

# A Systems-Oriented Approach to the Problem of Testimony

Kevin J.S. Zollman

April 2, 2014

I know that there is significantly more carbon dioxide in the Earth's atmosphere than there was one hundred years ago. I know this despite not having made any direct scientific observations that would count as evidence in favor of it. Instead, I know it from the testimony of others. The recognition that most of what we take ourselves to know comes, at least in part, through the testimony of others has prompted philosophical investigation into knowledge gained in this way.

The papers on this topic have mostly focused on a particular question, what is the epistemic warrant for believing the testimony of another? Suppose another individual tells us that some proposition,  $p$ , is true. Suppose further that we come to believe  $p$  as a result of that testimony. Finally, suppose it turns out to be the case that  $p$ . Under what conditions could we be said to know  $p$ ? Or put another way, when are we justified in believing that  $p$ ?<sup>1</sup>

There is another question that is closely related. When we are constrained as to the amount of testimony that we can solicit, how should we choose with whom to speak? For example, I may want to learn about some area of inquiry – say the Higgs-Boson – but I only have enough time to read one or two summaries of the state of research. What should I read? Shall I find specialists on the topic, or shall I listen to the first person willing to talk to me about it?

---

<sup>1</sup>In their papers on testimony, Burge (1995), Hinchman (2005), and others make a distinction between one having a justification for  $p$  and one being entitled to believe that  $p$ . This distinction is not critical for this paper, and I will use “justification” broadly to encompass any type of epistemic entitlement to come to believe  $p$ .

A third question is similar to the second. An interest in logical consistency might prevent me from believing the testimony of two different individuals one of whom tells me  $p$  and the other not- $p$ . Again, the source of our justification in believing testimony will have an important impact on how I resolve this disagreement between the testifiers (Goldman, 2001).

There are two considerations that relate the first question – regarding the source of justification for testimonial knowledge – to the second two – about whom to seek out and trust. First, if one ought to maximize one’s degree of justification (or alternatively maximize the chance that one is justified), then the source of justification will influence what norms govern those one believes. Second, if the topic is only of epistemic interest (if there are no additional concerns beyond purely epistemic ones) one might even go so far as to say the only normative constraints on whom one must seek out are those that come from the source of justification.

In broad strokes, there are three views regarding the source of justification for testimonial knowledge. The first view, which holds few modern adherents, declares that testimonial knowledge is never justified. One must come to know from direct observation or reasoning and not through the statements of others. As this view was stated most clearly by Descartes in his *Discourse on Method*, I will call this view the Cartesian view.

A second view, traced to David Hume, contends that testimonial knowledge is on a par with knowledge gained from measurement devices. I believe my thermometer because it has a history of providing me with reliable estimates of the temperature. Perhaps I have even tested it by placing it in boiling water or an ice water bath. Similarly for testimony, I can observe that the testimony of others tends to produce true beliefs, and I am thus justified in believing testimony on this basis. Because this view reduces the justification for believing in testimony to other, non-testimonial facts or beliefs, it is often called reductionism. Here we will refer to it as Humean view.

A final view, traced to Thomas Reid, balks at Hume’s reduction of other individuals to mere “truth gauges” (Hinchman, 2005, 580). Instead, under this view testimonial knowledge is of a different kind than knowledge gained from measurement devices, and one does not need an external reason for believing in the testimony of others. So long as the testifier is another rational agent, you have no positive reason to distrust what they say, and they offer you an assurance or an invitation to believe them, you are justified in believing them. While this view goes by many names, here we will call it the Redian view.

Although we have described each of these views in terms of their answer to the question of justification, they also have implications for the question of sought testimony. The Cartesian says that you should trust no one, and that rather than seeking out testimony you ought to spend your time inquiring on your own. The Humean says that you ought to seek out those who you have reason to suspect are most reliable. The Reidian, as we have described her so far, has no criteria to distinguish between rational individuals who have given you no reason to distrust them – indeed this fact has been used as a criticism against the Reidian view (Goldman, 2001). One might supplement the Reidian with a secondary criteria to allow them to seek out the more reliable, but for our purposes we will use the *extreme*-Reidian view that has not been so supplemented, while also noting that this view may be a caricature of the position taken by contemporary proponents.

