JOHN GRECO

DISCRIMINATION AND TESTIMONIAL KNOWLEDGE

ABSTRACT

Sanford Goldberg has called our attention to an interesting problem: How is it that young children can learn from the testimony of their caregivers (their parents, teachers, and nannies, for example) even when the children themselves are undiscriminating consumers of testimony? Part One describes the importance and scope of the problem, showing that it generalizes beyond tots and their caregivers. Part Two considers and rejects several strategies for solving the problem, including Goldberg's own. Part Three defends a solution, positing a previously unnoticed social dimension to knowledge.

In recent work, Sanford Goldberg has called our attention to an interesting problem in the epistemology of testimony.[1] The problem can be summarized as follows: How is it that young children can learn from the testimony of their caregivers (their parents, teachers, and nannies, for example) even when the children themselves are undiscriminating consumers of testimony? Another way to formulate the problem is to note an inconsistency among our pre-theoretical commitments about testimonial knowledge. Specifically, we want to say all of the following:

a. that young children can learn from the testimony of their caregivers; i.e. that they can come to *know* through such testimony,
b. that testimonial knowledge requires a reliable consumer of testimony; i.e. that the hearer can reliably discriminate between reliable and unreliable sources of testimony, and
c. that young children are not reliable consumers of testimony.

The problem, then, is to fix up the inconsistency in a plausible way. That is, to fix things up in a way that is consistent with both i) a plausible account of the nature of testimonial knowledge in general, and ii) a plausible account of the extent (or range) of testimonial knowledge.[2]

Goldberg thinks that solving the problem requires making substantive and even surprising commitments regarding the nature of testimonial knowledge. In particular, he argues, it requires that we recognize "social" dimensions to

DOI: 10.3366/E1742360007000123

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

testimonial knowledge that are *sui generis* and that have previously gone unnoticed. I agree with Goldberg that the problem he has identified is important and that addressing it commits us to substantive theses. However, I disagree that Goldberg's own proposals actually solve the problem. In the first part of the paper I further describe the problem and make some comments about its importance and scope. In the second part of the paper I consider several strategies for solving the problem, including Goldberg's own. I conclude with some considerations in favor of the solution I prefer.

1. THE PROBLEM OF (EPISTEMICALLY) UNDISCRIMINATING CHILDREN

Consider the following example from Goldberg, designed to elicit the intuition that testimonial knowledge requires a reliable consumer of testimony.

> **Room Full of Liars**. Sid is gullible in the extreme: he accepts anything anyone says merely in virtue of the fact that someone said so. Sid is in a room full of inveterate liars. He immediately and uncritically believes everything each of them says. At one point he happens to bump into Nancy, the only reliable person around. Nancy reliably tells Sid that *p*, and (as a matter of course) he believes her. (Goldberg forthcoming b)

Goldberg concludes, and I agree, that Sid does not know that *p* through Nancy's reliable testimony. The obvious diagnosis, Goldberg notes, is that Sid is not a reliable consumer of reliable testimony.

But now compare this with a second example:

> **Babe.** Babe is gullible in the extreme: like other small children, he accepts anything anyone says merely in virtue of the fact that someone said so. Babe is told by his mother that there is milk in the fridge, and this testimony is both true and reliable. However, Babe's notorious Uncle Marfman, who gulls little children every chance he has, is in the next room. He could easily have told Babe that there is milk in the refrigerator when there was none, and if he had done so, then Babe would have believed him.

Goldberg concludes, and I agree, that Babe knows that there is milk in the refrigerator. His mom told him so, and small children come to know things in this way. But as Goldberg notes, "there are nearby possible worlds in which [Babe] forms the testimonial belief that there is milk in the fridge, under conditions in which there is no milk in the fridge." (Goldberg forthcoming b) And therefore, it seems, we have a case of testimonial knowledge without a reliable consumer of testimony.

Here is a similar case.

> **Preschool.** Gabe is gullible in the extreme: like other preschool children, he accepts almost anything anyone says merely in virtue of the fact that someone said so. Gabe is told by his preschool teacher that frogs eat bugs, and this testimony is both true and reliable. However, Gabe is surrounded by his preschool classmates, who are notoriously

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

unreliable reporters. One of them, who happens to be confused on the issue, could easily have reported that frogs never eat bugs, and Gabe would have believed this.

The intuition here is that Gabe knows that frogs eat bugs because his teacher told him so, and this despite the fact that Gabe would have believed otherwise if his confused classmate had so testified.

I suggested above that the problem can be summarized as follows: How is it that young children can learn from the testimony of their caregivers (their parents, teachers, and nannies, for example) even when the children themselves are undiscriminating consumers of testimony? Stated in this way, the problem amounts to a call for explanation. Here is another way to make that call: How is it that small children can come to know from their caregiver's testimony, even if they are "locally unreliable" as consumers of testimony, i.e. even if they could easily go wrong by basing their beliefs on testimony?[3] How is it, given the *lack* of knowledge in Room Full of Liars? We can call this "The Problem of (Epistemically) Undiscriminating Children." But that label is somewhat cumbersome, so for short we can call it "Goldberg's Puzzle."

