paper. It is not clear to me how the new version of his theory would handle the Morris case, but I believe it will be more sympathetic to the view I am developing here than his older view.

A computer can exemplify high quality information processing by performing all of that processing itself or by being well connected to a network within which high quality processing occurs. When information processing is distributed across a network, the processing in other computers need not be taken to be a mere background condition for success, like the fact that electricity is available for the computers or that they haven't been infected with a computer virus.

Epistemic evaluation intersects both the domain of information processing and the domain of agential achievement. In the case of Morris and the testifier, Morris' question about the location of the Sears tower invites the testifier to form a temporary information network with Morris and to form it through intentional agential behavior. No doubt, one wants to require that some skill be exemplified by Morris in order for him to get credit for reaping the benefit of this network, but the ability can fall well below Greco's discrimination requirement if what is at issue is a team achievement rather than an individual achievement. The situation is no different from requiring that a role player on a sports team contribute some ability to the team instead of just being a free rider. If one shifts one's focus away from the normative features that Morris can ground solely with his own abilities, it becomes more plausible that he can hold up his end of the bargain.

In the case of Lydia, it is Lydia's skill and ability that ties her into a vast network of information-processing. She does not have to be able to recognize or be aware of the virtues of this network in order to be functionally integrated into the network. It would be a misconstrual of the situation to claim that Lydia's success is creditable to her abilities taken on their own. Her abilities are radically insufficient for that kind of success. Rather, Lydia forms a true belief as a result of the intersecting abilities of a group to which Lydia belongs. It is her participation in this group that justifies the claim that the success from ability is a success from ability for her, and this participation itself requires the exercise of abilities designed to hook into a broader network of ability. Her abilities allow her to get credit as a legitimate possessor of the good of true belief, but she might get only a modest amount of the credit for generating the good-making properties that her belief has. To switch to Sosa's terminology, Lydia's belief manifests her information-processing competence and this competence includes how one participates in the wider scientific community so as to reap the benefits of the collective scientific enterprise.

Before turning to some alternative accounts of distributed cognition, let us entertain an objection to the view.[11] One can imagine cases that are more extreme than Lydia's, cases where the agent in question does less to earn our approval but nonetheless obtains a reliable belief due to the efforts of others. Aren't we going to reach some point at which it is implausible to ascribe knowledge to the agent, especially on a *credit* view of knowledge? The short answer is "yes". There must be some minimal threshold of ability that an agent contributes in order to know on a credit view. Establishing where that minimal threshold should be placed is not easy.

For this paper,[12] it should be sufficient to point out that the threat of trivial accomplishments and the related need to define what counts as a minimal contribution to

[11] I thank an anonymous referee for prompting this addition to the paper.

[12] In a sequel paper, I develop the relationship between the group and the individual in team-like achievements at greater length and, in doing so, I also go to greater length in discussing this sort of objection.

distributed cognition are no different from challenges one faces, on the one hand, in giving an account of team-related success generally and, on the other hand, in developing any other epistemology. In athletics, for example, no one thinks that a player needs to be the focal point of the team for a success for the team to count as a success for the player. Nonetheless, the player needs to satisfy some minimal requirement in order for others to think that she is entitled to share the honor of victory. She needs to have contributed to the success in some way, usually some way specified by the role she is playing on the team. One can readily bring to mind what it might look like to contribute to the success of a sports team, and one can readily bring to mind what it might mean to be part of the team in name only. There will be a grey area, however, in which it is difficult to discern whether the contribution of the agent licenses the claim that the success of the team counts as a success from ability for her too. The existence of a grey area in the athletic domain does not call into question the applicability of a credit approach to athletic success.

Any epistemology that appeals to a value that comes in degree will similarly have trouble establishing a minimal level at which enough of that value is had for knowledge. Reliability and evidential support come in degrees, for instance, and there is no easy way of discerning exactly what the minimal threshold is that must be met for the evidentialist, the (un-extended) virtue epistemologist, or the process reliabilist. Moreover, ancillary concepts like cognitive integration, dispositional strength, and global or local reliability also come in degrees. To make matters worse, if contextualism is right, knowledge ascriptions and thus the minimal requirements for those ascriptions (properly) vary with context. Moreover, it is not uncommon for contextualists to affirm that the lower bound on knowledge ascriptions is rather undemanding.

Despite the fact that all these epistemologies have trouble giving a satisfying answer to where the minimal thresholds for their respective normative values lie, they all count as serious contenders in epistemology. This is partly due to the fact that they can provide clear exemplars of what it would look like to satisfy or fail to satisfy the values that ground their account and that they can use their account to diagnose what makes some of the grey cases grey. There is no reason to think that an extended credit theory will do any worse than the competition in this regard.