Instead of tackling the first question, and then later focusing on the second (and third), this paper will focus entirely on the second. It will take as a starting point an individual who wishes to learn more and has the opportunity to seek out the testimony of a few others. We will then determine (using criteria spelled out in the next section) which testimonial strategy is epistemically superior. Ultimately this analysis generates a novel justification for the Reidian view that is likely to be more convincing from the Humean perspective than those previously offered.

## 1 Criteria for resolution

Proponents of the different views regarding testimony sometimes talk past one another because they rely of different metrics for evaluation. Humeans point to the apparent benefit their view of testimony has for an individual's reliability (Goldman, 1999, Chapter 4). Reidians occasionally point to other criteria in support of their view (e.g. Burge, 1995; Hinchman, 2005).<sup>2</sup> Without claiming that this paper resolves the debate definitively, we will focus on a particular set of criteria and ask which of these policies for seeking out testimony is superior *on those criteria*. Of course, someone might regard these the wrong criteria, but nonetheless the debate is advanced by clarifying how the various view perform on given dimensions.

---

<sup>2</sup>Some Reidians do appeal to truth-oriented considerations, however. They point out, for instance, that if we were forced to justify all of our testimonial beliefs on non-testimonial grounds, we might collapse into skepticism (cf. Coady, 1992; Pritchard, 2004).

As mentioned already the Humeans tend to tout the superiority of their theory from the perspective of its contribution to the true beliefs of the subject, its *veritistic* performance. Reidians sometimes point to other, non-truth conducive criteria. Those who point to other criteria recognize that truth is important, but they argue that it is not the only norm governing our epistemic practices. With full transparency, we will here focus only on veritistic considerations and leave the non-veritistic considerations for another time. Reidians can fairly point out that we are excluding their primary criteria, and when we criticize the Reidian policy we may be begging the question against them. However, when we provide a limited endorsement of the Reidian view – as we will – this will provide a very strong argument for the Reidian against the Humean.

Even restricting ourselves to veritistic considerations, there is a well known tension in epistemology. One naïve epistemic norm enjoins the inquirer to believe most of what is true; another calls on the inquirer to avoid believing falsehoods. One could satisfy the former by believing every proposition (including its negation). One could satisfy the latter by only believing tautologies. Epistemologists have not, and probably could not, settle on the appropriate trade-off between believing many true things and restraining oneself from believing falsehoods, and this paper will not presume there is a correct way of comparing these two desiderata. Instead we will evaluate each testimonial norm (Cartesianism, Humeanism, or Reidianism) first in virtue of its ability to minimize believed falsehoods and secondly in virtue of its ability to maximize the number of true things believed. The reader can then weight these two criteria in whatever way the reader deems appropriate.

A final dimension where evaluations might differ is in choosing the unit of analysis. Much of epistemology focuses on the individual inquirer, even if it considers evidence that comes from another. However, one might wish to evaluate epistemic practices from further away – taking an epistemic group as a whole. Doing so, we will ask how does a group of Cartesians, a group of Humeans, or a group of Reidians fare when considered *qua* group. This level of analysis is called “systems-oriented social epistemology” by Goldman (2011) and will be our starting point. We will return to the individual question in Section 5.

## 2 A computer model of testimony

Determining which group will perform better is a task that cannot be solved by reflection alone. Instead, we will turn to a computer aided thought experiment. We will evaluate how idealized communities of Humean, Reidian, and Cartesians perform according to our criteria. Insofar as our idealized communities provide a guide to how real communities will perform, we can then infer norms for real groups of inquirers from the performance of our artificial ones. Our artificial communities leave many things out, and this paper will not hide this fact. However, like the traditional thought experiment in philosophy, this simple situation provides a test-bed which reveals important subtleties that have heretofore gone unnoticed.

We begin with a community of 100 individuals. Each of them is constrained to solicit the testimony of some small number of individuals (we will consider 2, 4, 6, and 8 people as constraints). There is a very large number of logically independent propositions that the individuals would like to learn. Our individuals are allowed to take one of three doxastic attitudes toward each proposition, they might believe the proposition, disbelieve it, or withhold judgment. They begin life with information about a small number of propositions which they either believe or disbelieve, about everything else they withhold judgment.

