

Group Knowledge Attributions

Jennifer Lackey
Northwestern University

A view growing in popularity in the current philosophical literature is that the purpose of knowledge attributions is to identify or flag good informants. Such a thesis has its origin in the work of Bernard Williams and Edward Craig. Williams, for instance, claims that the central point of the concept of knowledge is “to find somebody who is a source of reliable information about something” (1973, p. 146). And Craig maintains that knowledge ascriptions are made so as to “flag approved sources of information” (Craig 1990, p. 11). Let us call this the *Reliable Informant View of Knowledge Attributions* (hereafter, the RIVKA) and characterize it as follows:

RIVKA: The purpose of knowledge attributions is to identify or flag reliable informants. Recent writers not only endorse the RIVKA, but also claim to derive significant epistemological payoff from doing so. For instance, Duncan Pritchard argues that “...the central importance of the concept of knowledge resides in the practical need to pick-out reliable informants—informants that one can rely on” (Pritchard unpublished, p. 57 of ms.). Such a thesis, he claims, provides support for his preferred anti-luck virtue epistemology, which incorporates anti-luck and ability conditions and “accords each condition equal weight in the sense that they are each answering to a fundamental intuition about knowledge” (Pritchard unpublished, p. 51). Similarly, John Greco argues that “...knowledge attributions serve to flag good information and good sources of information for use in practical reasoning” (Greco 2007, p. 68). Endorsing this thesis, he argues, enables him to defend his “knowledge as creditable true belief” view, according to which:

...knowledge attributions can be understood as credit attributions: when we say that someone knows something, we credit them for getting it right. When we deny that someone knows something, we deny them credit for getting things right. In one sort of case, we deny

credit for success because there was no success—S’s belief is false. In other cases we deny credit because success was realized, but not through ability—S believes the truth but it was a lucky guess, or there was faulty reasoning. (Greco 2007, p. 57)¹

Likewise, Ram Neta writes:

Craig (1990) articulates and defends a plausible account of our practice of attributing knowledge, and I propose that we extend his theory to epistemological status generally.... To state Craig’s hypothesis...precisely...: the concept of knowledge is designed to flag sources of information that acquire their information in a way that we can tell is a reliable way of acquiring information, and that can take the credit for acquiring their information in this way. Let’s sum this up by saying that the concept of knowledge is designed to flag ‘creditable informants’.

I propose to generalize Craig’s hypothesis by claiming that the various terms of epistemological appraisal are designed to flag informants that are creditable to various levels, or in various ways. (Neta 2006, pp. 266-7)²

According to Neta, such a generalization of Craig’s view can explain (1) why the practice of assertion is governed by epistemological norms, (2) what the proper targets of epistemological assessment are, (3) what makes epistemological norms binding on epistemic agents, and (4) why the dispute between internalists and externalists in epistemology is so hardy and intractable. (Neta 2006, pp. 268-70)

Though there are some obvious differences among what these theorists endorse, they all agree both in their acceptance of the RIVKA—the purpose of knowledge attributions is to identify or flag reliable³ informants—and in their conclusion that such a view has enormous epistemological significance.⁴

One feature of our knowledge ascriptions⁵ that is of interest here is that knowledge is attributed, not only to individuals, but to groups as well. Just as I may say that my daughter knows

that my birthday is in September, so, too, I may claim that Borders knows that it is losing the battle over the e-book market with Barnes & Noble or that the Supreme Court knows the repercussions of its recent decision. This phenomenon has frequently been noted in the literature on group knowledge and belief. For instance, Alvin Goldman writes: “In common parlance, certainly, organizations are treated as subjects for knowledge attribution. In the wake of 9/11, there has been much commentary on what the C.I.A. and the F.B.I. did or didn’t know about terrorist plans before the event itself” (Goldman 2004, p. 12). Similarly, in arguing for the existence of group belief, Frederick F. Schmitt asks us to “[b]egin with the observation that we routinely ascribe knowledge ...to groups. We say things such as, ‘The Engineering Division of the Ford Motor Corporation knew that the Pinto was explosive,’ [etc.].... These are, on the face of it, ascriptions of cognitive states to groups” (Schmitt 1994, pp. 257-8).⁶ A common theme running through these passages is that we do in fact attribute knowledge to groups and that we do so systematically and in a widespread fashion. This observation—that knowledge is attributed to groups as well as to individuals—can be combined with the RIVKA, resulting in the following group version of the view:

G-RIVKA: The purpose of group knowledge attributions is to identify or flag reliable informants.

Since ascribing knowledge to groups is no less an attribution of knowledge than it is in the individual case, there is no reason for the proponent of the RIVKA to object to the formulation of this thesis in terms of groups. Indeed, were the G-RIVKA denied without substantial argumentation, the result would be that we can simply pick and choose which knowledge attributions to take seriously. Whenever a counterexample involving a knowledge attribution is raised to one’s view, one could simply respond, “well, that’s just one of those knowledge attributions that we aren’t taking seriously.” This arbitrary reliance on the linguistic data would seriously undermine the substantive

philosophical conclusions purportedly derived from them. Given the obvious unattractiveness of this result, combined with the fact that there are no arguments in the literature against the propriety of ascribing knowledge to groups, there is no reason to accept the RIVKA without also granting the G-RIVKA.

Further evidence for this conclusion comes from noticing that group knowledge attributions are made systematically and in a widespread fashion *in contexts where there is a heightened concern for speaking precisely*. Consider, for example, the following ascription taken from a recent article in the *Los Angeles Times*: “BP officials knew about a problem on a crucial well safety device at least three months before the catastrophic April 20 explosion in the Gulf of Mexico but failed to repair it” (*Los Angeles Times*, July 21, 2010). This attribution of knowledge to BP is regarded as proper not only when made in the *Los Angeles Times*, but also if it were later offered in a courtroom setting where critical issues of moral and legal responsibility are being debated. No one reading the newspaper or sitting in the courtroom would regard the attribution as merely metaphorical or as an exaggeration, as is the case when one says, “The injured mouse just sat there as I approached him—it was as if he knew that I was going to help him,” or “The wind blew the stick precisely to mark our location—it seemed to know that we needed to be rescued.” In these latter cases, there wouldn’t be a debate over whether the mouse or the wind really knew the propositions being attributed to them; instead, one would simply retract the knowledge attributions if challenged to be precise or careful. But clearly the situation is quite different in the case of claiming that BP knew about the faulty safety device, where such ascriptions would hold even in contexts where journalistic and legal standards call for precision and care. In what follows, then, I will assume that there are some group knowledge attributions that are proper, and I will focus on paradigmatic cases where knowledge is widely and systematically attributed to groups across a wide range of contexts as instances of them.

With this in mind, I will argue in this paper that the G-RIVKA, and therefore the RIVKA, is false. In particular, I will consider three different kinds of paradigmatic group knowledge attributions and show that on any reading of “reliable informant,” such attributions systematically fail to identify or flag reliable informants. There are two natural responses that may be offered to my arguments. First, it may be claimed that an inflationary account of group attitudes can avoid the problems raised against the G-RIVKA and, second, it may be held that the RIVKA should be restricted only to individuals. Along the way, I will raise considerations that block each of these moves. Specifically, I will show that there are reasons to reject an inflationary account of group attitudes and also that there are problematic cases for the RIVKA involving only individual knowledge attributions. Finally, I will propose alternatives to both the RIVKA and the G-RIVKA, which I call the *Reliable Source of Information View of Knowledge Attributions*—the RSIVKA—and its group counterpart—the G-RSIVKA—according to which a central purpose of knowledge attributions, individual or group, is to identify or flag reliable *sources of information*. These replacement theses provide conditions that are neither strictly necessary nor sufficient for proper knowledge attributions, but they purport to capture what is broadly defensible in spirit about the RIVKA and the G-RIVKA.

1. Reliable Informants as Reliable Believers

To begin, it will be helpful to get a clearer sense of what is meant by a reliable informant. Since Craig provides the most explicit and detailed account of this concept, his work is a natural place to start. The first condition that he offers, which is necessary, though not sufficient, for being a reliable informant, is characterized as follows:

BELIEF: A is a reliable informant with respect to the question whether p only if either p and A believes that p , or not- p and A believes that not- p . (Craig 1990, p. 12)

According to Craig, a person seeking a good informant wants someone “who will tell him the truth on [a given question]. The informant, we may assume, will not in general tell him the truth unless he (the informant) holds a true belief about it” (Craig 1990, p. 12). Given this, combined with the RIVKA, it follows that the purpose of knowledge attributions is to identify or flag informants who at least either believe that p when p is the case, or believe that not- p when not- p is the case. For ease of expression, let us say that according to BELIEF, reliable informants are understood as *reliable believers*. Thus, the G-RIVKA plus BELIEF yields the view that the purpose of group knowledge attributions is to identify or flag informants who are at least reliable believers.