How widespread is Goldberg's Puzzle? Goldberg formulates it as a problem about very young children – children who are too young to reliably discriminate among testimonial sources. But there are reasons for worrying that the problem generalizes far beyond undiscriminating tots.

Consider the structure of our examples above: A candidate for testimonial knowledge (the protagonist) receives testimony from a reliable source, but in an environment where unreliable testifiers are lurking. Moreover, the protagonist is "locally unreliable" as a consumer of testimony: she does not discriminate between reliable and unreliable sources of testimony in her environment, and so she could easily go wrong by forming beliefs on the basis of testimony. In Room Full of Liars, these salient features of the case generate the intuition that Sid does not know. In Babe and Preschool these features are likewise present, but don't generate the intuition that the protagonists do not know. The problem, as formulated so far, is that both good cases (where the protagonist does know) and bad cases (where the protagonist does not know) seem to have the same structure in all relevant respects.[4]

What I now want to suggest is that the problem generalizes: there are good cases with this same structure, and that do not involve very young children as protagonists. Older children, and even seasoned adults, can be locally unreliable in cases where they seem to gain knowledge. Here are two such cases.

High School. Troy is a high school student with normal cognitive capacities. One day Ms. Goodman, an expert in American history, tells the class that the U.S. Constitution was ratified in 1789, and Troy believes this on the basis of her saying so. However, Troy could easily have not been paying attention, and relied instead on his girlfriend Gabriella, who is an unreliable note taker. If Troy had asked Gabriella, she would have reported that the Constitution was ratified in 1879, and Troy would have believed her.

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

Professional Conference. Mary is a layperson interested in global warming and decides to attend a professional conference on the topic. During the conference she hears from Dr. Hothouse, the foremost expert on global warming in the world, something that is in fact common knowledge among experts in the field: that plans recently adopted by industrial nations to curb harmful gasses will not be sufficient to reverse the global warming trend. However, Dr. Crankpot, a convincing imposter, is in the next room pretending to be an expert on global warming. During his presentation he testifies that the recently adopted plans *will* be sufficient to reverse the global warming trend, reporting this as if it were common knowledge among the experts. Not being an expert on the topic herself, if Mary had heard this, she would have believed him.

***DSM-IV*.** Lizabeth picks up her copy of the latest Diagnostic Statistical Manual and reads that a diagnosis of Anorexia Nervosa requires a body mass index equal to or below 17.5. However, one of the copy-editors for this version of the manual suffers from that disease and disagrees that a body mass index of 17.5 is unhealthy. If she had been assigned the relevant pages to copy-edit, she would have changed the number to 16.5. And if she had done so, Lizabeth would have believed that this criterion for the disorder is the correct one.

My intuition is that there is knowledge in all three of these cases. And yet all of them have the structure of the Room Full of Liars case described above, where it seems clear that Sid does not have knowledge. In each case the protagonist learns from a reliable source, but in circumstances where he or she is locally unreliable.[5] But then why does the protagonist know in the good cases but not in the Room Full of Liars case? This is a more general problem than the one formulated above, because our good cases now involve older children and adults, and not just tiny tots. Moreover, our cases now involve protagonists with significant capacities for discriminating between reliable and unreliable sources of testimony. The problem is that people can be generally reliable in this respect and yet locally unreliable. And it seems that sometimes this local defect serves to undermine knowledge and sometimes it does not.

What is going on here? One thing to note, and I will elaborate on this later, is that in the good cases the source of testimony is more than *de facto* reliable – it is also socially approved as a source of information and knowledge. That is, in each of the good cases the source of testimony has a privileged social status as a source of information and knowledge: we have parents, teachers, experts, and "definitive sources" such as the DSM-IV. Here, then, is another way to describe the problem in its more general form: How is it that people can learn from the testimony of "socially approved" sources (parents, teachers, experts and definitive sources) even when the hearers themselves are locally unreliable consumers of testimony? This is the generalized form of "Goldberg's Puzzle."

2. STRATEGIES FOR SOLVING THE PROBLEM

a. Skepticism

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

One way of responding to the problem set out above is to take a skeptical stance toward testimonial knowledge. That is, one might deny (a) above: that young children can learn from the testimony of their caregivers; i.e. that they can come to know through such testimony. Taking this skeptical stance, one would deny knowledge in all of the cases above, including those associated with the generalized version of Goldberg's Puzzle. One need not embrace a general skepticism about knowledge, or even about testimonial knowledge, in order to adopt this approach. Rather, one need only insist that testimonial knowledge requires local reliability, and insist that this criterion is not filled in the cases at issue.

