

If that were true I would have heard about it by now

Sanford C. Goldberg
Department of Philosophy
Crowe 3-179
1880 Campus Drive
Northwestern University
Evanston, IL 60208-2214
s-goldberg@northwestern.edu

Social Epistemology Conference
University of Stirling
Scotland, UK
September 1, 2007

ABSTRACT

Most work on the epistemology of testimony focuses on issues pertaining to what I will call the *testimony-to-truth* transition a transition one employs whenever one moves from having observed a piece of testimony to the effect that p , to forming a belief (on that basis) in the truth of the proposition that p . Regarding this transition, various questions have been raised in the epistemological literature: when is a hearer entitled to make such a transition (in the course of belief-fixation)? Under what conditions does a belief fixed by way of this transition amount to *knowledge* that p ? But to my mind no one working on testimony has recognized the importance to epistemology of the converse transition. One employs this transition whenever one moves from the fact that one has not (yet) come across testimony that p , to the falsity (or likely falsity) of p . Although a straightforward reliance on this transition would exhibit the fallacy of ignorance, there can be little doubt that we often do reason in this way; and in this paper I will argue that we often do so to good epistemic effect. I develop this idea by speaking about our reliance on others, not just as providing us with sound (and reliable) testimony, but also as providing us with something approximating what I will call (*reliably*) *complete coverage* in various domains. A full accounting of the social nature of knowledge, and of testimonys role in reinforcing this social nature, is incomplete without an account of this sort of reliance on ones community.

1. Cases

Over lunch, you and a friend are having a discussion about US foreign policy.

She raises the question whether weapons of mass destruction (WMDs) have been

found in Iraq. You say that they have not. (You reason that if such weapons had been found, you would have heard about it by now.)

As his mind is wandering, the thought occurs to Smith out of the blue that the president announced a new major change in foreign policy last week. Believing that if the President had done so he (Smith) would have heard about it by now, Smith forms the believe that the president did not announce a new major foreign policy change last week.

Listening to the locals discuss the goings-on of various Hollywood celebrities, McSorley overhears a particularly juicy tidbit regarding Toothy Thompson, a particularly famous celebrity who is nearly universally regarded (including by McSorley) as a person of high integrity. According to the speaker, Toothy has actually lead the life of a degenerate who has barely escaped legal prosecution on various occasions. Given Toothys reputation, McSorley rejects the testimony: McSorley thinks to herself that if any of this were true, she would have heard about it by now (from some more familiar source).

The phenomenon illustrated in these cases is rather common. In them, the fact that a subject has never come across a piece of testimony to the effect that p is used as support for her belief in not- p . This support can take the form of a reason that offers (further) support for believing something she believed all along (the WMD case); it can take the form of a reason supporting the formation of a belief she did not previously have (the change in foreign policy case); or it can take the form of a reason to disbelieve some piece of presently-observed testimony opposing a belief she presently has (the Toothy Thompson case). I do not claim that these options are

exhaustive.

Regarding these kinds of case, various questions emerge. Under what conditions is this sort of reasoning offered? Under what conditions is there justification for a belief formed on the basis of this sort of reasoning? Under what conditions does this sort of reasoning justify the repudiation of a piece of testimony to the contrary? And what does this tell us about the nature of testimony? of our reliance on our peers? or of the organization of our epistemic communities?

2. Epistemic Coverage

I begin first with the descriptive task, that of saying when it is that reasoning of the above sort is offered.

First, it is clear that in a typical case a hearer will reason in the way above using the silence of a standard source (the newspaper; a trusted blog) to form or sustain a belief, or to repudiate a piece of testimony to the contrary only if she presupposes and/or explicitly believes certain things. In particular, the hearer will believe or presuppose that the proposition in question concerns a subject-matter regarding which there is some sub-group of members of her community (the standard source) who regularly report about such matters; that the standard source is, or is likely to be, reliable in uncovering and subsequently publicizing truths about such matters; that such a source had sufficient time to discover the relevant facts and report on them; that she herself (the hearer) would likely have come across such a report, had one been made; and that she has not come across such a report. If the hearer fails to believe or presuppose any of these five things, she is not likely to reason in the way indicated. (Or if she does, it

would appear that her doing so is unjustified; more on this in section 4, below.)

Let us call a belief that is supported by this sort of reasoning whether the belief is one that is newly-formed or sustained, and whether the supporting reasoning was explicit or only implicit a *coverage-supported* belief. Although I will have more to say on the matter below, for now I simply note that the relevant notion of coverage is seen in the believers reliance on a source to be both reliably apprized of the facts in a certain domain, and reliably disposed to report on the obtaining of facts in that domain (when those facts are known by the source).