Since we cannot determine whether a group is a reliable believer without knowing what group belief is, the next question to which we should turn is how to understand this concept. A standard view of group belief is *summative*, according to which a group’s believing that p can be understood in the minimal sense that all or some of the members of the group believe that p . Anthony Quinton (1975/1976) is responsible for coining this term and he defends a summative view of group attitudes and actions in general in the following passage:

We do, of course, speak freely of the mental properties and acts of a group in the way we do of individual people. Groups are said to have beliefs, emotions, and attitudes and to take decisions and make promises. But these ways of speaking are plainly metaphorical. To ascribe mental predicates to a group is always an indirect way of ascribing such predicates to its members. With such mental states as beliefs and attitudes, the ascriptions are of what I have called a summative kind. To say that the industrial working class is determined to resist anti-trade union laws is to say that all or most industrial workers are so minded. (Quinton 1975/1976, p. 17)⁷

Summative accounts of group belief also result from widely used aggregation procedures for combining individual beliefs held by the members of a group into collective ones. “Aggregation

procedures are mechanisms a multimember group can use to combine ('aggregate') the individual beliefs or judgments held by the group members into collective beliefs or judgments endorsed by the group as a whole" (List 2005, p. 25).⁸ For instance, a dictatorial procedure, "whereby the collective judgments are always those of some antecedently fixed group member (the 'dictator')" (List 2005, p. 28) reduces the belief of the group to the judgment of a single member—the dictator. A majority procedure, "whereby a group judges a given proposition to be true whenever a majority of group members judges it to be true," reduces the belief of the group to a majority of its individual members (List 2005, p. 27). A supermajority procedure, whereby a group judges a given proposition to be true whenever a supermajority of group members judges it to be true, reduces the belief of the group to a supermajority of its individual members. And a unanimity procedure, "whereby the group makes a judgment on a proposition if and only if the group members unanimously endorse that judgment," (List 2005, p. 30) reduces the belief of the group to all of its members when there is unanimous agreement. Though there are obvious differences between the reductive bases on such views, they are all summative: the beliefs of the group are reducible to the beliefs of all or some of the members that comprise it

Summative accounts of group belief are not only widely endorsed theoretically, they are also most in keeping with ordinary intuitions regarding such a notion. Margaret Gilbert makes this point when she writes:

What is it for *us* to believe that such-and-such, according to our everyday understanding? It is common to answer this question with some form of 'summative' account. For *us* to believe that *p* is for all or most of us to believe that *p*. Or perhaps a 'common knowledge' condition may be added: for *us* to believe that *p* is for all or most of us to believe that *p*, while this is common knowledge among us. Whatever the precise account given, the core of it is a number of individuals who personally believe that *p*. (Gilbert 1994, p. 235)⁹

Given the widespread support for a summative view of group belief, then, let us proceed for now with this conception of group belief in our evaluation of the G-RIVKA.

When reliable informants are characterized in terms of reliable believers and group belief is understood summatively, then the weakest possible version of the G-RIVKA holds that there must be at least one reliable believer in a given group when knowledge is attributed to it. But here is the problem: there are paradigmatic group knowledge attributions where not a single member of the group has the belief in question. Hence such group knowledge attributions are obviously not flagging reliable believers because, on a summative view, they are not flagging believers at all.

There are two different kinds of cases where there is a paradigmatic group knowledge attribution where not a single member of the group holds the relevant belief. The first arises when group knowledge is determined by *aggregating the judgments* of the individual members. Consider the following case due to Phillip Pettit:

JUDGMENT AGGREGATION: The employees of a company are deciding whether to forgo a pay-raise in order to spend the saved money on implementing a set of workplace safety measures. The employees are supposed to make their decision on the basis of considering three separable issues: “first, how serious the danger is; second, how effective the safety measures that a pay-sacrifice would buy is likely to be; and third, whether the pay-sacrifice is bearable for members individually. If an employee thinks that the danger is sufficiently serious, the safety measure sufficiently effective, and the pay-sacrifice sufficiently bearable, he or she will vote for the sacrifice; otherwise he or she will vote against” (Pettit 2003, p. 171). The company’s three employees vote in the following way:

	Serious danger?	Effective measure?	Bearable loss?	Pay sacrifice?
A.	Yes	No	Yes	No
B.	No	Yes	Yes	No

that the benefits of vaccines strongly outweigh the possibility that a very small percentage of children will become autistic as a result. Despite this, all of the doctors who are members of the American Academy of Pediatrics recognize that the evidence is inconclusive and so have some doubts that prevent them from believing that there is not a causal connection between the MMR and autism.

As we did in JUDGMENT AGGREGATION, we can assume that it is in fact true that there is no causal connection between the MMR and autism, that the “substantial, though not conclusive” evidence upon which the doctors’ conclusion is based is sufficient for epistemic justification, and that there are no relevant Gettier considerations. Thus, there is no reason to deny the attribution of knowledge to the American Academy of Pediatrics. Indeed, when the official position of the AAP is made public and millions of Americans have their children receive the MMR vaccine, surely knowledge attributions regarding the lack of a causal connection between the vaccine and autism will abound. Once again, then, there will be the attribution of knowledge that p to a group when not a single member of the group reliably believes that p , and thus such an attribution cannot be flagging a reliable group believer in a summative sense.¹³

The upshot of these considerations is that if the purpose of knowledge attributions is to flag reliable informants, where this latter concept is understood in terms of summatively characterized reliable group believers, then group knowledge attributions pose an obstacle to this view. For if knowledge that p is attributed to a group where not a single member of this group believes that p , and group belief is understood summatively, then this attribution cannot be flagging an informant who at least either believes that p when p is the case, or believes that not- p when not- p is the case.

The obvious move for the proponent of the G-RIVKA to make at this point, then, is to argue on behalf of a *non-summative, inflationary* account of group belief, according to which group belief is irreducible to the beliefs of all or some of its members. In particular, such a view holds that

in some very important sense, the group itself is the subject of the belief, where this is understood as over and above, or otherwise distinct from, the beliefs of any of its individual members. Such an inflationary account of group belief not only has the advantage of supporting the interpretation of G-RIVKA in terms of BELIEF, it also explains our knowledge attributions to groups in the cases described earlier.

The most widely accepted non-summative view of group belief is what we may call the *joint acceptance* account, a prominent expression of which is offered by Gilbert in the following passage:

A group *G* *believes* that *p* if and only if the members of *G* jointly accept that *p*.

The members of *G* *jointly accept* that *p* if and only if it is common knowledge in *G* that the members of *G* individually have intentionally and openly...expressed their willingness jointly to accept that *p* with the other members of *G*. (Gilbert 1989, p. 306)

A key aspect of such an account is that joint acceptance does not require belief on the part of a single member of the group in question. She writes:

It should be understood that: (1) Joint acceptance of a proposition *p* by a group whose members are *X*, *Y*, and *Z*, does not entail that there is some subset of the set comprising *X*, *Y*, and *Z* such that all the members of that subset individually believe that *p*. (2) One who participates in joint acceptance of *p* thereby accepts an obligation to do what he can to bring it about that any joint endeavors...among the members of *G* be conducted on the assumption that *p* is true. He is entitled to expect others' support in bringing this about. (3) One does not have to accept an obligation to believe or to try to believe that *p*. However, (4) if one does believe something that is inconsistent with *p*, one is required at least not to express that belief baldly. (Gilbert 1989, pp. 306-7)

Thus, according to Gilbert's non-summative view, so long as a group jointly accepts that *p*, such a group can be said to believe that *p* even if not a single one of its members holds this belief.¹⁴ This

account is perhaps best suited for explaining beliefs attributed to organized or established groups, such as that found in OFFICIAL POSITION. None of the doctors in the AAP believes that there is not a causal connection between the MMR and autism, but the members of the Board of Directors officially and jointly accept this conclusion on behalf of the Academy.