Each individual is endowed with an intrinsic reliability which determines how likely she is to believe true propositions and disbelieve false ones when she receives evidence directly from the world. This propensity is different for each individual in the community, but stays constant for that individual over time. The reliability is the same for all propositions – in this respect, we are considering a community that is contemplating only one domain. The individuals are assigned these propensities at random so that the community on average has a reliability of 60%.<sup>3</sup>

Time is divided into a series of discrete steps (an idealization that I believe should have no impact on the final results). At each time step each individual ...

1. ...with a 10% chance, receives some new piece of information from direct observation of the world.
2. ...chooses a group of people from whom to solicit testimony.

---

<sup>3</sup>Formally, each individual is assigned a reliability by performing an independent draw from a beta distribution with parameters  $\alpha = 1.5$  and  $\beta = 1$ .

3. ...asks each of the individuals in that group to tell them something that they either believe or disbelieve
4. ...receives an honest report from the asked individual
5. ...if told something about which they currently suspend judgment, believes what they are told. Otherwise they do not change their doxastic attitudes.<sup>4</sup>

This process is repeated five hundred times, at which point the model ends and the performance of individuals is evaluated according to the dual interests of maximizing the number of true beliefs and minimizing the number of false beliefs.

We will make the three views of testimony operational in the following way. The Cartesians, who do not think testimony is justified, will not listen to anyone – they will only modify their doxastic attitudes in response to their own evidence from the world. Because everyone in our world tells the truth the radical Reidians have no justificatory reason to stop listening to someone. They begin life by choosing a set of individuals to trust and never change.

Making the Humean strategy operational is slightly more difficult because there is more than one way for the Humean to determine if an individual is reliable. At one extreme, we might imagine that each Humean had a crystal ball that would tell her what proportion of another individual's beliefs were true. This represents a situation where one can calibrate another by directly comparing that individual to the world. Calibrating a weather forecaster might represent the most plausible situation where this would be possible. We will call this the “objective calibration” approach.

Often, however, we do not have direct access to the world. It is unlikely that I will ever be in a position to directly calibrate the predictions of cosmologists, for instance. In such a situation one must calibrate differently. In this case the best one can do is to see how many beliefs one shares with another. To be more precise, one might determine from the set of propositions about

---

<sup>4</sup>This means that in this model there is no “belief revision.” This is certainly an idealization. This assumption has been made for two reasons. First, following the literature on testimony this model focuses primarily on the acquisition of new beliefs not on belief revision. And second, there is no uncontroversial way to model belief revision especially in the context of qualitative beliefs. Important future work should tackle this question directly to determine how robust the findings are to modifications of this assumption.

which both individuals have opinions on what proportion of propositions the individuals agree. Call this the “subjective calibration” approach.<sup>5</sup>

In the model we assume something akin to Fricker’s (1995) “local” calibration. That is, each individual compares how a potential testifier’s total corpus of belief compares to her own total corpus of belief without distinguishing between those things that each has acquired through testimony or through direct experience. As both Fricker (1995) (a Humean) and Coady (1992) (a Reidian) point out, because so many of our beliefs come through testimony we cannot set aside all testimonial beliefs as a whole in order to evaluate them as compared to non-testimonial beliefs.

The subjective approach, from the third person perspective, has some obvious drawbacks. If an individual is often wrong, he will regard an individual who is often right as a unreliable partner instead preferring someone, like him, who believes many false things. But, from the first person perspective – which is all we have regarding ourselves – it is the best we can do. If we truly thought we were wrong more often than we were right, we would change our minds.<sup>6</sup> Whatever your attitudes about the normative appropriateness of this standard, it seems likely that this is in fact how people often calibrate others (Kahan et al., 2011) and as a result it is worthy of investigation.

By modeling the testimonial situation in this way, we have already excluded one of the central arguments in favor of the Reidian view. Coady (1992), Lackey (2010), Insole (2000) and others argue that the Reidian posi-

---

<sup>5</sup>Kitcher (1993, chapter 8) calls this “direct calibration” in contrast to “indirect calibration” where one relies on another to judge an individual’s reliability. For the purposes of this model subjective calibration is achieved in the following way. Each individual assigns every other individual an “agreement score” which is the number of propositions they both believe plus the number of propositions that they both disbelieve minus the number of propositions that one believes and the other disbelieves. The subjective Humean then seeks out those who are highest according to this score.