The sort of reliance exhibited by our hearer towards her standard source in cases of coverage-supported belief is not often discussed in connection with the epistemology of testimony. That literature is primarily focused such questions as: under what conditions is a hearer justified in accepting another speakers say-so? And under what conditions does a hearer acquire *knowledge* through such acceptance? We might say that, in asking such questions, we are examining our need (as hearers) for the reliable consumption of reliable testimony. In the case above, by contrast, we are going beyond this to examine a *further* type of reliance. I will designate this further type *coverage-reliance*. A hearer *H* exhibits coverage-reliance towards some source(s) in domain *D* when she relies on the existence of sources who are such that, if *p* is a true proposition in *D*, one or more of the sources will publicize or broadcast that *p*, in such a way that *H* herself will come across the report.

There is a very rough analogy here with the issues of the soundness and completeness of a formal system. We want a formal system to be sound: we want it to be the case that anything provable within the system is in fact valid. But we want a

formal system to be complete as well: we want it to be the case that any valid sentence is provable in the system. (This is why incompleteness results are so interesting and provocative.) In the rough analogy I am developing, one's standard source is analogous to proof in the formal system, in that the source's pronouncements offer a candidate way for getting at the relevant truths. We want that source to be sound: we want it to be the case that anything attested to by the source is in fact true. But often we would also like that source to be complete (or near-enough) as well: we want it to be the case that

whenever there is a (relevant) truth in that domain, it is reported by the source.

The analogy is rough along various dimensions. For one thing, talk of the soundness and completeness of the system of testimony would lead one to expect that this system can be precisely characterized and this is something that I doubt. In addition, whereas the question of the soundness and completeness of a given formal system may hold a real interest for us, no one would ever suppose that there is a sound testimonial source (even the most scrupulous of sources testifies to something false on occasion), nor would anyone suppose that there is an interesting domain in which some source is complete (interesting facts escape even those most disposed to publicize all that they notice in that domain). Nor are these the only flaws with the analogy.

But the analogy, though imperfect, is suggestive. It makes clear how the discussion in the epistemology of testimony has been almost universally one-sided. The literature has been devoted almost exclusively to issues pertaining to the soundness of the testimonial belief-fixing process, which is to say, (i) the reliability and trustworthiness of testimonies that are offered, and (ii) the reliability of hearers in distinguishing reliable from unreliable testimonies. What the literature has not focused

on indeed, what it has not been so much as raised as an issue to be investigated are issues pertaining to *coverage*, to the completeness of testimony within ones epistemic community. By focusing on these issues, we can bring to the fore yet another way in which the phenomenon of testimony pushes us ever more in the direction of an irreducibly social epistemology.

3. Testimony transitions

It will be helpful to have some more terminology on the table. Testimonial belief is belief formed through (and grounded on) an acceptance of a piece of testimony. In such cases a hearer observes a piece of testimony and, having accepted the speakers say-so, comes to form a belief in the attested content. Such cases instantiate what I will call a testimony-to-truth transition: the move is from the observing of a piece of testimony that p , to the (formation of a) belief in the truth of p . In saying this I do not mean that this transition is consciously undertaken, or that the hearer herself would articulate the transition as the basis of her belief. The point, rather, concerns the epistemic ground of the belief. The nature of this ground is brought out by the sort of *ex post facto* rationalization that could be offered (by an epistemologist, if not by the subject herself) for the subjects testimonial belief. The rationalization itself displays relations of epistemic dependence holding among pieces of information within a given subjects epistemic perspective. Stripped to its basics, such a rationalization would involve something like the following inference:

- (1) S testified that p .
- (2) If S testified that p , then p .

Therefore,

(3) p .

What I am calling the testimony-to-truth transition is manifested in the conditional statement (2). Epistemologists interested in testimony typically focus on issues pertaining to (2), and to a hearer's entitlement to rely on (2). Such questions include: does (2), or something like it, derive from a basic epistemological principle, or does it stand in need of further (presumably empirical) justification? And: under what conditions is a hearer entitled to employ or rely on (2) on a given occasion? I will not address these questions here; versions of them are addressed at length in the epistemological literature on testimony (and in several of the papers presented at this conference). My present point is one on which all parties to the debate should agree: that in an *ex post facto* rationalization of a testimonial belief, something like this stripped-down inference will be present. In this way we can say that the process involved in testimonial belief-fixation instantiates the testimony-to-truth transition.