However, while Gilbert's account of group belief is best supported by OFFICIAL POSITION, it is not entirely clear that her view can fully accommodate even this sort of case. For though the members of the Board jointly accept the conclusion that there is not a causal connection between the MMR and autism, it is not the case that each member of the AAP engages in such joint acceptance, despite the fact that the knowledge attribution is to the Academy as a whole. This is made clear by the fact that it is quite easy to imagine many individual members of the AAP baldly expressing their dissenting beliefs to their own patients, which is in violation of (4) above. This is not an uncommon phenomenon: the Supreme Court can issue a decision as a group while individual members nonetheless issue a dissenting report; a department can officially recommend as a group that one of its members be granted tenure even though individual members openly reject this conclusion.

Moreover, the situation is far worse when JUDGMENT AGGREGATION is considered, for here there is nothing even in the neighborhood of joint acceptance by the members of the group in question. Indeed, while the premise-based aggregation procedure results in the group's judging that they should accept the pay sacrifice, the individual members unanimously and explicitly reject this conclusion. Thus, even on a widely accepted non-summative account of group belief, OFFICIAL POSITION and JUDGMENT AGGREGATION fail to deliver the verdict that reliable group believers are being identified or flagged by the knowledge attributions in question.

These objections are specific to the joint acceptance account, which is by far the most widely accepted non-summative view of group belief. But recall that Pettit uses the premise-based

aggregation procedure to motivate embracing such an inflationary view in JUDGMENT AGGREGATION. So, it should be noted that there is also a serious problem with grounding a non-summative view of group belief in this sort of argument. To see this, notice that there are many different aggregation procedures that may be used to generate a group belief or judgment in any given case, some of which were discussed in the previous section. What is of interest for our purposes here, however, is that there is not a fact of the matter about which aggregation procedure should be used in any given case—each has its virtues and its vices. Fabrizio Cariani begins his overview of the literature on judgment aggregation with a similar point:

Judgment Aggregation studies how collective judgments arise from the aggregation of individual opinions. Its motivating observation is that *prima facie* plausible rules for aggregating judgments do not (and cannot) have all the features we take to be desirable. Judgment Aggregation, then, aims to classify the various aggregation rules by means of the properties they do satisfy and to select those that are, in some sense, best. (Cariani 2011, p. 1)¹⁵

Given that competing aggregation procedures can result in radically different group beliefs despite holding fixed the individuals and their respective beliefs, combined with the point that there is not always a fact of the matter about which procedure to use, there also does not seem to be a fact of the matter about what a given group's belief is in such cases. The group can be said to believe that *p* just as easily as it can be said to believe that not-*p*. Surely this conclusion provides a compelling reason to be doubtful of an inflationary account of group belief that is motivated via a premise-based aggregation procedure.

There are also concerns that can be raised with inflationary accounts of group belief in general. In particular, paradigmatic instances of group "belief" function differently in important ways than beliefs in individuals do. To see this, consider again OFFICIAL POSITION, which

involves the American Academy of Pediatrics denying that there is a connection between the MMR and autism. When group belief is determined by the official position arrived at by a decision-making power, group belief is far more directly voluntary than it is in the individual case.¹⁶ The power in question can simply *decide* that the official position of the group is that *p*, whereas individuals do not seem capable of just deciding to believe that *p* in this way.¹⁷ Similarly, group belief in this sense is far less governed by evidence than it is when an individual's doxastic states are at issue. For instance, K. Brad Wray argues that groups, unlike individual agents, can choose to believe based on their goals and Christopher McMahon claims that groups often defend as true positions that they adopt for purely instrumental reasons.¹⁸ In OFFICIAL POSITION, the goal of the AAP may be to produce the best health for the greatest number of children, and so as a result the Board of Directors may choose to downplay the possibility of some children becoming autistic as a result of the MMR since the benefits of vaccines are more likely to further this broader goal. Or the AAP may adopt the position that there is no causal link between the MMR and autism in an effort to promote their positive image in the public eye.¹⁹ Neither of these ways of "belief" formation seems available in the individual case, where doxastic attitudes are far more directly sensitive to evidence.²⁰ These considerations provide further reason to reject appealing to an inflationary account of group belief in an effort to defend the interpretation of the G-RIVKA in terms of BELIEF.

Thus, the G-RIVKA, when understood in terms of BELIEF, is supported by neither a summative nor a non-summative account of group belief.

A proponent of the G-RIVKA may respond to this conclusion by claiming that the RIVKA should be restricted to knowledge attributions involving only individuals. In particular, it may be argued that group knowledge attributions are metaphorical, "loose talk," or are otherwise not fully legitimate and hence the G-RIVKA should be rejected from the start.

By way of response to this move, the first point to notice is that in the absence of substantial argumentation on its behalf, it is utterly arbitrary. For as was noted at the outset of this paper, the fact is that we do attribute knowledge to groups and that we do so in a widespread and systematic way. It is entirely unclear, then, why we can pick and choose which knowledge attributions should be subsumed by our general account of knowledge attributions.

Moreover, there are independent reasons—ones that do not depend on considerations about groups at all—for doubting that belief in the relevant proposition is a necessary condition for being a reliable informant with respect to the question whether *p*. In particular, there are cases in which a speaker fails to believe a proposition to which she is testifying, but nevertheless reliably conveys the information in question through her testimony. For instance, suppose that Stella, an elementary school teacher who is also a creationist, regards it as her pedagogical duty to present material to her students that is best supported by the available evidence, which clearly includes the truth of evolutionary theory. As a result, after consulting reliable sources in the library and developing reliable lecture notes, she asserts to her students, “Modern day *Homo sapiens* evolved from *Homo erectus*,” while presenting her biology lesson today. Though Stella herself neither believes nor knows this proposition, she never shares her own views with her students, and so they form the corresponding true belief solely on the basis of her reliable testimony. Now, given Stella’s commitment to creationism, she does not believe that modern day *Homo sapiens* evolved from *Homo erectus*. But because she reliably conveys this information to her fourth-grade students, and does not provide them with any evidence to the contrary, they acquire the knowledge in question on the basis of her testimony. What this case reveals, then, is that an *unreliable believer* may nonetheless be a *reliable testifier*, and so may reliably convey information despite the fact that she fails to believe it herself. For although Stella ignores the relevant evidence with respect to her doxastic states concerning evolutionary theory—thereby leading her to lack the belief that modern day *Homo sapiens* evolved

from *Homo erectus*—she bases her testimony regarding this topic firmly on such evidence.²¹ This sort of case provides a compelling reason entirely separate from those involving groups for rejecting the view in which a reliable informant is understood in terms of BELIEF, thereby blocking the possibility of simply restricting the RIVKA to knowledge attributions involving individuals.

We have seen that when the notion of reliable informant operative in the G-RIVKA requires that the testifier believe that to which she is testifying, problems arise regarding knowledge attributions involving groups. If the G-RIVKA is to be defensible, then, we need to look beyond BELIEF for a proper characterization of the notion of a reliable informant.

2. Reliable Informants as Reliable Testifiers

In looking for a replacement for BELIEF to underwrite the G-RIVKA, an alternative suggestion presented itself in the previous section. Instead of focusing on reliable believers, as BELIEF does, the focus should be on *reliable testifiers*. This distinction is made clear in the above case, where Stella is an unreliable believer when it comes to evolutionary theory but a reliable testifier when it comes to conveying this sort of information to her students. More precisely, BELIEF should be replaced with:

TESTIMONY: A is a reliable informant with respect to the question whether p only if A is a reliable testifier with respect to whether p .

Being a reliable testifier can, in turn, be fleshed out in any number of ways. For instance, one way of understanding this reliability constraint is via Nozick's notion of *sensitivity*, i.e., A would not state that p if p were false.²² Alternatively, reliability can be fleshed out in terms of the *safety* requirement endorsed by Pritchard, Sosa, and Williamson, i.e., A would not state that p without it being so that p .²³ Or it may be the case that general facts about A's history as a testimony-producing source or as a virtuous testifier determine whether A's statement is reliable.²⁴ But regardless of how the details of

reliability are worked out, the central point that is of import here is that belief that p is not a necessary condition for being a reliable testifier and, thereby, a reliable informant regarding that p . Thus, when being a reliable informant is understood in terms of TESTIMONY, the G-RIVKA amounts to the following version of this view: the purpose of group knowledge attributions is to identify or flag reliable testifiers.

