LUND UNIVERSITY

A Simulation Approach to Veritistic Social Epistemology

Olsson, Erik J

Published in:
Episteme

2011

[Link to publication](#)

Citation for published version (APA):

Olsson, E. J. (2011). A Simulation Approach to Veritistic Social Epistemology. *Episteme*, 8(2), 127-143.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

A Simulation Approach to Veritistic Social Epistemology

Erik J. Olsson

Lund University

Abstract: In a seminal book, Alvin I. Goldman outlines a theory for how to evaluate social practices with respect to their "veritistic value", i.e., their tendency to promote the acquisition of true beliefs (and impede the acquisition of false beliefs) in society. In the same work, Goldman raises a number of serious worries for his account. Two of them concern the possibility of determining the veritistic value of a practice in a concrete case because (1) we often don't know what beliefs are actually true, and (2) even if we did, the task of determining the veritistic value would be computationally extremely difficult. Neither problem is specific to Goldman's theory but can be expected to arise for just about any account of veritistic value. It is argued here that the first problem does not pose a serious threat to large classes of interesting practices. The bulk of the paper is devoted to the computational problem which, it is submitted, can be addressed in promising terms by means of computer simulation. In an attempt to add vividness to this proposal, an up-and-running simulation environment (Laputa) is presented and put to some preliminary tests.

1. Introduction

In the words of its foremost practitioner, Alvin I. Goldman, “[v]eritistic social epistemology aims to evaluate social practices in terms of their veritistic outputs, where veritistic outputs includes states like knowledge, error and ignorance” (Goldman, 1999, p. 87). Examples of social practices are telling the truth, lying, bullshitting, trusting other people, asking a friend, engaging in inquiry, and so on. In his 1999 book, Goldman focuses on the tendency of practices to produce *true belief* in the participants, true belief representing in his view a weak form of knowledge. We will follow him in that respect.¹

In many cases it is pretty clear whether a practice promotes true belief. For instance, telling the truth promotes truth, whereas lying does not. In other cases, it is initially uncertain whether a given practice does or does not promote truth. Consider the practice of saying *p* just in case one considers *p* more likely than not. Does that practice promote truth in society? On the positive side, *p* is likely to be asserted and disseminated if it is true. On the negative, *p* is likely to be asserted and disseminated even if it is false. Or consider the two practices of lying and bullshitting, taking the latter in Harry Frankfurt’s sense of saying something regardless of its truth.² It is difficult to tell which of lying or

¹ Goldman acknowledges of course that there is also a stronger, reliabilist sense of knowledge according to which knowledge amounts to true belief plus reliable belief acquisition and an anti-Gettier condition. While Goldman is the most prominent advocate of the reliabilist theory, knowledge in that stronger sense does not play a substantial role in his writings on social epistemology. For the strong reliabilist sense of knowledge, see for example chapter 3 in Goldman (1986), and for a recent argument for the existence of a weak sense of knowledge Goldman and Olsson (2009).

² See Frankfurt (2005) and, for a discussion, Olsson (2008).

bullshitting is more likely to promote truth or, rather, less likely to impede truth. In order to handle such more complicated cases we need to have a firmer grasp of what it means for a practice to promote or impede truth. The discussion in this article will be based on Goldman's account of these notions, to which I now turn.

2. Goldman's veritistic framework

There is a lot to be said about Goldman's veritistic framework. However, for the purposes of this paper, it suffices to rehearse the most essential ideas. One such basic element is the claim that states like knowledge, error, and ignorance have *fundamental* veritistic value or disvalue, whereas practices have *instrumental* veritistic value insofar as they promote or impede the acquisition of fundamental veritistic value. Another key ingredient in Goldman's theory is the question and interest relativity of veritistic value. An agent S's belief states are said to have value or disvalue when they are responses to a question that interests S. For the sake of simplicity, Goldman chooses to focus much of his discussion on yes-no-questions, i.e., questions of the kind "Is it the case that p?".³