Now consider what we might call the phenomenon of *coverage-supported belief*, where one forms or sustains a belief in the negation of a proposition, on the basis of the fact that one's preferred source(s) has/have not (yet) so attested. What sort of *ex post facto* rationalization might a coverage-supported belief receive? Here I focus on a pure case: one where the belief in the negation-proposition is justified solely on the basis of coverage considerations (e.g. no other information is brought to bear on the likely truth of the negation-proposition in question). Stripped to its basics, a rationalization in such

a pure case would involve something like the following inference:

(4) None among the trusted sources Σ testified that p .

(5) If p , then one among the trusted sources Σ testified that p .

Therefore,

(6) It is not the case that p .

Here, what I would call the truth-to-testimony transition is manifested in the conditional (5).

It may seem curious that the truth-to-testimony conditional (5) is in the material mode. After all, the reasoning in cases of coverage-supported belief typically involves the subjunctive conditional (witness the title of this paper). To some degree this sense of curiosity can be alleviated by noting that (5) follows from the claim, at the heart of coverage-supported belief, that all (relevant) truths are/have been reported. It is also worth noting a parallel with the testimony-to-truth conditional. The reason we want trustworthy (reliable; competent; what-have-you) testimony is that we want it to be the case that we can come to know something about the world through accepting another's relevant say-so on the matter; and this, in turn, requires the truth of a conditional to the effect that if our source asserted that p , then p . Of course, the property of trustworthiness itself is not best characterized in terms of this material conditional. Trustworthiness appears to be more robustly modal: we would say that a piece of testimony is trustworthy (reliable; competent) when it is such that (e.g.) it *wouldnt* be offered if it *were* false. We want testimony to have this property precisely because we want the transition, from the observation of a piece of testimony, to the truth of the proposition attested to, to be warranted. In precisely the same fashion, I acknowledge that coverage is a more robustly modal property than that captured in the material conditional in (5); but the reason that we care that our community exhibit coverage

(when we form coverage-supported belief) is that we want it to be the case that the transition, from the absence of (observed) testimony to the effect that p , to not- p , is itself warranted. It is this transition that is captured by the material conditional (5).

It is perhaps unsurprising that, to date, the epistemology of testimony literature has focused on the testimony-to-truth transition. As noted above, this transition is instantiated whenever one forms the belief that p through one's acceptance of testimony that p . What I would like to do in the remainder of this paper is to frame the issues I see arising from the (implicit) employment of the truth-to-testimony transition. Regarding this transition, we can ask questions analogous to those asked in connection with the testimony-to-truth transition. These questions include the following: under what conditions is one epistemically entitled to rely on the truth of (5)? Under what conditions (if any) does one's reliance on (5), together with one's failure to have observed any of the trusted sources attest to p , suffice to underwrite one's justified belief (or knowledge) that not- p supposing that one does not have such justified belief (or knowledge) through other sources? In pursuing these questions, I hope to deepen our appreciation for the ways in which beliefs supported by *the noted absence of* particular testimony, like beliefs supported by testimony, demand an epistemological account that is interestingly social.

4. Community and Individual

Under what conditions is a hearer epistemically entitled to rely on the truth of the truth-to-testimony conditional, (5), in the course of belief-fixation? I want to approach this question by answering another question first: what sort of facts (regarding one's

community and ones place in it) would render it epistemically fruitful for a hearer to rely on the truth-to-testimony conditional? We might then answer our original question by saying that a hearer is entitled to rely on this conditional when she is entitled to believe that things are this way.

We do well to recall the points made above (in section 2) in connection with what is presupposed in the typical case by a hearer exhibiting coverage-reliance. Regarding the conditions on ones community, we noted that the hearer presupposes that the proposition in question concerns a subject-matter regarding which there is some sub-group of members of her community (the standard source) who regularly report about such matters; and that the standard source is, or is likely to be, reliable in uncovering and subsequently publicizing all relevant truths about such matters.

Let us begin with the presupposition that some sub-group of members of the subjects community regularly report about such matters. (Call this the *source-existence* condition.) This subgroup might be one that is traditionally recognized by virtually everyone in the community (the traditional print and TV media, for example). Or it might be that the subgroup in question is one that is specific to the subject herself (a group of her friends, say, whose members are particularly interested in, and disposed to publicize to the others what they have learned about, some subject-matter). The interesting question concerns the *regularity* of the relevant reports. Suppose you rely on me for what you know about the weather in Chicago. If my reports are less frequent than what youd like in the way of information about the Chicago weather, then I will often leave you without information you want. Clearly, the greater the discrepancy between the frequency of my reports and your demand for current information, the less

well-informed (relative to your information-needs) you will be at any given time.