<sup>6</sup>Kadane and Lichtenstein (1982) discuss this issue within the context of a Bayesian model of belief. They prove that any consistent Bayesian must regard themselves as well calibrated. This way of modeling subjective calibration stands in contrast to other approaches (cf. Lehrer and Wagner, 1981) which posit that individuals have second-order reliabilities – that they have some intrinsic ability to recognize the reliability of another individual. It is beyond the scope of this paper to engage in a philosophical debate about this methodology. But, it strikes me as strange to suppose that it normatively permissible that (a) Carlos believes that Jake is very reliable about some domain, (b) for most propositions  $p$  in this domain, Carlos knows that he and Jake disagree about the truth of  $p$ , and (c) Carlos refuses to change his mind about any of these propositions or about his view of Jake’s reliability.




Figure 1: *The average proportions of an individual’s non-abstention beliefs that are true for each of several testimonial strategies. The number on the x-axis are the number of people from whom a given individual can solicit testimony. The bars represents the means for the individuals from 100 runs of the simulation. The error bars represent the 95% confidence interval for the average performance of a community of 100 individuals.*

tion must be correct because we are simply not in a position to determine the another’s calibration (even subjectively). In this model everyone has, in principle access to at least a subjective estimate of the calibration of another. As before, we will note that we have been unfair to some of the Reidians considerations which will make our justification of Reidianism stronger.

### 3 Minimizing falsehoods

First we will consider how these different populations fare when judged by the degree to which they minimize false beliefs. The averages of 100 simulations for each setting of the parameters are presented in figure 1.<sup>7</sup> The y-axis represents the mean proportion of an individual’s beliefs that are true

<sup>7</sup>All data and simulation code will be provided via the author’s website and, if possible, via the journal publisher’s website as supplementary material.



(ignoring abstentions). On the x-axis is the number of individuals that each member of the population solicits for testimony.

At the far left we have the Cartesian view. Here, no one is solicited for testimony. Because the average reliability of individuals in the population is 0.6, we know that this will be the long-run reliability of the group. For each of the 2, 4, 6, and 8 entries we see a difference between the other three testimonial strategies. The red (right) bars represent the Humean strategy which is able to objectively calibrate others. Unsurprisingly such a strategy does very well in this context, far surpassing any of the other strategies considered.

As noted before, in many (perhaps even most) domains of inquiry, objective calibration is not possible. Instead, the Humean must turn to subjective calibration. The subjective Humean is represented by the blue (center) bars, and the Reidian strategy is represented by the yellow (left) bars. The reader may be surprised to see how little difference there is between the subjective Humean strategy and the Reidian strategy. Why should this be so? After all the Humean is paying some attention to the reliability of others while the Reidian is not.

This question can be answered by interrogating the social network formed by these trust relations. Figure 2 illustrates the trust relations formed by a community made up of subjectively calibrating Humeans. One can see immediately that the community has been broken up into many different sub-communities (formally known as *components*). These sub-communities are relatively homogenous with respect to the members' reliability, but differ radically from one another. The members of some communities have a large proportion of their beliefs that are true, while the members of others have very few. This phenomena – that relatively unreliable people come to trust other unreliable people – may strike many readers as familiar.<sup>8</sup> What this simulation shows is that this process balances out, almost perfectly, the countervailing process where relatively reliable people come to trust other reliable people.

That similar people come to associate disproportionately with one another is known as homophily. The investigation of alternative models with very different assumptions has shown that this harms the ability of communities to

---

<sup>8</sup>Pritchard (2004) argues that Fricker's (1995) version of Humeanism allows for circular justification of many beliefs acquired via testimonial. Beyond the philosophical concern with circularity, these simulation results illustrate that this circularity can have significant negative effects for certain people in an epistemic community.




Figure 2: *A social network formed by a community of subjective Humeans who can solicit testimony of two other individuals. The arrows represent the trust relation with the person at the tail of the arrow soliciting testimony from the person at the head of the arrow. Double headed arrows represent individuals who solicit testimony from one another. The color of the individuals represents their reliability with black representing very reliable and bright red indicating unreliable.*

successfully learn about the world (Hegselmann and Krause, 2006; Holme and Newman, 2006; Golub and Jackson, 2011). The convergence in the results of these models suggests that at least in this case, this result is not an artifact of this model.