3 Some distributed process views

Sandy Goldberg and Joe Shieber have each developed provocative accounts of testimony that privilege social phenomena. They have done so from within process reliabilism and thus without availing themselves of the resources of virtue epistemology or a credit theory of knowledge. One might think that focusing on sheer reliability instead of the efforts and abilities of individuals provides a more streamlined account. In this section, I will argue that both Goldberg and Shieber's process reliabilist views are thereby impoverished as accounts of distributed cognition.

Goldberg defends an anti-individualist position on testimony on which the mental processes of a testifier factor constitutively into the belief-forming process of a testimonial belief in the recipient of testimony (Goldberg 2010, p. 58). Much like the overall quality of an inference is dependent on the quality of the cognition that derived

the premises an inference proceeds from, so one's reception of testimony by itself is at most conditionally reliable (Goldberg 2010, p. 65). For the formation of a belief in testimony to be reliable overall, the cognitive doings that eventuate in an act of testimony must be reliable as well.

Goldberg does not much countenance more complicated or widely distributed processes, but it is not hard to see how one might extend the view so as to cover strongly distributed cases like that of Lydia.[13] In the standard testimony case, the belief-forming process must be regarded as extended because one's intra-mental process, as a receiver of testimony, is conditionally reliable, and the condition that must be satisfied is the reliability of the cognitive doings of the testifier. That testifier's cognitive doings, however, could be conditionally reliable on others as well. Thus, one could trace out a web of connections in cases of strongly distributed cognition so as to evaluate Lydia's belief in global warming by iterating Goldberg's base scenario.

No argument is needed to extend Shieber's view to distributed cognition. He argues explicitly that a process reliabilism that denies what he terms the "personalist requirement" can accommodate distributed cognition, including and especially what I have been calling strongly distributed cognition. Shieber cites Jennifer Lackey for his formulation of the personalist requirement to the effect that the putative possessor of testimonial knowledge must be "a reliable or properly functioning recipient of testimony" (Shieber, p. 267). He later describes the requirement as the requirement that "testimonial warrant involves, at least in part, (a) the exercise of a capacity or sensitivity to (b) properties of the testimony" (Shieber, p. 285).

Shieber thinks that cases of strongly distributed cognition undercut this requirement due to the "opacity of epistemically relevant properties" and the "non-locality of expertise" (Shieber, p. 275).[14] Put in terms of the case of Lydia with which we have been working, he is drawing attention both to the way in which Lydia is cognitively walled off from properties that make her belief reliable and to the relevance of expertise as a group-level feature. Lydia, as we have seen, might not be able to discern reliably the presence of group-level reliability vis-à-vis the proposition that global warming is true, but, if this group-level feature to which Lydia does not have access is necessary for Lydia's belief to be sufficiently reliable, then Shieber would say the personalist requirement is not satisfied in the case.

What follows according to Shieber? A recipient of testimony does not need to form her testimonial belief in virtue of recognizing or being sensitive to the presence of some special epistemic property in the testifier (Shieber, p. 288). Rather, the process that eventuates in a belief in the recipient of testimony must be reliable, and it does not

[13] He does pay attention to a very interesting case involving knowing through not having heard something that falls within a domain over which one's epistemic community has coverage. For instance, one can know that the president was not assassinated yesterday because one would have heard about it if that had happened. I set this kind of case aside, however, since Goldberg is explicit that an individual forming a belief due to coverage is not doing so through some kind of extended process but rather through an intra-mental inference (Goldberg 2010, pp. 154–155).

[14] Shieber also thinks that it is demonstrable from empirical literature that human beings do not have any competence that could satisfy the personalist requirement (Shieber 2012). This is an important part of Shieber's overall project. I cannot do it justice in this paper, but I take it up in Green forthcoming.

really matter what social or environmental factors undergird this reliability so long as it is there (Shieber, p. 289).

Goldberg and Shieber agree that epistemic evaluations can extend beyond the intramental processes of an individual so as to include the cognitive doings of others. The chief difference between their views and the one I have put forward is that, on my account, a network that supports knowledge needs to be built up out of skill and ability. The links between individuals need to be purposeful enough and skillful enough for the resultant true belief to be a sort of information-processing accomplishment of the epistemic agents who compose the group. As elaborated previously, it is no easy thing to set out minimal thresholds that must be surpassed to count as a success, but, fundamentally, the skill or ability of the members of the team needs to explain the occurrence of information-processing success. This is not required by Goldberg and Shieber. They only require reliability. They do not require that reliability be grounded in ability. It must simply be present in the extended belief-forming process.