At the same time, the notion of regularity in the regular reporting condition is not as straightforward as it may seem at first. *Regular* does not imply *frequent*. Suppose you want information regarding the weather each day, and I update you only quarterly. Even so, if each of my quarterly reports is exhaustive, reporting on the weather during every day in the interval, then this might satisfy your informational needs. (You might be a scientist studying weather patterns in Chicago: you need weather reports for each day, but they need not be timely.) In that case you might well be entitled to rely on my reports, so that if someone tells you something about the Chicago weather on a day that I have already reported on, but where my report did not report what she told you, you might well conclude by rejecting her testimony on the grounds that if that were true you would have heard about it (from me) by now. This example suggests that the regularity in the regular reporting condition appears to mean something like periodic, where the interval between reporting periods is one that falls within a standard expectation formed by the coverage-relying subject herself.

Such expectations of regular reporting might be formed in a variety of distinct ways. They might be formed on the basis of the sort of information in question: the average citizen expects the goings-on in the nations capitol to be reported daily, whereas she harbors no such expectation for the goings-on in the Physics Department at State U. Alternatively, the expectations might be formed on the basis of (ones beliefs regarding) the particular source in question: you expect your (daily) newspaper to be delivered on a daily basis, but you probably expect the news from a neighbor down the street on a much less regular basis (indeed, the intervals between these reports might

not be regular at all). Or the expectations may be the effect of some sort of ongoing coordination between two or more parties: I might give you the standing directive to report to me as soon as you observe any developments in the recent negotiations between labor and management (and we might agree to revise this directive, under various conditions, to make it more general, or more specific, in various ways).

In addition to the source-existence condition there must be a source that regularly reports on the facts in the domain in question there is a second presupposition that amounts to a condition on the community: the relied-upon source must be, or be likely to be, reliable in uncovering and subsequently publicizing truths about the domain in which the subject is exhibiting coverage-reliance. I will call this the *reliable-coverage* condition. Let D be a domain of interest to subject H , let p be any proposition in D regarding whose truth H might take an interest, and let α be some

source on whom H could rely on matters pertaining to D . Then we can say that
(CR) α is *coverage-reliable* in D relative to $H =_{\text{def}}$

it is likely that α (a) will (investigate and) reliably determine whether p , (b) will be reliable in reporting the outcome of that investigation, and (c) will satisfy both of the previous two conditions in a manner conforming to the regular reporting condition above.

With this as our basic notion, we can then go on to define other, related notions. For example, in many cases a subject does not rely on any *particular* source, but instead relies on the fact that the relevant information would be publicized by *some source or other*. We can capture this notion in terms of generic coverage-reliance, as follows:

(GC) There is generic coverage-reliability in D relative to $H =_{\text{def}}$

There is some source or other in H 's community that is coverage-reliable in D relative to H .

And I am confident that there will be other notions in the vicinity worth capturing. (I leave this for future work.)

As it is formulated (CR) captures what we might call a *non-attuned* sort of coverage-reliance, one in which the subject H coverage-relies on a source, α , who may or may not know that H is so relying, and who (even if α knows that H is so relying) may or may not know H 's specific informational needs and expectations. But there can be other cases with a source that is attuned to the scope and informational needs and expectations of its audience; and such a source will be one that explicitly aims to render itself coverage-reliable relative to that audience. I offer the following as capturing this

notion of attuned coverage-reliance:

(AC) α exhibits *attuned* coverage-reliability in D relative to $H =_{\text{def}}$

(I) α has knowledge of both the scope of the audience ϕ that relies on it for (some of) their informational needs, and the information-relevant expectations that members of ϕ have with respect to α itself; (II) H is in ϕ ; and (III) for any proposition p in D , if it is reasonable for α to suppose, *both* that (i) members of ϕ would be interested in the truth-value of p , *and* that (ii) members of ϕ are likely to rely on α for the information whether p , then it is likely that α (a) will (investigate and) reliably determine whether p , (b) will be reliable in reporting the outcome of that investigation, and (c) will satisfy both of the previous two conditions in a manner conforming to the regular reporting condition mentioned prior to the discussion of (CR).

What is more, we might distinguish the attuned coverage-reliability captured by (AC) with a sort of coverage-reliability where the scope and informational expectations of the relevant audience are *common knowledge*, had by both members of the audience and by the source itself (and where both sides know this of the other side etc.). The difference between *common-knowledge coverage reliability* (as we might call it) and the

sort of case captured by (AC) is that in the common-knowledge case the fact of attunement is known to both sides, not just to the source.