In contrast to the community of subjective Humeans, the Reidian’s community is far more homogenous. Sub-communities are not formed often, and when they do they are relatively large. As a result the reliability of every member of the entire community quickly approaches the mean of the entire community (which in this case is 0.6). While the Humean, Reidian, and Cartesian strategies have little to differentiate them in terms of *average* reliability, they are different in terms of the variation in individual’s reliability. If one is an epistemic egalitarian then one has a strong reason to prefer a community of Reidians to a community of Cartesians or a community of Humeans when the Humeans can only use subjective calibration.

As we turn to the second criterion for evaluation, maximizing the number of true things one believes, we find an even stronger endorsement for the Reidian strategy.

## 4 Maximizing truths

Figure 3 shows what happens when we count the number of true things an individual believes. First, one can now see an illustration of the obvious argument against the Cartesian strategy. By ignoring others as a source of evidence, Cartesians radically restrict their ability to learn facts about the world. There are, however, two less obvious results that deserve discussion. First is that Reidians outperform subjective Humeans, and second, that in many cases the subjective Humeans outperform Humeans who objective calibrate.

Before we turn to understanding the cause for this discrepancy, allow me to first say why it is *not* the case. As mentioned before, a number of Reidians argue that were we to require some external justification for the belief in the testimony in others we would have to believe significantly less (e.g. Coady, 1992; Lackey, 2010). This is not what is occurring here. The Reidians and both types of Humeans receive the same total amount of testimony. The difference is how they choose whom to contact. Rather than in the total amount of testimony, the explanation for the difference between the different testimonial strategies is to be found instead in the content of the testimony.




Figure 3: *The total number of true things that an individual believes. The number on the x-axis are the number of people from whom a given individual can solicit testimony. The bars represents the means for the individuals from 100 runs of the simulation. The error bars represent the 95% confidence interval for the average performance of a community of 100 individuals.*

In order to demonstrate this claim, we will first compare the Reidian strategy to the subjective Humean strategy. Again, figure 2 is illustrative. Notice that many of the sub-groups are very small. Many contain only three individuals. Once these sub-groups form, they are very stable, and as a result, for these individuals the epistemic community becomes very small – they are really only capable of learning from two others. Even the larger sub-groups represent a very small fraction of the total community. This is does not occur in a Reidian community, and every member continues to have access to a large fraction of the entire group.

Second, consider the subjective Humean as compared to the objective Humean. A objective Humean society is even more radically constrained, but it is less obvious why this happens. Consider for a moment a random individual in the society of objective Humeans, call him Randy. Randy connects to the two most reliable individuals in the community, call them Shannon and Julie. Now consider Shannon. She connects to the two most reliable people in the community (other than her). One of them must be Julie. Call the other one Carlos. Now who does Julie connect to? Shannon and Carlos. Who does Carlos connect to? Julie and Shannon. From Randy’s perspective the potential sources of information are relatively few, only Julie, Shannon, and Carlos.

In the subjective Humean society, on the other hand, the size of those small communities grows slightly faster. Already in the case where one can only solicit two individuals some sub-communities have more than four members. This is why, as the number of testimonial sources grows, the subjective Humean strategy comes to outperform the objective Humean strategy.

The failure of the objective Humean strategy in this context suggests an alternative testimonial strategy which has heretofore not been considered. These simulations point out how the Humean strategy focuses solely on minimizing the number of falsehoods, but like the Cartesian strategy, fails to take advantage of the benefits of social interaction. One might imagine an individual who seeks out, not the most reliable members of the community, but the one who has the most beliefs. This represents the maximize true beliefs dual to the Humean strategy.

## 5 Individual Choice

So far, we have focused entirely on how communities as a whole perform when compared against one another. As is now well known from discussions of social dilemmas like the Prisoner's dilemma and tragedy of the commons, it can be the case that individuals who maximize their own interests might make the community as a whole worse off. As a result, we cannot immediately move from the social level to the individual level.