This approach is unattractive, however, insofar as we mean to be explicating the concept of knowledge actually in use in ordinary speech. For it is plausible that Babe and Gabe represent paradigmatic cases of knowledge in that sense. Children believing their caregivers, such as their parents or teachers, constitute paradigmatic cases of coming to know. To put the point a different way, we make an easy distinction between learning that p from one's parents or teachers and believing that p on the say-so of less reliable sources, and we use our ordinary concept of knowledge to mark that distinction.

Some might insist, as some philosophers have, that they are not interested in what ordinary language calls "knowledge" – that they are interested in something more robust and harder to achieve. But we can, in turn, insist that we are interested in the ordinary concept of knowledge, i.e. the concept of knowledge that is actually in use and that serves to mark the distinctions that we actually make. *That* concept of knowledge seems to allow for social dimensions of various kinds, and the cases reviewed above suggest special social dimensions for testimonial knowledge in particular. Taking these cases seriously moves us away from the "rugged individualism" that often goes with skeptical inclinations.

There is, it seems, more room to deny knowledge in the other cases above than in those involving tots and their caregivers. For one, skeptical consequences do not threaten to multiply here as they do in the cases involving the tots. That is, the non-tot cases are likely to be rare, whereas the tot cases are likely ubiquitous, and so if there is no knowledge in the former we get less widespread skeptical results than if there is no knowledge in the latter. Importantly, there is less danger of a cascading effect, where we seem unable to account for knowledge in later life because there is too little knowledge early on. However, if we do consider the Babe and Preschool cases to be paradigmatic, then there is strong motivation for treating High School, Professional Conference and DSM-IV as cases of knowledge as well. That is, once we acknowledge special epistemic roles for socially approved sources such as parents and caregivers, there is motivation to extend the scope of this phenomenon to other sources as well.

b. Straightforward Reliabilism

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

A second approach to the present problem is straightforward reliabilism. That is, one might continue to hold that testimonial knowledge requires a reliable consumer of testimony, but insist that even small children fit the bill. This would be to deny (c): that young children are not reliable consumers of testimony. One might insist, for example, that such abilities are acquired early on, through the earliest processes of language training and other kinds of social development. This approach might even explain the special role of "socially approved sources" suggested above, insofar as reliance on such sources is the natural and expected outcome of the social training at issue.

One problem for this approach is that it is unsupported by the empirical evidence. Although there is evidence that development of relevant capacities begins early on, the development is slow and early performance is unimpressive.[6] Moreover, the present approach cannot handle the Babe and Preschool cases, where it is stipulated that the protagonists do not have the requisite capacities to discriminate reliable from unreliable testimony. Nor can it handle our other good cases, where it is stipulated that protagonists are locally unreliable.

Finally, we may return to the point that our ordinary concept of knowledge distinguishes between information learned from caretakers and information taken from other sources. Moreover, our common sense conception of young children is that they are gullible – it would come as a *surprise* that they are in fact highly reliable consumers of testimony. But then we can't explain our intuitions about Babe and Preschool by insisting that young children are, after all and despite appearances, reliable in this respect.

c. Goldberg's Solution

I now want to look at Goldberg's own solution to the problem we have been considering. I think there is a lot to learn from what Goldberg says here, although I don't think he actually solves the problem that he has raised.

Goldberg's strategy for addressing the problem is to re-conceive the process by which children form beliefs on the basis of testimony. Specifically, we should conceive that process as essentially involving the child's social environment. That social environment, in turn, is conceived as essentially involving epistemic caretakers, such as parents and teachers. These epistemic caretakers, Goldberg suggests, regularly monitor sources of testimony for reliability, and monitor uptake of testimony for understanding and accuracy. Babe's mother, for example, is there to see Uncle Marfman pulling Babe's leg, and there to correct misinformation if the practical joking goes too far. Likewise, Gabe's preschool teacher is there to monitor classmates and to correct misinformation taken from unreliable peers.

Goldberg explains the idea by means of an analogy.

> Samantha is a teenager who has just received her drivers' permit, which allows her to drive a car only under certain very restrictive conditions: she can only drive during the day, in good weather, and when in the company of an experienced driver. If we

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

> consider Samantha herself, independent of these restrictions on her driving, she is not a particularly safe driver (yet): she has difficulties making a left-turn into oncoming traffic, does not allow herself sufficient braking distance, has a tendency to drive too close to the cars ahead of her, rarely uses her turn signals, tends not to obey the speed limit, and cannot parallel park. At the same time, when she drives under the restricted conditions set forth by her permit – and these are the only conditions under which she drives – she drives very safely: she only drives during the day and in good weather; she is gradually desensitized to making a left into oncoming traffic; she is reminded to use her turn signals; she is urged to slow down when she starts to go too fast, and to speed up when she goes too slowly; and so forth. (Goldberg forthcoming b)

Goldberg asks: Is Samantha a good (reliable, safe) driver?

