

Modeling the social consequences of testimonial norms

Kevin J. S. Zollman

© Springer Science+Business Media Dordrecht 2014

Abstract This paper approaches the problem of testimony from a new direction. Rather than focusing on the epistemic grounds for testimony, it considers the problem from the perspective of an individual who must choose whom to trust from a population of many would-be testifiers. A computer simulation is presented which illustrates that in many plausible situations, those who trust without attempting to judge the reliability of testifiers outperform those who attempt to seek out the more reliable members of the community. In so doing, it presents a novel defense for the credulist position that argues one should trust testimony without considering the underlying reliability of the testifier.

Keywords Testimony · Social epistemology · Computer simulation

1 Introduction

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 testimony. The recognition that most of what we take ourselves to know comes through the testimony of others has prompted philosophical investigation into knowledge gained in this way.

Papers on this topic have mostly focused on a particular question: what is the epistemic warrant for believing the testimony of another? Suppose another

Electronic supplementary material The online version of this article (doi:[10.1007/s11098-014-0416-7](https://doi.org/10.1007/s11098-014-0416-7)) contains supplementary material, which is available to authorized users.

K. J. S. Zollman (✉)
Department of Philosophy, Carnegie Mellon University, Baker Hall 161,
Pittsburgh, PA 15213, USA
e-mail: kzollman@andrew.cmu.edu

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 ?¹

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 read the first summary I come across?

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 tells me not- p . Again, the source of my justification in believing testimony will have an important impact on how I resolve this disagreement between the testifiers (Goldman 2001).

The first question—regarding the source of justification for testimonial knowledge—is closely related to the second two—about whom to seek out and trust. 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. Furthermore, 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. I will call this the skeptical view.

A second view, often 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 calibrated 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.

A final view, traced to Thomas Reid, balks at Humes’ reduction of other individuals to mere “truth gauges” (Hinchman 2005, p. 580). Instead, on 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

¹ 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 .

positive reason to distrust what she says, and she offers you an assurance or an invitation to believe her, then you are justified in believing her. While this view goes by many names, here we will call it the credulist 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 skeptic says that you should trust no one, and rather than seeking out testimony you ought to spend your time inquiring on your own. The reductionist says that you ought to seek out those who you have reason to suspect are most reliable. The credulist, 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 credulist view (Goldman 2001). One might supplement credulism with a secondary criteria to allow inquirers to seek out more reliable people. But for our purposes we will use the *extreme-credulist* 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 credulist view that is likely to be more convincing from the perspective of a reductionist than those previously offered. This argument focuses on the effect that adopting particular testimonial strategies has on the reliability of the community as a whole instead of considering it just from the perspective of the individual. The results presented here illustrate several ways in which a reductionist community fares worse than a credulist community.

2 Criteria for resolution

Proponents of the different views regarding testimony sometimes talk past one another because they rely of different metrics for evaluation. Reductionists point to the apparent benefit their view of testimony has for an individual's reliability (Goldman 1999, Chapt. 4). Credulists occasionally point to other criteria in support of their view (e.g. Burge 1995; Hinchman 2005).² 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.

With full transparency, we will here focus only on veritistic considerations and leave the non-veritistic considerations for another time. Some credulists can fairly point out that we are excluding their primary criteria, and when we criticize the

² Some credulists 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).

credulist policy we may be begging the question against them. However, when we provide a limited endorsement of the credulist view—as we will—this will provide a very strong argument for the credulist against the reductionist.

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 later by only believing tautologies. This debate quite easily reproduces itself in testimony. Those who accept testimony more readily will believe more with higher risk of error, while those who do not run the risk of knowing less (Gelfert 2014).

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 (skepticism, reductionism, and credulism) 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 what ever 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 a different level by taking an epistemic group as a unit of analysis. Doing so, we will ask how does a group of skeptics, a group of reductionists, or a group of credulists fare when considered qua group. This level of analysis is called “systems-oriented social epistemology” by Goldman (2011) and will our be starting point. We will return to the individual question in Sect. 6.