In order to investigate how individuals fare, heterogeneous communities were evaluated. Simulations were run where part of the community followed the Reidian strategy and part followed the subjective Humean strategy. A second set considered hybrid communities made up of Redians and objective Humeans. In these cases the performance of the different strategies was compared to see which individuals would prefer to switch in these cases.<sup>9</sup>

No social dilemmas were found. In all cases, the strategy which made the group better was also better for the individual. And, in the case where the Reidian and the subjective Humean appear approximately equal, there was no strong indication that one is better for the individual. Thus, we can conclude that not only is the Reidian strategy sometimes better for the group when judged on veritistic grounds but it is better for the individual as well.

## 6 Conclusion

Our central focus for analysis has been on the question: when one can only seek out a limited number of individuals who should one solicit for testimony? We restricted our analysis to only veritistic considerations, which we recognized might elicit a complaint from the Reidian camp who often focus on other metrics of evaluation. Despite prejudicing the matter in this way, we found several endorsements for the Reidian view.

First, when one cannot objectively calibrate testifiers, one has little reason to prefer the Humean strategy to the Reidian strategy on veritistic grounds. Furthermore, if one cares about maintaining epistemic equality in addition to average performance, one now has a positive reason to prefer the Reidian strategy to the subjective Humean one. When one turns to attempting maximizing the number of true things one believes one has a strong rea-

---

<sup>9</sup>Simulations considered heterogeneous populations made up of individuals who could only solicit two others for testimony.

son to prefer the Reidian strategy to the Humean strategy regardless of one's method of calibration. This is surprising precisely because the Humean strategy seems so directly targeted at veritistic considerations.

We have not yet endeavored to answer the more commonly discussed question about the grounds for justification of testimonial knowledge. This question is more abstract than the one addressed here, and its connection to action is more remote. Truth oriented considerations like the one here are, of course, related to questions of justification. In particular, if one is to argue that the justification for testimony is grounded in one's judgment of the reliability of the testifier, even in contexts where that judgment is made via subjective calibration, one must now explain why seeking justification will make one less reliable. I will not claim that no one could make such a case, but it should now seem like a difficult hill to climb.

## References

- Burge, T. (1995). Content Preservation. *Philosophical Issues* 6, 271–300.
- Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford: Oxford University Press.
- Fricker, E. (1995). Telling and Trusting : Reductionism and Anti-Reductionism in the Epistemology of Testimony. *Mind* 104(414).
- Goldman, A. (1999). *Knowledge in a Social World*. Oxford: Clarendon Press.
- Goldman, A. (2001). Experts: Which Ones Should You Trust ? *Philosophy and Phenomenological Research* 63(1), 85–110.
- Goldman, A. (2011). Systems Oriented Social Epistemology. In A. Goldman and D. Whitecomb (Eds.), *Social Epistemology: Essential Readings*. Oxford University Press.
- Golub, B. and M. O. Jackson (2011). How Homophily Affects the Speed of Learning and Best-Response Dynamics.
- Hegselmann, R. and U. Krause (2006). Truth and Cognitive Division of Labor: First Steps Toward a Computer Aided Social Epistemology. *Journal of Artificial Societies and Social Simulation* 9(3), 1–29.

- Hinchman, E. S. (2005). Telling as Inviting to Trust. *Philosophy and Phenomenological Research* LXX(3), 562–587.
- Holme, P. and M. E. J. Newman (2006). Nonequilibrium Phase Transition in the Coevolution of Networks and Opinions. *Physical Review E* 74(056108), 1–5.
- Insole, C. J. (2000, January). Seeing off the Local Threat to Irreducible Knowledge by Testimony. *The Philosophical Quarterly* 50(198), 44–56.
- Kadane, J. B. and S. Lichtenstein (1982, February). A Subjectivist View of Calibration.
- Kahan, D. M., H. JenkinsSmith, and D. Braman (2011, February). Cultural cognition of scientific consensus. *Journal of Risk Research* 14(2), 147–174.
- Kitcher, P. (1993). *The Advancement of Science*. New York: Oxford University Press.
- Lackey, J. (2010). *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- Lehrer, K. and C. Wagner (1981). *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*. Dordrecht: D. Reidel.
- Pritchard, D. (2004). The epistemology of testimony. *Philosophical Issues* 14.