I will focus here on developing one problem with extended process reliabilism, which has to do with the way it exacerbates an objection that credit and virtue theories already level at process reliabilism having to do with "strange and fleeting processes" (Cf. Greco 1999, pp. 286–289).[15]

For an objection to *simple* process reliabilism from strange and fleeting processes, take Alvin Plantinga's case of the brain tumor that produces in oneself the belief that one has a brain tumor (Plantinga 1988, p. 30). It does not do so via some experience of pain or pressure in one's head that thereby causes one to infer the presence of a brain tumor. Rather, it produces the belief all at once and fully formed. Beliefs formed through the interpolations of tumors don't seem to be cases of information-processing excellence even if they happen to be reliable, and most people deny that knowledge is had in this case. So, the objection goes, there must be more to knowing than just forming a belief in a reliable way. Plantinga's case is illustrative of a general recipe. To produce a counterexample to process reliabilism, one needs only to find a case in which reliability is had but for a reason that intuitively has little to do with the quality of one's cognition. One should look for a case in which elements foreign to our epistemic interests underwrite the reliability of the process.

A simple process reliabilist may object to many cases of strange and fleeting processes for reasons peculiar to that particular case, or she may bite the bullet, welcoming any odd belief-forming process that happens to be reliable. Arguably, a more satisfying general response to counterexamples fitting this recipe would include a way of tying belief-forming processes more closely to the cognitive processes of the agent who is being evaluated.[16] If the simple process reliabilist were to insist that epistemic evaluations have to do with the reliability of an agent's cognitive processes, she can resist many of the putative counterexamples. Strange processes that have to remain within the recognized cognitive system of an agent tend to be much less obviously problematic. For instance, one might dismiss the brain tumor case on the grounds that

[15] In Green (2012), I interact with Goldberg's view at length, and I develop a separate objection to his view there that would complement the line I'm urging here.

[16] Cf. Goldman (1993), especially p. 282. It is interesting that Goldman's response to strange and fleeting processes in this essay is combined with some explicit flirting with virtue epistemology.

the brain tumor's inner workings are not part of the agent's cognition, and thus, the reliable output of true beliefs by the conjunction of the brain tumor and the agent's brain does not automatically get valorized by the simple process reliabilist. Furthermore, it is very unlikely that the cognitive processes interfered with by the brain tumor will be reliable taken on their own. It is much more likely that the interference of the tumor will make them more unreliable. Thus, one might think a constrained simple process reliabilism can diffuse the example.

The simple process reliabilist can justify a focus on the intra-mental cognitive processes of the agent by appealing to intuitions about the scope of epistemic evaluation that are independent of reliabilism. As Alvin Goldman says in a discussion of justification, what is being evaluated are "'*cognitive*' operations," which "are most plausibly construed as operations of the cognitive faculties, i.e., 'information-processing' equipment *internal* to the organism" (Goldman 1979, p. 97).[17] The extended process reliabilist, by abandoning the traditional focus on only intra-mental processes, likewise removes this means of defending herself against cases of extrinsically supported, reliable cognition.

Compare a version of the case of Lydia with a science fiction case involving Dido. Let us specify that, in both Lydia's case and the case of Dido, the human members of the respective social networks are connected only by computer and that the extended belief-forming processes that eventuate in the beliefs of Lydia and Dido are equally reliable. Mapping the human and digital events that bring about the true beliefs of Lydia and Dido will look very similar in the abstract.

...[Human]→[Computer]→[Human]→...[Computer]→[Lydia's forming a true belief]
...[Human]→[Computer]→[Human]→...[Computer]→[Dido's forming a true belief]

Dido's situation is much like Lydia's except that, as Dido inhabits a science fiction world, let us specify that this world contains supercomputers with special capabilities and that Lydia's does not contain these computers. In Lydia's case, the reliability of Lydia's belief is grounded in the information-processing abilities of the humans in the extended chain. Some of these abilities are abilities to use computers and other equipment, but one cannot explain what is going on in Lydia's chain without making reference to the exercise of epistemic agency by the people in Lydia's social network. In Dido's case, the human scientists are tools used by the supercomputers (which, let us specify, are not conscious agents). In particular, the supercomputers causally manipulate the vast majority of what the human researchers do in the conducting of research. The humans are used in ways that allow the supercomputers to make progress reliably, and the research results are sometimes stored in true beliefs the scientists form under the supervision of the computers.

To make the example more vivid, imagine that all the human beings in the case of Dido are children. The supercomputers arrange all the research tasks so that they appear to be games. The children do various actions that mimic what the researchers

[17] The page number is from the reprinted version. Goldberg cites this same passage on Goldberg (2010, p. 43.)

in the Lydia case do. In the case of Lydia, the researchers perform their actions as a manifestation of abilities intentionally directed at getting at the way things are by doing scientific research, not so in the case of Dido. Dido and the rest of the children respond reliably to their manipulated environment due, not to cognitive skill employed in an effort to get at the way things are, but due to the way the supercomputers shape game-playing behavior. In this environment, Dido finds herself believing p as the terminus of an extended process involving the manipulated doings of the other children. We can specify that Dido's cognition is not interfered with directly and that her intra-mental cognitive process is conditionally reliable, conditional upon the unrecognized manipulation of the other children in her social network by the machines. Dido's conditionally reliable belief would not be reliable on its own, but the extended process that produced this belief is reliable.