3 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 skeptics, reductionists, and credulists 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. These 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 two, four, six, and eight 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.³

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%.⁴

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.
3. ...asks each of the individuals in that group to tell them something that the testifier either believes or disbelieves.
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.⁵

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 skeptics, 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 is reporting their honest belief the radical credulists have no justificatory reason to stop listening to someone. They begin life by choosing a set of individuals to trust and never change.

Making reductionism operational is slightly more difficult because there is more than one way for the reductionist to determine if an individual is reliable. At one extreme, we might imagine that each individual 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.

³ In the simulations presented here, there are 1,500 propositions and individuals begin life with non-abstention beliefs about 15 of them.

⁴ Formally, each individual is assigned a reliability by performing an independent draw from a beta distribution with parameters $\alpha = 1.5$ and $\beta = 1$.

⁵ This means that in this model there is no “belief revision.” This is certainly an idealization, which 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. The later issue, called peer disagreement, has an extensive literature which will not be addressed here. 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.

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. 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 which both individuals have opinions on what proportion of propositions the individuals agree. Call this the “subjective calibration” approach.⁶

This subjective approach is 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 reductionist) and Coady (1992) (a credulist) 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 another 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.⁷ 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 credulist view. Coady (1992), Lackey (2010), Insole (2000) and others argue that the credulist position must be correct because we are simply not in a position to calibrate another (even subjectively). In this model everyone has in-principle access to at least a subjective estimate of the reliability of another. As before, we will note that we have been unfair to some considerations in favor of credulism which will make our justification of credulism stronger.

⁶ Kitcher (1993, chapt. 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 reductionist then seeks out those who are highest according to this score.

⁷ 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 is 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. In the context of qualitative belief it might be possible to maintain this attitude in settings similar to the lottery paradox.

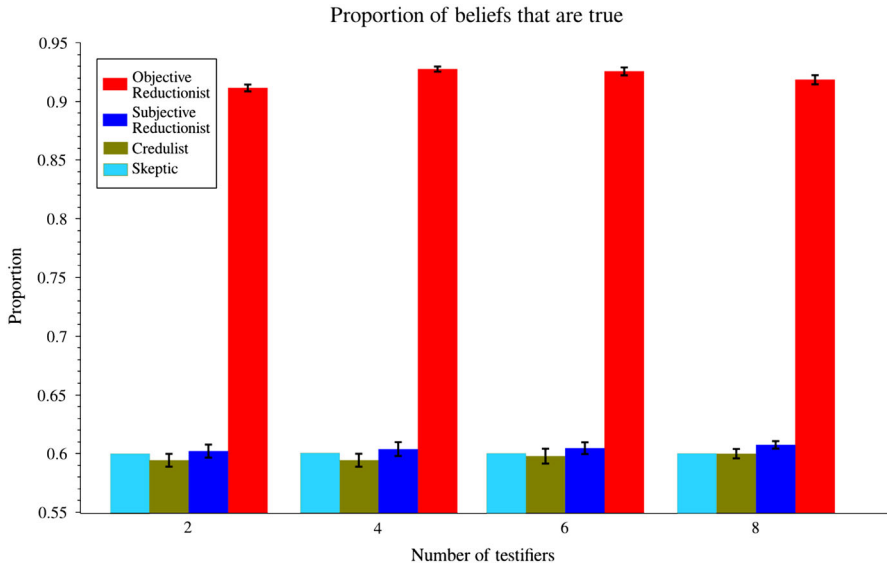


Fig. 1 The average proportions of an individual's non-abstention beliefs that are true for each of several testimonial strategies. The *number* on the *x-axis* are the number of people from whom a given individual can solicit testimony. The *bars* represent the means for the individuals from 100 runs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

4 Minimizing false beliefs

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 Fig. 1.⁸ The *y-axis* represents the mean proportion of an individual's beliefs that are true (ignoring abstentions). On the *x-axis* is the number of individuals that each member of the population solicits for testimony.