> In one sense, she is not: if she were to drive alone, she would likely exhibit many serious driving flaws. In another sense, she is: as a matter of fact, she only drives in the restrictive conditions set forth by her permit, with the result that her environment is structured in such a way as to reduce the occasions on which she exhibits those flaws, and to minimize her flaws' effects on those occasions when she does begin to exhibit them. We might then describe this state of affairs by saying that, although Samantha herself is not yet a safe driver, she drives safely in the very restrictive conditions under which she drives (i.e. those set forth by her permit).

The analogy to small children and their reception of testimony should be obvious. Goldberg writes,

> My claim is that, just as we do not adequately characterize Samantha's driving without taking into account the social context of her driving (in particular, the laws restricting when she can drive, and with whom), so too we do not adequately characterize the process involved in young children's consumption of testimony without taking into account the child's *social environment*. Rather, we should consider the role that the child's adult guardians play in the process that eventuates in the young child's consumption of testimony.

With that role considered, Goldberg argues, we can say that small children often satisfy the reliable consumer condition on testimonial knowledge. Small children are often reliable consumers of testimony, relative to the social environments they are actually in.

I said above that there is much to recommend in Goldberg's proposal. One advantage of the proposal is that it highlights a social dimension of testimonial knowledge that can easily go unnoticed. Consider: knowledge in general requires a "friendly" environment. For example, perceptual knowledge requires physical conditions that enable the proper functioning of our perceptual faculties. Such conditions include appropriate lighting, opaque objects, and other "normal" conditions necessary to prevent various kinds of optical illusion. But so too does testimonial knowledge require a friendly (i.e. enabling) environment, in this case a friendly social environment. The most obvious case in point is that one needs reliable testimonial sources. But as Goldberg points out, environments can be friendly

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

or enabling in other ways as well. And it makes perfect sense to recognize this in the ways we conceive testimonial knowledge and the process of testimonial exchange.

Nevertheless, it seems to me that these insights do not solve the problem at hand. Goldberg's solution to the Problem of Undiscriminating Children amounts to this: Given adult monitoring in a friendly social environment, even small children are reliable consumers of testimony. The problem with this solution is that adult monitoring does not do the job required. That is, adult monitoring does not produce reliability in the full range of cases where we want to count young children as knowing from testimony.

Consider, for example, Goldberg's own case of Babe. Suppose that Mom is too busy to keep a watchful eye on Uncle Marfman. She can't be on him every minute, and Marfman is willing to wait for his opportunity. In that case, there are close possible worlds where Babe believes on the basis of testimony that there is milk in the refrigerator, but there is none. Babe is not reliable in this respect. Nevertheless, Babe knows there is milk in the refrigerator when *Mom* tells him so. It doesn't matter that Babe could easily have been fooled by Uncle Marfman if *he* told him so. We can make the same point about Preschool. Teachers are busy people – they can't stay on top of all the little brats all the time. Nevertheless, Gabe knows that frogs eat bugs when his teacher tells him so. It doesn't matter that Gabe could easily have been fooled by classmates if they had told him something different.

In sum, epistemic caregivers cannot be so vigilant as to insure reliability all the time. Nevertheless, learning from testimony often takes place anyway. In particular, young children can learn from Mom and Teacher, even in an environment where the children are unreliable consumers of testimony.

d. Disjunctivism

The forgoing consideration suggests that there is something special about moms and teachers. Well, we already knew *that*. But the forgoing suggests that there is something *epistemically* special about moms and teachers, and about certain other people as well. Goldberg's idea was that such people have "a distinct interest" in their charges' consumption of testimony, and as a result they take steps to bolster the reliability of that consumption. But it now looks as if their role is more special still. Moms and teachers, and some others as well, seem to play a special role as *sources* of testimony: they seem to be sources of testimonial knowledge in ways that others are not.

One might suppose that they play this role by virtue of their special social status. And in fact we suggested as much in our generalized formulation of "Goldberg's Puzzle": How is it that people can learn from the testimony of "socially approved" sources (parents, teachers, experts, and definitive sources) even when the hearers themselves are locally unreliable consumers of testimony? "Disjunctivism" about testimonial knowledge takes this idea seriously. On a disjunctivist account, there are two routes to testimonial knowledge:

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

S has testimonial knowledge that p if and only if p is true and either

a) S's testimonial belief that p satisfies the usual conditions for knowledge (for example, S's belief that p is reliably formed)
or
b) S believes p on the basis of a socially approved testimonial source.

A disjunctivist approach to testimonial knowledge would make quick work of Goldberg's Puzzle, easily explaining how young children and others can learn (i.e. gain knowledge) from socially approved sources. One problem, of course, is that the approach seems *ad hoc*. The view seems tailor made to explain just the cases we are having trouble with, but otherwise unmotivated. In the remainder of this section I will argue that this problem is not as bad as it seems. There are good and independent reasons for thinking that testimonial knowledge should have this disjunctive nature.