As was the case with group belief, the standard view of group testimony is summative, according to which a group's testifying that p can be understood in the minimal sense that all or some members of the group would testify that p were the relevant opportunity to arise. Deborah Tollefsen characterizes the view as follows: "...when a group offers testimony it is really understood as the testimony of all or some of the members—what they would testify to if given the opportunity. Another way to understand this summative approach would be to say that a group's testimony p expresses the views of all or some of the members of the group" (2007, p. 300).²⁵ Since this is the classic view of group testimony, let us evaluate the G-RIVKA with a summative conception of group testimony in mind.

Now, if reliable informants are understood in terms of reliable testifiers and group testimony is understood summatively, then the weakest possible version of the G-RIVKA holds that there must be at least one reliable testifier in a given group when knowledge is attributed to it. But here is the problem: knowledge can plausibly be attributed to groups even when not a single member of the group is reliable with respect to the testimony in question. Hence such group knowledge attributions cannot be flagging reliable testifiers.

To see this, consider the following paradigmatic instance of group testimony, though one that represents a kind different than that found in either JUDGMENT AGGREGATION or OFFICIAL POSITION:

DISTRIBUTED INFORMATION: The UN Population Commission, which is comprised of forty-seven individual members, issues a report entitled *Charting the Progress of Populations*.²⁶ Each member of the group was responsible for collecting information about a different segment of the population represented in the document, and their respective work was done entirely independently from one another. The information contained in the report is, then, widely distributed across the members of the group. Sam, who is not a member of the UN Population Commission, was hired to interpret and compile all of the data contributed by the members of this group into the published report and to serve as the group's spokesperson. One of the statements in this report is, "the birth rate of Latinos in the US is on the rise," of which not a single member of the UN Population Commission is aware.

We can assume that the statement in question is true, that the testimony is justified, and that there are no relevant Gettier considerations, thereby leaving the door open to the attribution of knowledge that the birth rate of Latinos is on the rise to the UN Population Commission.

According to the scenario envisaged above, however, not a single member of the group is privy to this information because the data upon which this conclusion is based was widely distributed among those in the group. In particular, Sam, who is not a member of the group, interprets and compiles the data contributed by the members of this group into the published report and serves as the group's spokesperson. When knowledge is attributed to the UN Population Commission, then, the G-RIVKA holds that its purpose is to identify or flag a reliable informant. But here not only is there the absence of the relevant belief on the part of the group's members, as we saw in the cases from the previous section, there is not even a reliable testifier in the group regarding the information in question. In other words, the attribution of knowledge is over *here*—where the group is—and the reliable testifier is over *there*—where the spokesperson is. Thus, DISTRIBUTED INFORMATION shows that when being a reliable testifier is understood summatively, the G-RIVKA is false.

It is worth mentioning that DISTRIBUTED INFORMATION also raises problems for a slightly more specific version of the RIVKA mentioned earlier. Recall that John Greco holds that “knowledge attributions can be understood as credit attributions: when we say that someone knows something, we credit them for getting it right. When we deny that someone knows something, we deny them credit for getting things right” (Greco 2007, p. 57). Similarly, Ram Neta claims that “the concept of knowledge is designed to flag ‘creditable informants’” (Neta 2006, p. 267). But when knowledge that the birth rate of Latinos is on the rise is attributed to the UN Population Commission in DISTRIBUTED INFORMATION, does the group itself deserve credit for getting this right? It is hard to see how this can be true, at least in any straightforward way. No one in the group believes this proposition, no one in the group is privy to all of the information grounding this result, and no one in the group is even aware of this fact itself. It is Sam who compiled and interpreted all of the relevant data and arrived at the conclusion that the birth rate of Latinos is on the rise. Thus, if anyone deserves credit for getting this right, it is Sam, not the subject of the attribution of knowledge in question. Of course, the members of the group are clearly relevant to Sam’s getting it right. But this is no different than the common phenomenon of an individual drawing a new conclusion or arriving at a new bit of knowledge through reliance on the testimony of others. I may rely on the data of countless historians to arrive at the entirely new view that, say, the unresolved tension and resentments resulting from WWI did not in fact contribute to the start of WWII. In such a case, we would surely still credit me with getting this conclusion right—assuming, for the sake of argument, that it is—even if my knowledge is significantly indebted to the testimony of many others. The situation with Sam is no different. Sam could not have arrived at the conclusion that the birth rate of the Latinos is on the rise without the data contributed by the members of the UN Population Commission. But it does not follow from this that the UN Population Commission, and not Sam, deserves credit for getting this right. Thus, DISTRIBUTED

INFORMATION shows that knowledge attributions are not always credit attributions since we plausibly attribute knowledge to the UN Population Commission in this case despite the fact that they do not deserve credit for the conclusion that the birth rate of Latinos is on the rise.

It may be objected, however, that Sam is not in fact the testifier of the information in DISTRIBUTED INFORMATION—the UN Population Commission is. In particular, it may be argued that a spokesperson such as Sam functions as a parrot of the thoughts of the group and, since parrots do not offer testimony, neither does the spokesperson. If the spokesperson is not testifying, then we need to look elsewhere for the testimonial source. And, of course, the group that the spokesperson is representing is the most natural candidate for occupying this role. Thus, we have reason to deny that DISTRIBUTED INFORMATION is a counterexample to the G-RIVKA when understood in terms of TESTIMONY.

By way of response to this objection, the first point to notice is that Sam's statement that the birth rate of Latinos in the US is on the rise satisfies every existing theory of the nature of testimony offered in the literature. For instance, Elizabeth Fricker characterizes testimony as "tellings generally" with "no restrictions either on subject matter, or on the speaker's epistemic relation to it" (Fricker 1995, pp. 396-7). Robert Audi claims that we must understand testimony as "...people's telling us things" (Audi 1997, p. 406). Catherine Elgin characterizes testimony as "...utterances and inscriptions that purport to convey information and transmit warrant for the information they convey" (Elgin 2002, p. 292). C.A.J. Coady argues that S testifies by making some statement that *p* if and only if: (i) S's stating that *p* is evidence that *p* and is offered as evidence that *p*; (ii) S has the relevant competence, authority, or credentials to state truly that *p*; and (iii) S's statement that *p* is relevant to some disputed or unresolved question (which may or may not be whether *p*) and is directed to those who are in need of evidence on the matter. (Coady 1992, p. 42) And according to my view of testimony, "S testifies that *p* by making an act of communication *a* if and only if (in part)

in virtue of *a*'s communicable content, (1) *S* reasonably intends to convey the information that *p* or²⁷ (2) *a* is reasonably taken as conveying the information that *p*" (Lackey 2008, p. 30).²⁸ Now, obviously, Sam's report is plausibly regarded as a "telling," and so it fulfills the broad characterizations of testimony offered by Fricker and Audi. Sam's statement is also an utterance that purports to convey information and transmit warrant for the information that it conveys, thereby qualifying as testimony on Elgin's view. Moreover, she is a competent speaker who offers a statement that is evidence in an objective sense, and we can certainly imagine that it is relevant to some disputed or unresolved question and is directed to those who are in need of evidence on the matter. Thus, there is no obstacle to Sam's report qualifying as testimony on Coady's view as well. And finally, Sam reasonably intends to convey information in virtue of the communicable content of her assertion, and hence her statement is testimony on my view, too. According to every view of testimony offered in the literature, then, Sam satisfies the requisite conditions for testifying.

Moreover, Sam is unlike a parrot in various crucial respects. When Sam states that the birth rate of Latinos in the US is on the rise, she is offering an act of communication, which requires that she have the *intention to express communicable content*. To clarify this notion, consider a case in which, unbeknownst to me, Chloe has headphones on and is bopping her head to the beat of the music. I walk into the room, ask her if there is any cake left, and, seeing her bop her head, think that she has intended to communicate to me that there is cake left in the kitchen. This sort of case may be an example of *ostensible testimony*, but it should not qualify as genuine testimony in any sense of the word. The reason for this is that Chloe's head-bopping is not an act of communication *since she did not intend to express communicable content*. This is true of a parrot's offering a string of sounds as well. When a parrot merely repeats what it hears, it does not have the intention to express communicable content. In this way, the parrot *mimics* an act of communication, but does not *make* one. So, there is no reason to compare Sam's statement about the birth rate of Latinos in the US to the sounds

produced by a parrot. Hence, I see no compelling reason to deny that a spokesperson is offering testimony.