We have been exploring the conditions on the subject's community, if her coverage-reliance tendencies are to be epistemically fruitful. What can we say regarding the conditions on the coverage-relying subject, *H*, herself? Once again we can begin with one of the presuppositions listed in section 2: (the subject presupposes that) she herself must be likely to have come across whatever relevant reports were offered by the source(s) on whom she was relying. (I will call this the *reception condition*.) To be sure, the subject can satisfy the reception condition without having to receive the relevant reports directly from the source itself: it may be that there is a more extensive chain of communication linking her to the source. So long as the communication chain is itself both sufficiently reliable (preserving whatever relevant content there was in the original report(s)), sufficiently complete (passing on all of the relevant reports), and sufficiently well-publicized (reaching at least some of the places where the subject is likely to encounter the transmitted message), the subject is likely to come across the reports from the source(s) on which she is relying.

There are other presuppositions concerning *H* herself: to date, she must not have encountered any report attesting to the truth of the proposition in question, over a period of time long enough to allow for the discovery and reporting of any relevant fact. We can break this up into two conditions, the *silence* condition (no relevant report observed) and the *sufficient time* condition (time was sufficient for the source to discover and publicize relevant information, had there been any.)

We have, then, what I allege are five jointly sufficient conditions on the epistemic

fruitfulness of coverage-supported belief: the source-existence condition, the coverage-reliance condition, the reception condition, the silence condition, and the sufficient time condition. But precisely what is the relevance of these conditions to the justification of coverage-supported belief? On the assumption that a belief can be justified without being true, we need not suppose that all five conditions must be true in order for a coverage-supported belief to be justified. Rather, the plausible initial proposal is that a coverage-supported belief is justified so long as the subject is (or would be) justified in believing that each of the conditions holds. With this, of course, we are staring in the face of some vexed matters about the nature of justification generally (and perhaps of testimonial justification in particular). For example, need the subject, *H*, have *positive reasons* to regard these five conditions as satisfied? Or does it suffice that her coverage-supported beliefs are e.g. reliable as they would be if she were disposed to form coverage-supported beliefs only when these conditions are in fact satisfied (and when she has no reasons to think that they are not satisfied)? This issue, of course, is a special case of a much larger question, concerning the nature of justification itself: does something like reliability suffice for justification, or does justification require the presence of some sort of adequate positive reasons? Although my sympathies lie with the reliabilist, I do not have the space to address the larger question here; I can only suggest how matters will be seen according to whether one's theory of justification is reasons-centered or reliabilist or, more generally, internalist or externalist.

Our question concerns the conditions on a subject's epistemic entitlement to rely on the truth of the truth-to-testimony conditional, (5), in the course of the fixation of a coverage-supported belief. Suppose one is an epistemic internalist regarding

justification. Then it would seem that justification for coverage-supported belief is not a matter of one's sensitivity (in the fixation of coverage-supported belief) to the actual satisfaction of the five conditions above, so much as it is a matter of the hearer's *reasons* for thinking that these conditions are satisfied in the present case: her

coverage-supported belief is justified if and only if these reasons are adequate.

Such a view puts all of the epistemic burden on the hearer, rather than on her community: if her coverage-supported belief is to be justified, she must have good reasons to think that there is a relevant coverage-reliable source whose reports she would likely have come across by now. It is tempting to suppose that, on such a view, there will be no need to acknowledge any distinctly social dimension to the epistemology of coverage-supported belief. Though tempting, such a conclusion should be resisted. For even if one's theory of justification is reasons-centered and internalist, one ought to acknowledge that facts regarding one's community can affect the full epistemological assessment of a coverage-supported belief if not in connection with justification, then in connection with some other epistemological status. This can be brought out in connection with pairs of cases, alike as to the hearer's reasons and the proper functioning of their cognitive processes, but differing in the satisfaction of one or more of the five conditions noted above. My claim is that in such cases, the hearer's coverage-supported beliefs are not equally well-off, epistemologically speaking and that this epistemic difference must be accounted for in terms of differences in the social context of their coverage-supported beliefs.

The following example is meant to illustrate. Let Don and Don* be doppelgängers, both of whom form the coverage-supported belief that the President has

not announced any new major policy initiative. (Neither has heard mention of such an initiative, and each reasons that if the President had announced a new major policy initiative, he would have heard about it by now.) Both Don and Don* would avow the very same reasons for thinking that the five coverage conditions are satisfied in their respective cases; and both are in excellent cognitive condition, with their faculties working properly. In fact, the cases are exactly alike, save with the following difference: where the sources on which Don is relying for coverage continue to operate efficiently and well, as they always have, the sources on which Don* is relying for coverage, though historically as reliable as those on which Don is relying, are (unbeknownst to anyone except the sources themselves) in serious disarray, and have not been operating efficiently or reliably for the past several days. To be sure, all of this has escaped Don*s (and virtually everyone else's) notice: Don* is non-culpably ignorant of these facts. But the result is that, although it is true that the President has not announced any new major policy changes, if he had and this could easily have happened the sources on which Don* was relying for coverage would have failed to report it.