Here I draw on some important ideas from Edward Craig in his *Knowledge and the State of Nature* (1990). The first idea is methodological. Craig suggests that we can gain insight into what knowledge *is* by asking what the concept of knowledge is *for*. That is, we should ask what roles that concept plays in our conceptual economy and our linguistic practices. Craig describes the methodology as follows:

> Instead of beginning with ordinary usage, we begin with an ordinary situation. We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application. (2)

The second idea I want to take from Craig is substantive. Specifically, Craig defends a particular hypothesis regarding what role the concept of knowledge plays. At least one important function of that concept, he argues, is to flag good sources of information.

> Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths . . . So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use. The hypothesis I wish to try out is that the concept of knowledge is one of them. To put it briefly and roughly, the concept of knowledge is used to flag approved sources of information. (11)

I take it that Craig's "hypothesis" about the concept of knowledge is plausible. Suppose it is right. Then it would be no surprise if that concept privileged some sources of information over others as vehicles of information flow. We would expect some sources to get special status as "originating" sources of information. But we should also expect others to get special status as transmitters of information.

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

On the proposed picture, our concept of knowledge would perform two related duties. First, it would identify "primary" sources of information, or sources where information originates in a reliable way. Presumably, one's status of this sort would depend both on one's general capacities, such as perception and reasoning, and on the particularities of one's history and circumstances, such as one's particular location in time and space. For example, I become an approved source of information about where your keys are (at least partly) because I have reliable perception and I am well positioned to see them. Saying that I know where your keys are (at least partly) functions to identify me as a reliable source regarding that information.

But we should also expect our concept of knowledge to have a second function: that of identifying "secondary" sources of information, or sources where information is reliably stored for general purposes of retrieval and flow. One's status of this sort would again depend on both one's general capacities and special circumstances, but now of a different sort. It would depend on those characteristics associated with one's social role as caregiver, teacher, expert, etc. Saying that S knows p serves to identify her as a reliable source of information regarding p, either a primary source or a secondary source. Primary sources originate a flow of information. Secondary sources keep information flowing.

Again, suppose that these speculations about our concept of knowledge are correct. It would then be reasonable to expect a disjunctive concept of testimonial knowledge. One way to get testimonial knowledge would be from an original source. But that would require a capacity to identify an original source as such; i.e. it would require the capacity to discriminate reliable from unreliable testimony. But a second way to get testimonial knowledge would be to tap into the socially approved flow of information. This second source of testimonial knowledge would come for free, or at least cheaper, courtesy of the social system.

There are independent reasons, then, in favor of a disjunctivist account of testimonial knowledge. And in fact, I do not think that the resulting view is unattractive. I do think there is a better approach available, however – one that retains the attractive features of disjunctivism while adding some more of its own.

e. A Generality Problem?

The last approach to Goldberg's Puzzle that I want to consider treats it as a kind of generality problem. To see the motivation for this approach, it will be helpful to first consider a more famous generality problem.

"The Generality Problem for Reliabilism" was first raised as a problem for reliabilist theories of justification.[7] Such theories want to understand epistemic justification in terms of reliable cognitive processes. Roughly formulated, the idea is that a belief is justified just in case it is produced by a reliable cognitive process type. But how do we specify the process type that is relevant for theoretical purposes? For example, it seems plausible to understand perceptual justification as true belief

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

arising from reliable perception. But how are we to specify the relevant process of perception? If we specify that process type very narrowly, in terms of very detailed conditions, for example, then nearly any true perceptual belief will be the result of reliable perception. But if we specify the process type very broadly, in terms of any possible conditions, for example, then no perceptual beliefs will be the result of reliable perception. The problem, then, is to specify the levels of generality that are relevant for theoretical purposes, i.e. for purposes of understanding which beliefs count as epistemically justified and which do not. Put differently, the problem is to specify "justification-relevant" cognitive process types.

Consider now the Problem of (Epistemically) Undiscriminating Children. We described that problem as follows: How is it that small children can come to know from their caregiver's testimony, even if they are themselves unreliable consumers of testimony? How is it, given the *lack* of knowledge in Room Full of Liars? One way to answer the question is to be specific about the processes used to form the beliefs in question. Babe forms his belief on the basis of testimony from his mom. That is a locally reliable process in Babe's environment: not easily could Babe go wrong, using that process, in the environment that Babe is in. Gabe forms his belief on the basis of testimony from his teacher. That too is a locally reliable process in Gabe's environment. Sid forms his belief on the basis of testimony from someone in the room. That is a locally unreliable process in Sid's environment, since Sid is in a room full of liars. The present point is that we can discriminate the good cases from the bad cases by being specific about the process types used to form the various beliefs in question. To put things differently, we can discriminate the good cases from the bad cases by adjusting the level of generality by which we describe the process types in question. Of course that by itself is *ad hoc*. The trick is to make the case that we are specifying the *right* levels of generality for theoretical purposes, i.e. for purposes of understanding which beliefs count as knowledge and which do not. What we need now is a theoretical motivation for specifying knowledge-relevant processes as we have.