The obvious move for the proponent of the G-RIVKA to make at this point, then, is to argue on behalf of a *non-summative* account of group testimony, according to which group testimony is irreducible to the testimony of all or some of its member. While there are virtually no such accounts in the literature,²⁹ I have recently developed a view of group testimony that is non-summative, though it is not inflationary. According to my *deflationary account of group testimony*, the reliability of the testimony of a given group is explained using only the resources needed in the individual case.³⁰ On my view, then, the reliability of the testimony of the UN Population Commission in DISTRIBUTED INFORMATION simply reduces to that of Sam's testimony, and thus there is no testimony over and above the spokesperson's to appeal to here. In particular, whether the statement that the birth rate of Latinos is on the rise is reliable on my view will depend upon facts about Sam as a testifier of this information since, *qua* spokesperson, she is the one offering the testimony. As we saw above, this reliability constraint can be fleshed out in a number of ways, e.g., via Nozick's notion of *sensitivity*—Sam would not state that p if p were false—or the *safety* requirement endorsed by Pritchard, Sosa, and Williamson—Sam would not state that p without it being so that p , and so on. But since *Sam* is the one issuing the statement in question, whether hearers acquire knowledge on the basis of her testimony will depend on whether *she* is a reliable testifier of this information. Thus, on my view of group testimony, the epistemic status of the testimony of the UN Population Commission on the birth rate of Latinos in the US reduces to the epistemic status of Sam's statement about this fact. In this sense, then, my view of group testimony may be described as *reductionist*. It is, however, perhaps even better described as *deflationary* since there is no special epistemology of group testimony at all—it is simply subsumed by my general account of individual testimony.

What should be emphasized here, however, is that my deflationary account of group testimony is non-summative in nature. Recall that on a summative view, one's general epistemology of testimony can simply be applied to all or some of the individual members whose testimony constitutes that of the group. On my view, however, the epistemic status of the testimony of a group can reduce to the epistemic status of an individual *who is not a member of the group in question*. This is the case with Sam who, *ex hypothesi*, is not a member of the UN Population Commission but nonetheless issues a statement on its behalf about the birth rate of Latinos in the US. In such a case, my view holds that we treat Sam's report, which conveys the testimony of the group, just like any other instance of testimony; hence, whether a hearer can come to know this fact on the basis of her report depends largely on the reliability of her statement. Given this, my deflationary view is broadly reductionist, though non-summative: the epistemic status of a group's testimony is reducible to that of one or more individuals, though not necessarily ones who are members of the group in question. Specifically, the requirements in the case of individual testimony apply to the testifier of the group's statement, regardless of whether this spokesperson is a member of the group.³¹ Thus my deflationary view of group testimony provides a non-summative framework for understanding how Sam is the reliable testifier of the information that the birth rate of Latinos is on the rise, despite the fact that she is not a member of the group to which the knowledge in question is attributed. Of course, there need not be an official "spokesperson" in order to understand group testimony on my view: the knowledge of the relevant hearers is explained via the reliability of the process issuing the testimony in question, whether this is via a spokesperson, a written report, a collective briefing, and so on. I focus on DISTRIBUTED INFORMATION in fleshing out my account since this sort of case purports to provide one of the most difficult scenarios for accounts of group testimony to handle.

It should be noted, however, that once group testimony is understood in terms of my deflationary account, there is a sense in which DISTRIBUTED INFORMATION does not falsify the version of the G-RIVKA under consideration. For according to the current characterization of the G-RIVKA, the purpose of group knowledge attributions is to identify or flag reliable testifiers. But if the testimony of the UN Population Commission in DISTRIBUTED INFORMATION reduces to Sam's testimony, as my account of group testimony holds, then the UN Population Commission *is* a reliable testifier with respect to whether the birth rate of Latinos in the US is on the rise so long as Sam is. Given that, *ex hypothesi*, Sam is reliably testifying to this fact in the report that she issues, attributing knowledge to the UN Population Commission in DISTRIBUTED INFORMATION does identify or flag a reliable testifier. Strictly speaking, then, the G-RIVKA is not falsified.

While this is true, my deflationary view is the only reductionist, non-summative account of group testimony in the literature and hence it is the only one that can accommodate DISTRIBUTED INFORMATION without directly falsifying the G-RIVKA. On any other available view of group testimony, DISTRIBUTED INFORMATION is a straightforward counterexample. On a summative view, for instance, we saw that the testimony of a group reduces to that of some or all of its members. Since Sam is not a member of the UN Population Commission, the group's testimony cannot reduce to hers. Thus, when knowledge is attributed to the UN Population Commission, there is no relevant reliable testifier being flagged or identified when group testimony is understood in such a summative way. On an inflationary account, the testimony of a group is over and above, or otherwise distinct from, that of its members. But I know of no such inflationary account according to which the reliability of the testimony of the UN Population Commission is explained via Sam's reliability as a testifier. Instead, inflationary accounts take group testimony to issue from the group as a whole, not from a single non-member of it.

Hence, unless one accepts my deflationary account of group testimony, DISTRIBUTED INFORMATION does falsify the G-RIVKA.

Moreover, once one appreciates the force of DISTRIBUTED INFORMATION, a slightly modified version of this case poses a direct counterexample to the G-RIVKA, even when group testimony is understood in terms of my deflationary account. To see this, consider the following:

DISTRIBUTED INFORMATION2: The UN Population Commission, which is comprised of forty-seven individual members, has a central office located in New York City in which all of the data gathered by the Commission members are stored. Each member of the group was responsible for collecting information about a different segment of the population, and their respective work was done entirely independently from one another. Sam, who is not a member of the UN Population Commission, was hired to interpret and compile all of the data contributed by the members of this group into a single document, which she completed but did not yet publish. Maria, who is an investigative journalist working for an independently owned newspaper, suspects that the Commission is radically underreporting the number of Latinos currently living in the US for political purposes, and so she breaks into the New York office to research the matter. Once inside, Maria interprets and compiles all of the data contributed by the members of this group into a single document just as Sam did and publishes it in the newspaper. One of the statements in this report is, “the birth rate of Latinos in the US is on the rise,” of which not a single member of the UN Population Commission is aware.

As we did in the original scenario, we can assume that the statement in question is true, that the testimony is justified, and that there are no relevant Gettier considerations. In DISTRIBUTED INFORMATION2, then, everything that is epistemologically relevant is identical to the situation found in DISTRIBUTED INFORMATION. Given this, if it is plausible to attribute knowledge of

the fact that the birth rate of Latinos is on the rise to the UN Population Commission in the latter case, it should be no less plausible in the former case. However, though Maria is the source of the testimony that the birth rate of Latinos is on the rise, there is no reasonable sense in which Maria's testimony is, or can be, the reductive base for that of the UN Population Commission. She arrived at her conclusion in a way that is causally independent from Sam's process, she is not a member of the UN Population Commission, she is not their official spokesperson, and she is not testifying on their behalf. In DISTRIBUTED INFORMATION₂, then, we have a case in which knowledge is properly attributed to the UN Population Commission, despite the fact that there is no sense whatsoever on any account of group testimony in which a reliable testifier has been identified or flagged.

Thus, when the notion of reliable informant is understood in terms of being a reliable testifier, DISTRIBUTED INFORMATION and DISTRIBUTED INFORMATION₂ show that the G-RIVKA is false. Given this, combined with the considerations from the previous section, we may need to look in an altogether different place for an account of group knowledge attributions.

3. Reliable Informants versus Reliable Sources of Information

One lesson that may be learned from the previous section is that there is a difference between an *informant*, on the one hand, and a *source of information*, on the other. In DISTRIBUTED INFORMATION₂, Maria is clearly the informant when it comes to learning that the birth rate of Latinos is on the rise. But it does not follow from this that the UN Population is not a source of this information in a broader sense. If one wished to find out whether the birth rate of Latinos is on the rise, all of the relevant data needed to arrive at this conclusion is located within the group and its members. If one had the requisite skills to compile and interpret the relevant data, one could come

to know this fact without the help of Maria. In this way, the UN Population Commission is clearly a source of this information, even if it is not properly regarded as an informant with respect to it.