Now I submit that Don*s coverage-related belief is less well-off, epistemically speaking, than is Don's; and I submit that this point should be granted even by those epistemic internalists who regard Don's and Don*s respective coverage-related beliefs as on a par, justification-wise. Nor is this epistemic difference between their cases merely one of Gettierization (Don*s belief being Gettiered, whereas Don's belief is not). It might appear that way, given how I have described Don*s case. But the point I wish to make, regarding a difference in epistemic goodness between their respective

coverage-related beliefs, could be made in a different case, where it is clear that the epistemic differences between the twins beliefs is not to be chalked up to Gettierization. Suppose that, instead of being in serious disarray, the sources on which Don* is relying for coverage are only slightly worse-off than those on which Don is relying for coverage. As it might be: Dons sources would pick up and report on all of the facts that Don*s sources would pick up and report on, and then a few more as well (which would escape the notice of Don*s sources). Here it should still seem clear that Don*s coverage-related beliefs are not quite as well-off, epistemically speaking, as are Dons and yet there should be no temptation to ascribe this difference to Gettier considerations. (We can imagine that this difference in the coverage-reliability of their respective sources makes for a difference in the reliability of their respective coverage-supported belief, but this difference does not make for a difference at the level of knowledge both know, or fail to know.) This is the sort of difference, I submit, that reflects the distinctly *social* contribution in coverage-reliance cases even for those who favor an internalist, reasons-based theory of justification.

Of course, the social dimension of the epistemology of coverage-supported belief will be even clearer if ones theory of justification is not internalist. To see this it will be helpful to work with a standard externalist theory of justification. I propose to use

standard reliabilism, according to which
(JR) A belief is justified iff it is formed and sustained through processes that are (generally) reliable.

One who is a justification-reliabilist will hold that beliefs formed through coverage-reliance on a source are justified iff they are formed and sustained through processes

that are (generally) reliable. And it would seem that they are reliably formed just in case she is disposed to form coverage-supported beliefs only when the five conditions above are satisfied.

In spelling out the reliabilist picture here, it is interesting to focus exclusively on the coverage-reliability condition. (This will enable us to bring out what the reliabilist might see as the epistemic significance of the different sorts of coverage-reliability I mentioned above.) Let us suppose that, in general, sources are more likely to be coverage-reliable when they are attuned to the informational needs of their audience that is, when the case is one satisfying (AC). In that case, there may be a difference, justification-wise, between cases satisfying (CR) but not (AC), and cases satisfying the more demanding (AC). Assuming that we have a situation in which the beliefs in both cases attain a degree of (general) reliability that suffices for justification, beliefs in both cases will be justified; but beliefs formed under (AC)-conditions will in general be better off, justification-wise, than those formed under (CR). What is more, these results do not depend on *Hs* having justified beliefs regarding her sources: so long as she is in fact disposed to rely on sources that in fact are coverage-reliable (and the other conditions are satisfied), her coverage-supported beliefs will be justified. Or so it would seem on a standard externalist (reliabilist) framework.

It is worth making the same point from a slightly different perspective. Many reliabilists will want to know what explains *Hs* reliance (in a given situation) on a source that in fact is coverage-reliable relative to *H*. If the explanation is dumb luck perhaps there were many sources that were not coverage-reliable on the topic at hand, *H* just happened to rely on one that was many reliabilists will hold that *Hs* connection to the

truth is still too lucky to count as reliable in the relevant sense. (This might be the analogue, for coverage-supported belief, of what Goldmans barn façade case is for perceptual belief.) But now suppose that the sources in Σ (the set of sources available to H) are not merely coverage-reliable, but *attunedly* so (and so satisfy (AC)). In that case the explanations for H 's success in relying on a source that is in fact coverage-reliable might be the high percentage of coverage-reliable sources in Σ where this might be explained, in turn, in terms of the (social and economic) pressures for coverage-reliability in the various domains regarding which people in H 's community assign a high utility to having reliable and (relatively) complete information. In this case even the more demanding reliabilist has grounds for regarding as justified the beliefs H forms or sustains through coverage-reliance on sources in Σ . For even if H herself does not have explicit reasons for regarding as coverage-reliable the various sources on whom she relies for coverage, even so this being a case of *attuned* coverage-reliability those sources aim to satisfy the informational needs of subjects like H , and they have a profile of what those needs are. Here, the burden (on ensuring that the conditions on justified coverage-supported belief are satisfied) is not entirely H 's; on the contrary, *her sources themselves* assume some of this burden, since (as part of their very aim as information-sources) these sources assume the burden of rendering themselves coverage-reliable to those who rely on them.