That motivation comes from two plausible lines of thought. The first is that relevant levels of generality are at least partly a function of our interests and purposes as information-sharing beings. This line of thought is expressed in the following passage from Ernest Sosa.

> We care about justification because it tends to indicate a state of the subject that is important and of interest to his community, a state of great interest and importance to an information-sharing social species. What sort of state? Presumably, the state of being a dependable source of information over a certain field in certain circumstances. In order for this information to be obtainable and to be of later use, however, the sort of field F and the sort of circumstances C must be projectible, and must have some minimal objective likelihood of being repeated in the careers of normal members of the epistemic community. (1991, 281–2)

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

For present purposes we can sum up the idea as follows: Our concept of knowledge answers to our interests and purposes as social, information-sharing beings. In order to do so, it must track "fields and circumstances" that are relevant to those interests and purposes, i.e. fields and circumstances that cut up our cognitive activities at useful joints. Again, the first line of thought is this: Relevant levels of generality (i.e. those specifying knowledge-relevant cognitive processes) are at least partly a function of our interests and purposes as information-sharing beings.

The second line of thought is the one we have already seen from Craig: that our "interests and purposes" as information-sharing beings include our interest in identifying reliable sources of information. More specifically, they include our interest in identifying secondary sources of information, or those sources that facilitate reliable information flow. Wedding these two lines of thought, we get the following result: Relevant levels of generality are at least partly a function of our interest in identifying reliable secondary sources of information.

To complete the argument we need one more plausible idea: that reliable secondary sources are (at least sometimes) specified by virtue of their social status. That is, one way we identify people as reliable sources of information is by virtue of their social roles as parent, nanny, teacher, expert, etc. This gives us the theoretical motivation that we were looking for. That is, it gives us a motivation for incorporating social roles into the specification of knowledge-relevant cognitive processes.

We can summarize the present solution to Goldberg's Puzzles as follows. First, recall the Problem of (Epistemically) Undiscriminating Children: How is it that small children can come to know from their caregiver's testimony, even if undiscriminating vis-à-vis unreliable testifiers in the vicinity? Answer: By forming their beliefs on the basis of testimony from their caregiver specifically, which is a reliable process. In effect, we deny (c) above: that young children are not reliable consumers of testimony. They are, so long as they are consuming the testimony of their caregivers. Moreover, the solution generalizes. How is it that people can learn from the testimony of "socially approved" sources in general (parents, teachers, experts and definitive sources) even when the hearers themselves are locally unreliable consumers of testimony? Answer: By forming their beliefs on the basis of testimony from approved sources specifically, which is a reliable process. Finally, the present solution is not *ad hoc*. There are independent reasons for specifying knowledge-relevant cognitive processes in the way required by the solution.

3. CONCLUSION

Why prefer the present approach, which treats Goldberg's Puzzle as a kind of generality problem, to disjunctivism as described above? One reason is that the present approach is more elegant. Specifically, it allows a unified account of testimonial knowledge rather than a disjunctive account. If we are reliabilists about

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

knowledge in general, then the approach is more elegant still, since it understands testimonial knowledge within a more general reliabilist framework.

Finally, we may note that the present approach is consistent with agent reliabilism in particular, and therefore with a virtue-theoretic version of reliabilism. On an agent reliabilist account, knowledge in general is the product of agent reliability. Put differently, it is the result of agent powers or abilities. Testimonial knowledge is a species of this, the result of abilities involved in forming testimonial beliefs. On a virtue-theoretic account, knowledge in general is a kind of success from ability, and testimonial knowledge is a kind of success from ability in cooperation with others. As such, knowledge in general and testimonial knowledge in particular can be understood as instances of a more general normative phenomenon, inheriting various aspects of that familiar kind of normativity. Insofar as these approaches to knowledge and epistemic normativity have their own advantages, the present proposal inherits these as well.[8]

In sum, one point in favor of the present proposal is theoretical elegance, in virtue of the account of testimonial knowledge here offered, and in virtue of its easy assimilation within attractive broader frameworks. Nevertheless, the present approach faces a number of important objections. I will end by considering three of these. Answering these objections requires clarification of the proposal as so far presented.

The first objection comes in the form of a purported counterexample.

> **Room Full of Liars with Expert**. As in Room Full of Liars, gullible Sid is in a room full of inveterate liars. He immediately and uncritically believes everything each of them says. At one point he happens to bump into Dr. Hothouse, an expert on global warming and the only reliable person around. Hothouse reliably tells Sid some facts about global warming trends, and (as a matter of course) Sid believes him.[9]

The objection is that Sid here forms his belief on the basis of testimony from an expert, i.e. someone with socially approved status, and yet the intuition remains that Sid does not know.