First consider the skeptical view where testimony is totally ignored. 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 two, four, six, and eight entries we see a difference between the other three testimonial strategies. The red (far right) bars represent the reductionist 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 reductionist must turn to subjective calibration. The subjective reductionist is represented by the dark blue (right center) bars, and the credulist strategy is represented by the yellow (center left) bars. The

⁸ All data and simulation code will be available as electronic supplementary material.

reader may be surprised to see how little difference there is between the subjective reductionist strategy and the credulist strategy. Why should this be so? After all the reductionist is paying some attention to the reliability of others while the credulist 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 reductionists. 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.⁹ 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 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 reductionists, the credulist 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 subjectively reductionist, credulist, and skeptic strategies have little to differentiate them in terms of *average* reliability, they are different in terms of the variation in individual's reliability. In a credulist community most individuals end up with approximately 60 % of true beliefs. While the overall community average is the same in a reductionist community, there are more individual differences. Some individuals end up with a large proportion of their beliefs being true, while others are left with mostly false beliefs.

When situations like this are encountered with wealth—where two countries have the same average wealth level, but one has more rich and more poor than the other—people often appeal to egalitarian considerations to argue for one community over another. A similar argument might be given in social epistemology. If one is an epistemic egalitarian then one has a strong reason to prefer a community of credulists to a community of skeptics or a community of reductionists when the reductionists can only use subjective calibration. In the former, there is more epistemic equality than in the later two.

⁹ Pritchard (2004) argues that Fricker's (1995) version of reductionism allows for circular justification of many beliefs acquired via testimony. Beyond the philosophical concern with circularity, these simulation results illustrate that this can have significant negative effects for certain people in an epistemic community.




Fig. 2 A social network formed by a community of subjective reductionists who can solicit testimony from 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. (Color figure online)

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 credulist strategy.

5 Maximizing truths

Figure 3 shows what happens when we count the number of true things an individual believes. One can now see an illustration of the obvious argument against the skeptical strategy. By ignoring others as a source of evidence, skeptics radically restrict their ability to learn facts about the world. There are, however, two less obvious results that deserve discussion. First is that credulists outperform subjective reductionists, and second, that in many cases the subjective reductionists outperform reductionists 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 philosophers have argued that were we to require some external justification for the belief in the testimony in others we would believe significantly less (e.g. Coady 1992; Lackey 2010; Gelfert 2014). This is not what is occurring here. The credulists and both

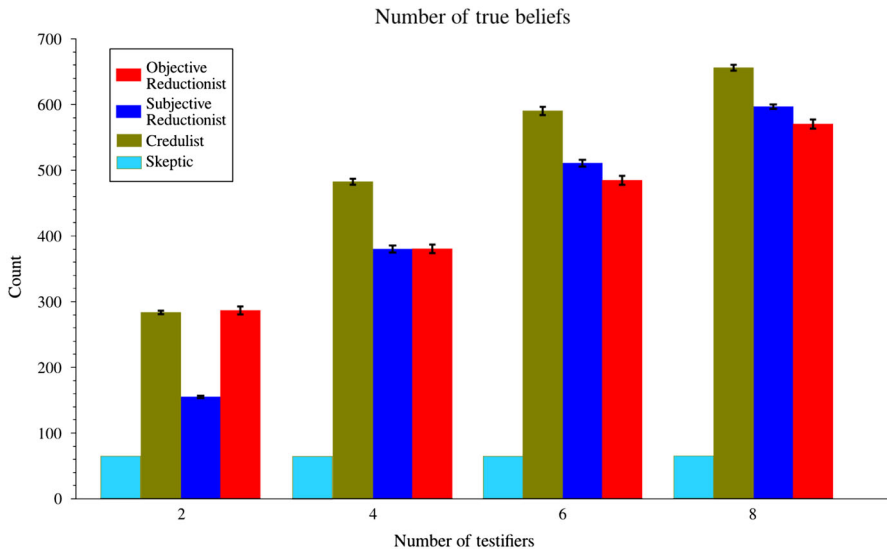


Fig. 3 The total number of true things that an individual believes. The *number* on the *x-axis* are the number of people from whom a given individual can solicit testimony. The *bars* represents the means for the individuals from 100 runs of the simulation. The *error bars* represent the 95 % confidence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

types of reductionists 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.