This distinction between informant and source of information is not lost on proponents of the G-RIVKA. Indeed, Craig himself calls attention to it, though he does so specifically to tie knowledge attributions to the former and not to the latter. He writes:

There are informants, and there are sources of information. Or, to arrange the terminology differently, among the various sources of information there are on the one hand informants who give information; and on the other there are states of affairs, some of which involve states of human beings and their behaviour, which have evidential value: information can be gleaned from them. Roughly, the distinction is that between a person's telling me something and my being able to tell something from observation of him.... In general terms it can be said that the concept of knowledge, as we operate it in everyday practice, is tied to informants rather than to sources of information in the sense just...characterized. (Craig 1990, p. 35)

According to Craig, then, informants are, roughly speaking, testifiers who offer reports or statements to others, while sources of information can be any states of affairs that have evidential value such that information can be gleaned from them. Now, as should be clear, I reject Craig's claim that our knowledge attributions identify or flag reliable informants rather than reliable sources of information. In *DISTRIBUTED INFORMATION2*, our attribution of knowledge that the birth rate of Latinos is on the rise to the UN Population Commission clearly fails to identify or flag a reliable informant since not a single member of the group is privy to the relevant information. Moreover, Maria, who *is* the reliable informant in this case, is not the subject of the relevant knowledge attribution.³² So, we have a knowledge attribution without a reliable informant, and a reliable informant without a knowledge attribution. Despite this, however, the attribution of

knowledge in DISTRIBUTED INFORMATION2 *does* identify or flag a reliable source of information in the UN Population Commission.

Given this, the problem posed by DISTRIBUTED INFORMATION2 can be avoided if the RIVKA and the G-RIVKA were replaced with the following *Reliable Source of Information View of Knowledge Attributions* (RSIVKA) and its group counterpart (G-RSIVKA):

RSIVKA: *A central purpose* of knowledge attributions is to identify or flag reliable sources of information.

G-RSIVKA: *A central purpose* of group knowledge attributions is to identify or flag reliable sources of information.

We saw above how the G-RSIVKA, as opposed to the G-RIVKA, can handle DISTRIBUTED INFORMATION2 with ease. The same is true of JUDGMENT AGGREGATION and OFFICIAL POSITION. If one wished to discern whether in JUDGMENT AGGREGATION the pay sacrifice is in the employees' best interest or whether there is not a causal connection between the MMR and autism, the groups in question and their members are surely reliable sources of information on these matters. In particular, all of the data needed to arrive at the relevant conclusions are located within the groups and their members. So, even if the groups are not properly regarded as reliable informants of the facts in question, the attributions of knowledge nonetheless identify or flag them as reliable sources of information.

There are other sorts of cases, which do not involve groups at all, that further support these replacement theses. Such cases again block the move of simply restricting the scope of the RIVKA to individuals. To see this, consider the following:

PRE-LINGUISTIC BABY: Tommy is a fifteen-month-old pre-linguistic baby whose cognitive abilities far surpass his verbal abilities. While talking to a friend, his mother says, "Tommy knows that when I pull out the octopus sippy cup, I will fill it with his beloved

milk whereas when I pull out the crocodile sippy cup, I will fill it with his dreaded ear infection medicine.”

Here, Tommy’s mother is clearly attributing knowledge to him—namely, that he knows which sippy cup will hold which drink. At the same time, however, she does not seem to be identifying or flagging a reliable informant when she does so. Tommy is a pre-linguistic baby and so there is an obvious sense in which he is not an informant at all. In particular, because he has not mastered any meaningful words in a human language, he is not yet in the realm of being a testifier capable of offering testimony to an audience. Indeed, we can imagine that Tommy has also failed to master any non-verbal signals that are distinctively communicative. For instance, he has not yet learned to nod, point, or gesture as a means of communicating with those around him. Thus, when Tommy’s mother attributes knowledge to him, it is purely on the basis of observing him and his behavior. Perhaps he smiles and drinks voraciously upon being handed the octopus sippy cup, but frowns and turns his head away when handed the crocodile sippy cup. Regardless of the details, the point that is here relevant is that Tommy is not yet capable of being an informant, let alone a reliable one. But clearly this is not the case when he is considered as a source of information on the topic of his attitudes toward sippy cups and their contents. When his mother attributes knowledge to him, I can certainly observe his behavior just as his mother can. I can watch his reactions when different cups and liquids are within his visual field and I can draw inferences about this data. Assuming that Tommy’s mother made an accurate assessment of his cognitive abilities, then, the attribution of knowledge to him clearly identifies or flags a reliable source of information. Hence, while PRE-LINGUISTIC BABY poses a problem for the RIVKA, it can easily be accommodated by the RSIVKA.

Now it should be noted that Craig is aware of this sort of problem and he responds by *objectivising* the concept of being a good informant.³³ In particular, he argues that in order for the

concept to have value in communication, we need to subtract from it what is required for being a good informant for a particular person on a specific occasion and focus instead on the detectability of an informant's reliability *in principle*. Thus, so long as an informant's reliability is objectively detectable by some possible means—even if it is not detectable by the means possessed by any actual person—this will suffice for being a reliable informant in the relevant sense.

Once the notion of being a good informant is objectivised in this way, however, the concept becomes so weak that it bears little resemblance to its intuitive self. For instance, even though Tommy is a pre-linguistic baby in the above case, we can certainly imagine a person with mind-reading powers who could discern the extent to which the baby, were he capable of communicating, would be reliable as an informant. Indeed, such powers could even enable one to read off the reliability of an entirely non-linguistic, non-human animal.³⁴ Craig's view, then, counts as reliable informants those who currently are not informants at all—such as pre-linguistic babies—and even those who never will be informants—such as some non-human animals. This strikes me as an unwelcome consequence. Thus, it is a virtue of the RSIVKA over the RIVKA that it can accommodate cases such as PRE-LINGUISTIC BABY while altogether eliminating the need for the appeal to objectivisation.

However, replacing the RIVKA and the G-RIVKA with the RSIVKA and the G-RSIVKA will not be met with enthusiasm by all. Ram Neta, for instance, claims that Craig avoids the problem of “wrongly classif[y]ng] thermometers as knowers” only by arguing that knowledge attributions flag informants, not mere sources of information (Neta 2006, p. 266). He writes:

Craig avoids [this] problem by appealing to the distinction between reliance upon an informant and reliance upon other sources of information. This distinction is grounded in our interest in cooperation and communication: reliance upon an informant can be part of cooperation or communication with that informant, but no such thing is true of reliance

upon other sources of information. Informants can take the credit or blame for the information that they provide, but sources of information cannot. (Neta 2006, pp. 266-7).

Neta's reasoning seems to go as follows: first, he seems to endorse what may be called the Knowledge Attribution/Knowledge Thesis (hereafter, the KA/KT):

KA/KT: If knowledge that p is properly attributed to S , then S knows that p .

Under one interpretation, the KA/KT is trivially true—to be a proper knowledge attribution *just is* to correctly identify a knower. But Craig and his followers intend a more substantive reading of this thesis where, instead of providing analyses of the concept of knowledge, we look at our practices and use paradigmatic attributions of knowledge as a guide to the presence of knowledge.

With this in mind, suppose that the RSIVKA, rather than the RIVKA, is true. Given this, knowledge attributions are proper when the subject of the attribution is a reliable source of information. But here is the problem: thermometers clearly can be reliable sources of information. Thus, the RSIVKA has the consequence that the following knowledge attribution can be proper: “The thermometer knows that it is 98 degrees outside.” If this is then combined with the KA/KT, we have the following absurd result: thermometers are knowers. According to Craig and Neta, this absurdity can be avoided by rejecting the RSIVKA in favor of the RIVKA. On this view, knowledge attributions are proper when the subject of the attribution is a reliable informant. Given that thermometers are not informants, knowledge cannot be properly attributed to them. Thus, we avoid wrongly classifying thermometers as knowers.

While replacing the RSIVKA with the RIVKA is one way of avoiding the absurd result that thermometers are knowers, it is certainly not the only one. What I propose instead is that the KA/KT be modified as follows:

KA/KT*: If knowledge that p is properly attributed to S , then either S knows that p or S is a group that has members that possess epistemic properties that ground the propriety of the attribution of knowledge that p to S .