To handle this case we have to be clear about how the present approach proposes to specify the knowledge-relevant processes in question. The idea is that social roles are tied to social practices, which practices constitute reliable and socially important ways of transferring information. Accordingly, the process types we want to specify are those embedded in these practices. Put another way, the mere fact that a source has some particular social status is not itself epistemically relevant. What is relevant is that the source has that status in the context of a reliable social practice. Since no such practice is in play in Room Full of Liars with Expert, the account need not count Sid as knowing in that case.

Does this mean that experts such as Hothouse have special epistemic significance only at conferences and other formal venues? No, because not all of the practices in which experts play a special epistemic role are formal. There are

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

informal practices involving the transmission of knowledge through experts (and other socially approved sources) as well.

The same consideration addresses another kind of example.[10]

Professional Conference with Gullible Bystander. Gullible Sid walks into a hotel where, unknown to him, a professional conference on global warming is taking place. He overhears Hothouse reliably reporting some facts about global warming and, as usual, believes what he hears.

The intuition that Sid does not know can be explained as above: Although Hothouse has special status as a source of testimony, Sid is not participating in a practice where that status is relevant. Accordingly, it would be misleading to describe Sid as "forming his belief on the basis of expert testimony." In a sense he does, but not in the sense that is epistemically relevant.

The next objection to be considered employs a different sort of case.

Bad High School. Hank is a high school student with normal cognitive capacities. One day Mr. Burnout, who is somewhat fuzzy about the details of American history, tells his class that the U.S. Constitution was ratified in 1789, and Hank believes him. But today was a good day. Mr. Burnout could easily have told his students that the date was 1776, and Hank would have believed that.

The intuition is that Hank does not know on the basis of Mr. Burnout's testimony. Yet the right sort of social status and the right sort of social practice are in place.[11] The case can be explained, however, by noting that Hank is not locally reliable in the present environment. That is, using the socially approved process he does, Hank could easily go wrong in the environment he is in.

This example helps to clarify the nature of the proposal. The proposal is not that all true belief based on socially approved testimony, or on an approved testimonial source, counts as knowledge. Rather, it is that we properly specify knowledge-relevant processes in terms of social roles and social practices. Thus specified, it remains an open question whether any such process issues in locally reliable belief. The answer to that question depends on further facts about the case, including facts about the believer's circumstances and environment. I take it that this is in line with reliabilist proposals in general; that is, reliabilists do not hold that knowledge is true belief grounded in a process that is generally reliable. Rather, they hold that knowledge is true belief that is grounded in a process that is reliable relative to S's environment.

The final objection to be considered is that generality problems remain.[12] To see that this is so, consider the way we treated High School and Bad High School. In the former case, we said that Michael knows because he forms his true belief on the basis of a locally reliable process. That is, he believes what a history teacher tells him, in an environment where that process reliably leads to true belief. In the latter case, however, we said that Hank does not know, because he forms his belief on the basis of a locally unreliable process. That is, he believes what a history teacher

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

tells him, in an environment where that process does not reliably lead to true belief. But consider, the cases came out right only because we specified environments in a particular way. Suppose that Michael and Hank go to school in the same school district, or in the same state. If we specified environments in terms of entire school districts, or entire states, then Michael and Hank would have to be considered as in the same environment, and their beliefs considered as equally locally reliable.

A similar point can be raised about the way we are specifying method or process. We specified the process used by Michel and Hank as "forming one's belief on the basis of what one's teacher says." But we could have specified it in terms of "history teacher," "history teacher with an M.A.," "history teacher who is not burned out," etc. Depending on how we specify processes, we will get different levels of reliability and therefore different judgments about which cases count as knowledge. In short, the proposal that we specify processes in terms of social roles and social practices does not by itself solve this kind of generality problem.

The answer to this objection is that the present proposal was never meant to solve all generality problems. Rather, it was meant to solve a more restricted problem: How is it that young children can learn from the testimony of their caregivers (their parents, teachers, and nannies, for example) even when the children themselves are undiscriminating consumers of testimony? Or more generally, How is it that people can learn from the testimony of "socially approved" sources (parents, teachers, experts and definitive sources) even when the hearers themselves are locally unreliable consumers of testimony? The answer proposed is this: We can see the knowers in such cases as locally reliable after all, if we are careful in specifying the method or process by which they are consuming testimony. In short, such methods or processes should be seen as constituting social practices that serve reliable information flow. This answer does not suppose that all generality problems are thereby solved. Rather, it supposes that, however such problems are solved, distinguishing cognitive practices in social terms will be part of the story.[13]

REFERENCES

Conee, E. and R. Feldman. 1998. "The Generality Problem for Reliabilism." *Philosophical Studies* 89: 1–29.

Craig, E. 1990. *Knowledge and the State of Nature*. New York: Clarendon Press.