It would seem, then, that justification-internalists and -externalists alike should agree that there is a social dimension to the epistemology of coverage-supported belief. To be sure, the two sides will account for this dimension differently. The main

disagreement will concern whether social considerations can affect a beliefs status as justified, in a way that is independent of being represented in the subjects belief corpus itself: the externalist can accept such a view, while an internalist will dispute it. What is more, while an internalist might have no grounds for treating cases as distinct according to whether the coverage reliance at play is attuned, the externalist may well want to do so, on grounds pertaining to the difference in reliability in cases of attuned vs. non-attuned coverage-reliability.

5. The social epistemology of coverage: some remaining questions

Let us step back and review our phenomenon from the perspective of social epistemology. I propose two notions, both of which emerge from the foregoing discussion of the epistemology of coverage, and both of which are naturally construed as notions from a social epistemology. These are the notions of *newsworthiness* and of *epistemic environments*.

Let us begin with the notion of *newsworthiness*. Above we asked after the conditions under which coverage-reliance is to be an epistemically fruitful way of forming or sustaining belief. Although it does not say so explicitly, it is reasonable to think that the coverage-reliability condition one of the conditions in our jointly-sufficient set requires that there be a meshing between the propositions regarding whose truth community members *assume* there to be adequate (reliable) coverage in the community, and the propositions regarding whose truth there actually *is* adequate (reliable) coverage in the community. It is here that the notion of *newsworthiness* can be of service: the individuals sense of what is newsworthy must eventuate in judgments

that agree, more or less, with the standards of newsworthiness that inform the publishing decisions of the relevant news group(s) in her epistemic community.

It should be clear that one's sense of newsworthiness is relevant to the justification of one's coverage-supported belief: a subject whose sense of newsworthiness is radically different from that of the sources on which she is relying risks forming coverage-related beliefs in propositions regarding which she has no reliable coverage. Following this, I suggest that the notion of newsworthiness has an important, and heretofore largely neglected, role to play in epistemology. The relevant notion is not just that possessed by those who *report* the news, but also by those who *consume* it at least if the latter hope to be able to exploit the coverage they receive. The relevant point here is simply that the individual's notion must be *synchronized* with that of those who are or would be most likely to acquire and disseminate the information in question. No doubt, this process of synchronization is one that develops over time, as an interplay between consumers and producers of news. Individual consumers develop a sense of the sorts of media (and other information-transmitters) with which they regularly interact, of the sorts of news that is regularly reported by those sources, and of the communication channels through which the news is reported. But it can also work the other way, as when the demand for news in a certain domain is unmet, leading either existing media (and other transmission sources) to extend what they consider newsworthy, or to the creation of new media designed specifically to fill the lacuna. (I take it that the internet, and the relatively recent phenomenon of blogging, has only sped up this media-spawning process.)

Much more work could be done investigating the epistemic dimension of the

evolution (within a community) of the notion of newsworthiness. I submit that this will capture part of the *cognitive* dimension of the epistemology of coverage that part of the epistemology of coverage that takes stock of the beliefs and presuppositions of the participating parties (hearers and sources alike). However, the considerations from 4 also suggest that part of the epistemology of coverage will not be in cognitive terms at all, but instead will have to do with features of how well-placed, information-wise, an individual is within her community. Here we see the relevance of the notion of an *epistemic environment*.

Epistemic environments can be evaluated along various dimensions. But if we restrict our attention to coverage-reliance, two such dimensions are particularly salient. Once again, these dimensions reflect the main types of role played by community and individual in the process of coverage reliance.

One of these dimensions concerns an attempt to assess the coverage that is provided within that community. How many domains D are such that there is one or more community groups that take an active interest in discovering and reporting the facts in D ? Within such a domain, how likely is it that if p is true, one or more of these community groups will discover and publicize that p ? What is the nature of the publication process? (Is it through traditional media? On the internet? Etc.)

The second dimension along which to assess an epistemic environment (for the purpose of characterizing coverage-reliance in that community) regards the factors relating to the likelihood that a community member interested in knowing whether p will come across relevant reports made within her community (access assessment). How accessible is the source publication? How likely is it that an individual with

informational interests not served by the mainstream media (newspapers, TV) will find those more specific news groups devoted to meeting informational needs of the sort she has? How efficient and reliable are the reporters downstream in the chain of communication those who observe either a first-hand or a second-hand report and go on to report to others what they've heard? Are there any institutions or social practices that can serve as correctives, either when a false report is transmitted along a chain of communication, or else when a true report has its contents distorted in transmission? How effective are these corrective institutions or social practices? One can imagine assessing both the relative information-saturation of a community, and also the reliability of any arbitrary communication one comes across in that community.