While a full discussion of the second disjunct of the above consequent lies outside the scope of this paper, let me offer a few brief remarks on its behalf. To begin, consider DISTRIBUTED INFORMATION2: on my view, each member of the UN Population Commission possesses not only evidence about population growth in the United States, but also awareness of their roles in the group as a whole, which together justifies the attribution of knowledge to them that the birth rate of Latinos in the US is on the rise. In particular, all of these bits of evidence are precisely the epistemic properties that ground the propriety of the knowledge attribution in question. Of course, given the arguments raised earlier in the paper, there may be problems with attributing belief to the UN Population Commission that the birth rate of Latinos in the US is on the rise, where group belief is understood either summatively or non-summatively. So, my view here is compatible with either the possibility that groups may know that p without believing that p or with the possibility that proper group knowledge attributions may fail to pick out genuine instances of knowledge. I will not here choose between these options since my central purpose is to show that one can countenance the RSIVKA while nonetheless ruling out thermometers as knowers. The second disjunct of the KA/KT* does precisely this, and also hints at what guides our attribution of knowledge to groups in the first place.

Now, it should be emphasized that I am not offering the RSIVKA and the G-RSIVKA as straightforward replacements of the RIVKA and the G-RIVKA since I regard indentifying or flagging a reliable source of information as neither necessary nor sufficient for being a proper knowledge attribution. Instead, I take it that knowledge attributions have various functions, with indentifying or flagging reliable sources of information being simply one central purpose. For

instance, attributing knowledge sometimes has a *purely descriptive purpose*—e.g., “My deceased aunt knew that her daughter would eventually move to France.” Here, we can assume that my aunt did not leave any testimonial records that indicated her belief about her daughter’s future intentions. Thus, there may not be any sense in which my attribution of knowledge is meant to flag a reliable source of information since, given that my aunt is deceased, she may not be a relevant source at all. I may simply be describing a state of affairs. An attribution of knowledge also sometimes has the purpose of *identifying who is morally or legally responsible* in a given case. Consider, for example, the ascription discussed earlier from the *Los Angeles Times*: “BP officials knew about a problem on a crucial well safety device at least three months before the catastrophic April 20 explosion in the Gulf of Mexico but failed to repair it” (*Los Angeles Times*, July 21, 2010). We can assume for the sake of argument that it is generally regarded that BP officials are a highly unreliable source of information on matters relating to the oil spill in the Gulf of Mexico since they are trying to avoid any sort of culpability. Now, if this attribution were made, say, in a courtroom, its purpose may be to identify who is the bearer of either moral or legal responsibility for the oil spill, despite the fact that BP officials are unreliable about the very matter at issue. A knowledge attribution may also have the purpose of trying to *convince one’s interlocutor that knowledge is in fact present*. For instance, suppose that I say to my friend, who strongly denies that non-human animals are capable of possessing knowledge, “My dog clearly knows that we are going for a walk when I take her leash out of the drawer.” Here the central purpose of my knowledge attribution is to provide evidence to a naysayer that dogs are capable of being knowers in the first place.

Given that all of these attributions of knowledge are proper, it follows that identifying or flagging a reliable source of information is *only one* of the central purposes of attributing knowledge, an emphasis that is made clear in my formulations of the RSIVKA and G-RSIVKA. Moreover, there are obviously some reliable sources of information that are not at all appropriate subjects of

knowledge attributions. Footprints in the sand may be a reliable source of information about who was last on the beach, but we would not be at all inclined to ascribe knowledge to such a source. Perhaps, then, the only reliable sources of information that are here relevant are those that are subject or subject-like entities. But regardless of how the details are fleshed out, the point that is worth emphasizing is that the RSIVKA and the G-RSIVKA are meant to capture what made the RIVKA and the G-RIVKA intuitively plausible in the first place, without endorsing the overall project of the proponents of these views.

4. Conclusion

We have seen that paradigmatic instances of group knowledge attributions systematically fail to identify or flag reliable informants, whether this latter notion is understood in terms of being reliable believers or reliable testifiers. Since there are compelling reasons to reject an inflationary account of group attitudes and there are problematic cases for the RIVKA involving only individuals, the two most natural responses to my arguments are blocked. Thus, the G-RIVKA, and therefore the RIVKA, is false. We have also seen that these considerations pose problems for the related thesis that knowledge attributions are credit attributions since there are paradigmatic group knowledge attributions where there is no corresponding credit for the truth that is deserved. The RSIVKA and the G-RSIVKA not only avoid the problems posed by JUDGMENT AGGREGATION, OFFICIAL POSITION, DISTRIBUTED INFORMATION, and DISTRIBUTED INFORMATION², they also accommodate cases involving pre-linguistic babies or non-linguistic animals with ease. To the extent that a central purpose of knowledge attributions is to identify or flag anything, then, the focus should be on reliable sources of information.³⁵

References

- Alston, William P. 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257-99.
- Audi, Robert. 1997. "The Place of Testimony in the Fabric of Knowledge and Justification." *American Philosophical Quarterly* 34: 405-22.
- Bird, Alexander. Forthcoming. "Social Knowing: The Social Sense of 'Scientific Knowledge'." *Philosophical Perspectives*.
- Cariani, Fabrizio. 2011. "Judgment Aggregation." *Philosophy Compass* 6: 22-32.
- Coady, C.A.J. 1992. *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- Corlett, J. Angelo. 1996. *Analyzing Social Knowledge*. Lanham, MD: Rowman & Littlefield.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Elgin, Catherine Z. 2002. "Take It from Me: The Epistemological Status of Testimony." *Philosophy and Phenomenological Research* 65: 291-308.
- Fricker, Elizabeth. 1995. "Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony." *Mind* 104: 393-411.
- Fricker, Miranda. 2007. *Epistemic Injustice: Power & the Ethics of Knowing*. Oxford: Oxford University Press.
- Gilbert, Margaret. 1989. *On Social Facts*. London and New York: Routledge.
- Goldman, Alvin I. 2004. "Group Knowledge Versus Group Rationality: Two Approaches to Social Epistemology." *Episteme* 1:11-22.
- (2009). "Social Epistemology." *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2009/entries/epistemology-social/>>.
- Greco, John. 2007. "The Nature of Ability and the Purpose of Knowledge." *Philosophical Issues* 17: 57-69.

- , 2009. "Knowledge and Success from Ability." *Philosophical Studies* 142: 17-26.
- Hakli, Raul. 2007. "On the Possibility of Group Knowledge without Group Belief." *Social Epistemology* 21: 249-66.
- Kusch, Martin. 2009. "Testimony and the Value of Knowledge," in Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.), *Epistemic Value*. Oxford: Oxford University Press: 60-94.
- Lackey, Jennifer. 2006. "Learning from Words." *Philosophy and Phenomenological Research* 73: 77-101.
- , 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- , Unpublished. "A Deflationary Account of Group Testimony."
- McMahon, Christopher. 2003. "Two Modes of Collective Belief." *Protosociology* 18/19: 347-362.
- Meijers, A.W.M. 2002. "Collective Agents and Cognitive Attitudes." *Protosociology* 16: 70-85.
- Millgram, Elijah. 2009. "D'où venons-nous... Que sommes nous... Où allons-nous?" in Daniel Callcut (ed.), *Reading Bernard Williams*. New York: Routledge: 141-165.
- Neta, Ram. 2006. "Epistemology Factualized: New Contractarian Foundations for Epistemology." *Synthese* 150: 247-80.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, MA: The Belknap Press.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford: Oxford University Press.
- , Unpublished. "Anti-Luck Virtue Epistemology."
- Schaffer, Jonathan. 2007. "Knowing the Answer." *Philosophy and Phenomenological Research* 75: 383-403.
- Schmitt, Frederick F. 1992. "Review of *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*." *Mind* 101: 555-9.
- , 1994. "The Justification of Group Beliefs," in Frederick F. Schmitt (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham, MD: Rowman & Littlefield: 257-87.
- Sosa, Ernest. 1991. *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.

- , 1996. "Postscript to 'Proper Functionalism and Virtue Epistemology,'" in John L. Kvanvig (ed.), *Warrant in Contemporary Epistemology*. Lanham, MD: Rowman & Littlefield: 271-81.
- , 1999. "How Must Knowledge Be Modally Related to What Is Known?" *Philosophical Topics* 26: 373-84.
- , 2000. "Contextualism and Skepticism." *Philosophical Issues* 34: 94-107.
- , 2002. "Tracking, Competence, and Knowledge," in Paul Moser (ed.), *The Oxford Handbook of Epistemology*. Oxford: Oxford University Press: 264-87.
- Tollefsen, Deborah. 2007. "Group Testimony." *Social Epistemology* 21: 299-311.
- Tuomela, R. 2004. "Group Knowledge Analyzed." *Episteme* 1: 109-27.
- Williams, Bernard. 1973. *Problems of the Self*. Cambridge: Cambridge University Press.
- Wray, K. Brad. 2001. "Collective Belief and Acceptance." *Synthese* 129: 319-33.
- , 2003. "What Really Divides Gilbert and the Rejectionists." *Protosociology* 18/19: 363–376.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

¹ See also Greco (2009).