Goldberg, S. Forthcoming a. "Testimonial Knowledge in Early Childhood, Revisited." *Philosophy and Phenomenological Research*.

Goldberg, S. Forthcoming b. *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge: Cambridge University Press.

Goldman, A. 1979. "What is Justified Belief?" In G. Pappas (ed.), *Justification and Knowledge*, pp. 1–23. Dordrecht: D. Reidel Publishing.

Greco, J. Forthcoming. "The Nature of Ability and the Purpose of Knowledge," *Philosophical Issues*.

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

Greco, J. 2003. "Knowledge as Credit for True Belief." In M. DePaul and L. Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 111–34. Oxford: Oxford University Press.

Green, A. Unpublished. "The Epistemic Scaffold: Reductionism, Testimony, and the Infant/Child Objection." Saint Louis University.

Koenig, M. and P. Harris. This issue. "The Basis of Epistemic Trust: Reliable Testimony or Reliable Sources?" *Episteme, A Journal of Social Epistemology* 4(3): 264–84.

Lackey, J. 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.

Sosa, E. 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.

NOTES

1 Goldberg (forthcoming a) and (forthcoming b, especially ch. 8).

2 The two criteria are not independent, of course, since a plausible account of the nature of testimonial knowledge must yield a plausible account of the range of testimonial knowledge.

3 This is Goldberg's term, which he defines as follows:

> Let us say that a subject *S*'s belief that *p*, formed on occasion *O* via method *M*, was formed in a manner that is *locally reliable* when the following counterfactual holds: if *S* were to believe that *p* via *M* in a situation relevantly like *O*, *p* would be true.

Goldberg goes on to note,

> The local reliability of hearer *H*'s testimonial belief that *p,* formed on an occasion *O*, would then require the following: if the hearer *H* were to believe that *p* through testimony in a relevantly similar situation, *p* would be true. (Goldberg forthcoming b, ch. 5).

So defined, local reliability pertains to beliefs. We can say that a method M is locally reliable, relative to environment E, just in case M would give rise to locally reliable beliefs in E. We can say that a subject S is locally reliable (when using method M), relative to an environment E, just in case S would form locally reliable beliefs (when using M) in E.

Local reliability amounts to what I have elsewhere called "weak safety." Keeping method of belief formation and other relevant features fixed,

> S's belief that p is *weakly safe* just in case, in close worlds, usually if S believes p then p is true. Alternatively, in close worlds, almost never does S believe p and p is false. (Greco 2003)

4 The examples do not have the same structure in *all* respects, of course. One difference between the good cases and the bad cases is that the protagonists in the good cases are unreliable discriminators of testimony in general, whereas the protagonist in the bad case is merely locally unreliable.Sid fails to discriminate between reliable and unreliable testifiers in his local environment, a room full of liars, although presumably he has normal capacities for discriminating testimony in general. But notice that this makes the protagonists in the good cases *worse* off epistemically, not *better* off. So this is an

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press

odd feature of the cases considered so far – the protagonists who seem worse off epistemically (young children with limited cognitive capacities) are the ones that come out knowing. The protagonist who seems better off (an adult with normal capacities) comes out not knowing.

5 If we think of environments as constructed out of actual states of affairs, then Lizabeth is "locally reliable" in the sense of "reliable in her actual environment." For example, she easily discriminates between the DSM-IV, popular self-help books, and fashion magazines for their reliability regarding eating disorders. Nevertheless, there are close possible worlds where the DSM-IV is not a reliable source, and yet Lizabeth continues to treat it as such. She is not "locally reliable" in this broader sense.

6 For a nice overview of relevant empirical evidence, see Koenig & Harris (this issue). See also sources cited in Goldberg (forthcoming b), Lackey (2008), and Green (unpublished).

7 Cf. Goldman (1979) and Conee and Feldman (1998).

8 I discuss some advantages of this approach to knowledge, including testimonial knowledge in particular, in Greco (forthcoming).

9 This example was raised by Sandy Goldberg in discussion.

10 Also suggested by Goldberg.

11 This sort of concern was raised by Al Casullo in discussion.

12 Concerns of this sort were raised by Alvin Goldman and Peter Lipton in discussion.

13 This paper benefited from discussion at the fourth annual *Episteme* conference at Rutgers University. Special thanks to Al Casullo, Paul Faulkner, Sandy Goldberg, Alvin Goldman, Jennifer Lackey, and Peter Lipton for helpful comments and discussion there. Thanks also to an anonymous referee, who provided helpful comments on a draft of the paper.

John Greco is the Leonard and Elizabeth Eslick Chair in Philosophy at Saint Louis University. He received his Ph.D. from Brown University in 1989. Recent publications include: *Sosa and His Critics* (ed.) (Blackwell, 2004) and *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry* (Cambridge University Press, 2000).

https://doi.org/10.3366/E1742360007000123 Published online by Cambridge University Press