These remarks are abstract in the extreme, and any successful assessment of the epistemology of coverage will have to descend from the heights of such abstraction to take a look at the details of particular cases. But as we do it is worth keeping the big picture in mind. Take any information-rich community, where there are groups dedicated to the investigation and reporting of facts in some domain. In any such community, it can come to pass that community members begin to rely on one or more of these groups for what they know in that domain, in the sense that such groups are the main, and perhaps the only, source that the individual has regarding information in the domain in question. It can also come to pass that individuals will begin to rely on such groups for adequate (reliable) coverage. Such a moment presents both an epistemic opportunity and a danger.

The opportunity lies in having some epistemically adequate way to separate the informational wheat from the chaff in our information-rich society. Take cases where the

information we receive involves something we cannot independently confirm (at least not without great cost and effort). In such cases it may well be an epistemically-wise policy, not just to accept *everything* the sanctioned sources tell us (unless there is/are positive reason(s) not to), but also to accept *only* information from such sources (unless there is/are positive reason(s) to accept information from a not-yet-recognized source). However, no sooner is such a policy formulated, than we see the very real possibility for abuse by the powerful interests at play in the dissemination of news which brings us to the danger.

The danger is two-fold. For one thing, the coverage-relying individual runs the risk of committing the fallacy of ignorance: believing not- p merely because she has not come across evidence (testimony) that p . For another, an individual who exhibits coverage-reliance towards a particular source can become blind to new or unrecognized but still reliable sources of information, when these offer reports going beyond what the relied-upon sources have said about some matter. In this respect the individual is both persisting in retaining false belief in the face of what is in fact a reliable indication to the contrary, and also losing out on an opportunity to acquire reliable information.

I see the epistemology of coverage as assessing how well individuals and communities manage these risks as they aim to reap the benefits of their information-saturated environment.

6. Conclusion

In standard discussions of the epistemology of testimony, it has been noted that

hearers aim to acquire testimony-based beliefs from testimony that satisfies the testimony-to-truth conditional,

(2) If S testified that p , then p .

But here I have been suggesting that something close to the converse is also a desirable feature of the institution of testimony. In particular, we want to be in an epistemic community where the newsworthy propositions satisfy the coverage (or truth-to-testimony) conditional,

(5) If p , then one among the trusted sources Σ testified that p .

While it has been clear to everyone working in epistemology that there is something epistemologically beneficial with being in a community whose members regularly satisfy (2), it has been less remarked though it should be equally clear that there is something epistemologically beneficial with being in a community that satisfies (5) as well. It is not for nothing that *The New York Times* presents itself as publishing *All* of the news that fits to Print.

Work Cited

Audi, R. 1982: Believing and Affirming. *Mind* 91, 115-20.

Burge, T. 1993: Content Preservation. *Philosophical Review* 102:4, 457-488.

Coady, C. 1992: Testimony: A Philosophical Study. (Oxford: Oxford University Press).

Fricker, E. 1987: The Epistemology of Testimony. *Proceedings of the Aristotelian Society*, Supplemental Vol. 61, 57-83.

Fricker, E. 1994: Against Gullibility. In B.K. Matilal and A. Chakrabarti, eds., Knowing from Words (Amsterdam: Kluwer Academic Publishers), 125-61.

Goldberg, S. 2005: Testimonial knowledge from unsafe testimony. *Analysis* 65:4, 302-11 .

Goldberg, S. 2006: Reductionism and the Distinctiveness of Testimonial Knowledge, in Lackey, J. and Sosa, E., eds. The Epistemology of Testimony (Oxford: Oxford University Press, 2006), pp. 127-44.

Goldberg, S. 2007: Anti-Individualism: Mind and Language, Knowledge and Justification. (Cambridge: Cambridge University Press.)

Goldberg, S. Forthcoming. Testimony as Evidence. *Philosophica*.

Goldberg, S. and Henderson, D. 2007: Monitoring and Anti-Reductionism in the Epistemology of Testimony. *Philosophy and Phenomenological Research* 72:3, 600-17.

Graham, P. 2000: Transferring Knowledge. *Nous* 34:1, 131-152.

Reid, T. 1872/1993: The Complete Works of Thomas Reid. Sir William Hamilton (ed.). (Edinburgh: Maclachlan and Stewart.)

Ross, A. 1986: Why do we believe what we are told? *Ratio* 28:1, 69-88.