² Talk about the purpose of the concept of the knowledge, on the one hand, and the purpose of knowledge attributions, on the other, is used interchangeably in this context since it is the concept of knowledge that features in our attributions of knowledge. This is evident in the quoted passages above, where Williams, Craig, and Pritchard focus on the concept of knowledge while Greco and Neta, who explicitly follow Craig in their views, discuss knowledge attributions. Moreover, despite the fact that Craig himself focuses on the purpose of the concept of knowledge, his view is frequently cast in terms of knowledge attributions. For instance, Jonathan Shaffer writes, "The idea that knowledge ascriptions identify experts is due to Craig, who says that the role of the knowledge ascription is 'to flag approved sources of information'" (Schaffer, p. 391, note 12). Similarly, Martin Kusch says, "Craig rejects an analysis of *knowledge* in terms of necessary and sufficient conditions; he focuses primarily on the attribution of knowledge to others...." (Kusch 2009, p. 71).

And, according to Elijah Millgram, “[Edward] Craig proposed sidestepping the exercises in conceptual analysis that still dominate epistemology by asking what function ascriptions of knowledge serve.... The need he identified was a generic and transmissible certificate for information. To say ‘I know that p ’ conveys roughly that p is good enough to go on....” (Millgram 2009, p. 148).

³ Reliability seems to lie at the heart of what all the quoted theorists have in mind when talking about a “good” informant. Williams and Pritchard explicitly mention reliability in the passages quoted, and Greco defends a version of agent reliabilism in the papers in which he endorses Craig’s thesis. Moreover, in characterizing the notion of a good informant, Craig writes:

All of this is going to edge us towards the idea of someone who is a good informant as to whether p whatever the particular circumstances of the inquirer, whatever rewards and penalties hang over him and whatever his attitude to them. That means someone with a very high degree of reliability, someone who is very likely to be right—for he must be acceptable even to a very demanding inquirer. (Craig 1990, p. 91).

Given this, I will focus specifically on reliable informants rather than the generic “good informant” in what follows.

⁴ Fricker (2007) and Kusch (2009) are also sympathetic to the general spirit of the RIVKA.

⁵ I shall use knowledge “attributions” and “ascriptions” interchangeably in this paper.

⁶ See also Hakli (2007, p. 249).

⁷ There are two different views of group belief suggested in this passage. On the one hand, there is the *eliminativist* view, according to which it is literally false that groups believe things and hence group belief attributions are simply metaphorical. On the other hand, there is the *reductionist* view, according to which it is literally true that groups believe things, but such claims are made true by individual members of the group believing things. The summative account, as I am understanding it in this paper, is reductionist, not eliminativist.

⁸ For more on the theory of judgment aggregation, see List and Pettit (2002, 2004), Dietrich (2005), List (2005), and Pauly and van Hees (2006).

⁹ As will be clear later in this paper, Gilbert herself defends a non-summative account of group belief. In addition to the citations in the text, Corlett (1996) espouses a summative view of group belief.

¹⁰ I replaced “belief” with “judgment” here only to avoid confusion while we are considering a summative model of group belief.

¹¹ A similar argument is put forth by Frederick Schmitt in the following passage:

Suppose the Library Committee and the Budget Committee must each judge whether last year’s library budget is adequate for this year. The Budget Committee judges so, while the Library Committee judges not. Yet their members are the same (and let us suppose, to make matters simple, judge alike). It follows that one of the two committees believes a proposition p that no member believes. (Schmitt 1994, p. 261)

¹² This is how “hierarchical groups” in the sense characterized in Goldman (2004) function.

¹³ We can also imagine a scenario similar to the one found in OFFICIAL POSITION, but where all of the members of the group in question jointly accept that p *qua* group, despite the fact that each individual member has doubts and hence fails to believe that p . This case would still involve an “official position,” just not one decided by a board of the group. This type of case arises most frequently with what Gilbert calls *established* groups (2004, p. 96) since it is arguable that *non-established* groups do not have official positions or engage in joint acceptance in these sorts of ways. For a very nice discussion of non-established group knowledge, see Bird (forthcoming).

¹⁴ Schmitt (1994) and Tuomela (2004) also hold a joint acceptance view of group belief. For instance, Tuomela claims that “...when a group knows as a group...the members...are collectively committed to the content of the knowledge and to each other relative to it” (Tuomela 2004, p. 115).

¹⁵ Similarly, List and Pettit (2002) defend the following Impossibility Theorem:

There exists no aggregation procedure generating complete and consistent collective judgments that satisfies the following three conditions simultaneously:

Universal domain. The procedure accepts as admissible input any logically possible combinations of complete and consistent individual judgments on the propositions.

Anonymity. The judgments of all individuals have equal weight in determining collective judgments.

Systematicity. The collective judgment on each proposition depends only on the individual judgments on the proposition, and the same pattern of dependence holds for all propositions.

¹⁶ For arguments that group belief is voluntary in ways that individual doxastic states are not, see Wray (2001 and 2003), Meijers (2002), and Hakli (2007).

¹⁷ For a classic defense of the view that individuals lack direct voluntary control over their beliefs, see Alston (1988).

¹⁸ See Wray (2003) and McMahon (2003), respectively.

¹⁹ Goldman (2009) also provides the example of the tobacco companies adopting the position that smoking does not cause cancer in an effort to promote financial success.

²⁰ It should be noted that individuals often *do* have beliefs for purely instrumental reasons, e.g., because it is better for their image, expedient, and so on. The point that should be emphasized here is that individuals cannot generally be aware that their beliefs are rooted in such factors, but groups can be.

²¹ This case, along with others that tell against characterizing reliable informant in terms of BELIEF, can be found in my (2006 and 2008).

²² See Nozick (1981).

²³ See Pritchard (2005), Sosa (1996, 1999, 2000, and 2002), and Williamson (2000).

²⁴ See, for instance, Goldman (1976 and 1979) and Sosa (1991).

²⁵ It should be noted that Tollefsen herself does not espouse a summative view of group testimony.

²⁶ The example of the UN Population Commission issuing the *Charting the Progress of Populations* report is found in Tollefsen (2007), though all of the details and arguments to follow are my own. I also discuss this sort of case in detail in my (unpublished).

²⁷ This, of course, is not an exclusive 'or'; both (1) and (2) could be satisfied simultaneously.

²⁸ There are views of testimony in the literature other than those cited in the text. For instance, Peter Graham maintains that S testifies by making some statement that *p* if and only if: (i) S's stating that *p* is offered as evidence that *p*; (ii) S intends that his audience believe that he has the relevant competence,

authority or credentials to state truly that p ; and (iii) S 's statement that p is believed by S to be relevant to some question that he believes is disputed or unresolved (which may or may not be whether p) and is directed at those whom he believes to be in need of evidence on the matter. (Graham 1997, p. 227) And James Ross offers a definition of testimony that bears some similarities to Graham's account. According to Ross, testimony is "...any verbalized reporting of a purported state of affairs where the reporter intends that the hearer (reader, viewer, etc.) will take it on his report that the state of affairs is *as reported*" (Ross 1975, p. 3).

²⁹ Tollefsen (2007) is an exception here since she rejects a summative account of group testimony and argues that a reductionist view in the epistemology of testimony can be applied to groups themselves. For objections to Tollefsen's arguments, see my (unpublished).

³⁰ See Lackey (unpublished).

³¹ Of course, features of the group itself may largely determine whether the spokesperson is a reliable testifier, such as how thorough the members are in acquiring their data, how trustworthy they are in conveying it, and so on.

³² This is not to say that knowledge couldn't also be attributed to Maria.

³³ See Craig (1990). For an endorsement of this response, see Neta (2006).

³⁴ Additional problems have been raised to Craig's appeal to objectivisation in Schmitt (1992).

³⁵ For helpful comments on earlier drafts of this paper, I am grateful to Jessica Brown, Mikkel Gerken, Sandy Goldberg, Baron Reed, and audience members at the 2010 Knowledge Ascriptions Conference at the University of St. Andrews, Scotland.