In order to demonstrate this claim, we will first compare the credulist strategy to the subjective reductionist strategy. Again, Fig. 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 does not occur in a credulist community; every member continues to have access to a large fraction of the entire group.

Second, consider the subjective reductionist as compared to the objective reductionist. A objective reductionist 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 reductionists, 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 reductionist 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 reductionist strategy comes to outperform the objective reductionist strategy.

These considerations also illustrate another important methodological point. Most of the discussion of testimony in the philosophical literature has focused on the impact of those strategies for the individual. From the perspective of a focal individual things are very similar in the communities that adopt different testimonial strategies; everyone listens to two people. However, when we focus on the group, one discovers the importance of seeing the whole chain of testimony, not just one step. This would not be identifiable if we did not consider the testimonial strategy of would-be testifiers.

The failure of the objective reductionism in this context suggests an alternative testimonial strategy which has heretofore not been considered. These simulations point out how the reductionist strategy focuses solely on minimizing the number of falsehoods, but like the skeptic, 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 reductionist strategy.

6 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 credulist strategy and part followed the subjective reductionist strategy. A second set considered hybrid communities made up of credulists and objective reductionists. In these cases the performance of the different strategies was compared to see which individuals would prefer to switch in these cases.¹⁰

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 credulism and the subjective reductionism appear approximately equal, there was no strong indication that one is better for the individual. Thus, we can conclude that not only is the credulist strategy sometimes better for the group when judged on veritistic grounds

¹⁰ Simulations considered heterogeneous populations made up of individuals who could only solicit two others for testimony.

but it is better for the individual as well regardless of the strategies adopted by others in her community (at least within the parameters of these simulations).

7 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 credulist camp that often focuses on other metrics of evaluation. Despite prejudicing the matter in this way, we found several endorsements for credulism.

First, when one cannot objectively calibrate testifiers, one has little reason to prefer the reductionist strategy to credulism 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 credulist strategy to subjective reductionism. When one turns to attempting maximizing the number of true things one believes one has a strong reason to prefer the credulist strategy to the reductionist strategy regardless of one's method of calibration. This is surprising precisely because reductionism 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 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.

Acknowledgments The author would like to thank an anonymous reviewer and audiences in Pittsburgh, Groningen, Munich, and Düsseldorf for helpful comments. This research was supported by National Science Foundation grants SES 1026586 and SES 1254291.

References

- Burge, T. (1995). Content preservation. *Philosophical Issues*, 6, 271–300.
- Coady, C. A. J. (1992). *Testimony: A philosophical study*. Oxford: Oxford University Press.
- Fricke, E. (1995). Telling and trusting : Reductionism and anti-reductionism in the epistemology of testimony. *Mind*, 104, 393–414.
- Gelfert, A. (2014). *A critical introduction to testimony*. London: Bloomsbury.
- 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 & D. Whitecomb (Eds.), *Social epistemology: Essential readings*. New York: Oxford University Press.

- Golub, B., & Jackson, M. O. (2011). How homophily affects the speed of learning and best-response dynamics. *Quarterly Journal of Economics*, *127*(3), 1287–1338.
- Hegselmann, R., & Krause, U. (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., & Newman, M. E. J. (2006). Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E*, *74*(056108), 1–5.
- Insole, C. J. (2000). Seeing off the local threat to irreducible knowledge by testimony. *The Philosophical Quarterly*, *50*(198), 44–56.
- Kadane, J. B., & Lichtenstein, S. (1982). A subjectivist view of calibration. *American Statistician*, *36*, 240–255.
- Kahan, D. M., JenkinsSmith, H., & Braman, D. (2011). 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., & Wagner, C. (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*, 326–348.