Intuitively, there is a large normative difference between the two cases. Lydia's case is reflective of information-processing done by agents having gone well, and the information-processing of agents having gone well is what explains why Lydia forms a true belief. Dido's case is one in which the epistemic agents are mere simulacra of agents engaged in good information-processing. What the agents are doing is externally supported by the supercomputers so as to make it reliably end in a true belief. Moreover, Dido does not use the other agents like one might use a thermometer or a clock. She just happens to be the beneficiary of a reliable process conducted by the machines.

The extended process reliabilist is in a difficult position. On the one hand, Dido and her coworkers form true beliefs through reliable processes. One can even say that the supercomputers use their human tools in a way such that Goldberg's requirement of conditional reliability is met. The human beings might not do what they do in virtue of a sensitivity to the epistemic properties of the efforts of others. This is true. Such a sensitivity is irrelevant for Shieber, however, and due to the high quality of the programming of the supercomputers, it is unnecessary to get the conditional reliability that Goldberg demands.

On the other hand, it seems absurd to claim that there is no relevant normative difference between these two cases for the human cognizers involved. Information-processing excellence may be involved in both cases, but it's human in one case and digital in the other. Neither overall reliability nor the conditional reliability of the parts can draw the line between the two cases appropriately. Furthermore, even if the extended process reliabilist found some excuse for excluding belief-forming processes that include causal chains that run through computer intermediaries, it could not do so without ruling out forms of the Lydia case that cannot be plausibly excluded.

The extended credit theory fairs better. The extended credit theory can point out that in the Lydia case, but not the Dido case, it is the skill of the human researchers that explains why true beliefs are reached. In the Lydia case, what computers do is explained by the way in which the scientists co-op the computers to extend the scientists' abilities. The Lydia case is a cooperative endeavor, albeit one built out of local interactions that can be blind to the form of the global process being participated in. The Dido case is one of manipulation so far as its epistemic dimension is concerned. In the Lydia case, the human scientists deserve credit for success. In the Dido case,

they do not. Thus, from an extended credit perspective, one has reason to think that Lydia is much better positioned than Dido epistemically.

In conclusion, cases of strongly distributed cognition put pressure on credit theories of knowledge to allow the abilities of other agents to factor constitutively into epistemic evaluation. The resulting framework is one in which knowledge can be a team-like accomplishment. An extended credit theory also does a better job of describing the normative terrain of distributed cognition than some more stream-lined process reliabilist accounts, which founder when it comes to cases of reliable distributed cognition where the information-processing abilities of agents are not what ground reliability.

References

Coady, C. A. J. (1992). *Testimony: A philosophical study*. Oxford: Clarendon Press.

Giere, R. N. (2006). *Scientific perspectivism*. Chicago: University of Chicago Press.

Giere, R. N. (2012). Scientific cognition: Human centered but not human bound. *Philosophical Explorations*, *15*(2), 199–206.

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

Goldman, A. (1979/2000). What is justified belief? In Sosa, E., & Kim, J. (Eds.) *Epistemology: An anthology* (pp. 89–104). Oxford: Blackwell.

Goldman, A. (1993). Epistemic folkways and scientific epistemology. *Philosophical Issues*, *3*, 271–285.

Greco, J. (1999). Agent reliabilism. *Nous*, *33*(13), 273–296.

Greco, J. (2003). Knowledge as credit for true belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (pp. 111–134). New York: Oxford University Press.

Greco, J. (2007). Discrimination and testimonial knowledge. *Episteme*, *4*(3), 335–351.

Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. New York: Cambridge University Press.

Green, A. (2012). Extending the credit theory of knowledge. *Philosophical Explorations*, *15*(2), 121–132.

Green, A. Monitoring, testimony, and a challenge from social psychology. American Philosophical Quarterly (forthcoming).

Hutchins, E. (1995). *Cognition in the wild*. London: MIT Press.

Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese*, *158*, 345–361.

List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. New York: Oxford University Press.

Plantinga, A. (1988). Positive epistemic status and proper function. *Philosophical Perspectives*, *2*, 1–50.

Ramachandran, V. S. (2009). The neurons that shaped civilization. TED talks. Retrieved April 18, 2012 from www.ted.com/talks/lang/en/vs_ramachandran_the_neurons_that_shaped_civilization.html.

Shieber, J. (2012). Against credibility. *Australasian Journal of Philosophy*, *90*, 1–18.

Shieber, J. (2013). Toward a truly social epistemology: Babbage, the division of mental labor, and the possibility of socially distributed warrant. *Philosophy & Phenomenological Research, 86*(2), 266–294.

Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. I). New York: Oxford University Press.