Let us now turn to the very concept of veritistic value. Goldman's main proposal is that degrees of belief (DB) have veritistic value relative to a question Q, so that any DB in the true answer to Q has the same amount of V-value as the strength of the DB. In Goldman's terminology, $V\text{-value of } DB_X(\text{true}) = X$. Suppose, for example, that Mary is interested in the question whether it will rain tomorrow. If the strength of Mary's belief

³ This section is based on Goldman (1999), pp. 87-100. For criticism of Goldman's veritistic framework, see Maffie (2000) and Schmitt (2000). For Goldman's responses, see Goldman (2000).

that it will rain tomorrow is .8, and it will in fact rain tomorrow, then the V-value of Mary's state of belief vis-à-vis the rain issue is .8.⁴

As we saw, practices have instrumental veritistic value to the extent that they promote or impede the acquisition of states that have fundamental veritistic value. Suppose that a question begins to interest agent S at time t_1 , and S applies a certain practice π in order to answer the question. The practice might consist, for instance, in a certain perceptual investigation or in asking a friend. If the result of applying π is to increase the V-value of the belief states from t_1 to t_2 , then π deserves positive credit. If it lowers the V-value it deserves negative credit. If it does neither, it is neutral with respect to instrumental V-value.

The matter does not end here, however. In evaluating the V-value of a practice, we usually cannot focus merely on the one agent scenario. As Goldman notes, “[m]any social practices aim to disseminate information to multiple agents, and their success should be judged by their propensity to increase the V-value of many agents’ belief states, not just the belief states of a single agent” (1999, p. 93). This is why we should be interested in the *aggregate* level of knowledge, or true belief, of an entire community (or a subset thereof).

Goldman gives the following example. Consider a small community of four agents: S_1 - S_4 . Suppose that the question of interest is whether p or not-p is true, and that p is in

⁴ Goldman also mentions an alternative “trichotomous” model of V-value. Suppose S takes interest in the question whether p. The basic principles of this model are: If S believes the true proposition, the V-value is 1; if S rejects the true proposition, the V-value is 0; and if S withholds judgment, the V-value is .5. This alternative way of thinking about V-value will play no role in this article.

fact true. At time t_1 , the several agents have DBs vis-à-vis p as shown in the corresponding column (see Table 1). Practice π is then applied, with the result that the agents acquire new DBs vis-à-vis P at t_2 as shown in the column under t_2 .

	t_1	t_2
S_1	DB(p) = .40	DB(p) = .70
S_2	DB(p) = .70	DB(p) = .90
S_3	DB(p) = .90	DB(p) = .60
S_4	DB(p) = .20	DB(p) = .80

Table 1

At t_1 the group's mean DB in p is .55, so that .55 is their aggregate V-value at t_1 . At t_2 , the group's mean DB in p is .75, so that this is their new aggregate V-value. Thus the group displays an increase of .20 in its aggregate V-value. Hence the practice π displays positive V-value in this application.

A further complication is that there is a need to consider not just one application of a practice but many such applications. In evaluating a practice, we are interested in its performance across a wide range of applications. In order to determine the V-value of the

practice π in our example we would have to study how well it fares in other applications as well. This would presumably mean, among other things, varying the size of the population of inquirers as well as allowing it to operate on other initial degrees of belief. Once we have isolated the relevant set of applications against which the practice is to be measured, we can take its average performance as a measure of its V-value.

It follows from these considerations that, when assessing the V-value of a practice, we need to “average” twice. For each application A_i of the practice, we need to assess the average effect E_i it had on the degrees of belief of the members of the society. The V-value of the practice is then computed as the average over all the E_i s.

Having provided the essentials of Goldman’s theory, I now move on to what appears to be a serious problem for that approach. Indeed, the problem I will raise (of which Goldman himself is acutely aware) threatens the very idea of veritistic social epistemology because it sheds doubt on the notion that one could ever determine the veritistic value of interesting practices.

3. The determination problem

The problem of determining the veritistic value of a practice has two faces that I will choose to treat separately. I will refer to them as the *truth objection* and the *computational objection*. The truth objection is stated as follows by Goldman (1999, p. 91):

In defining the V-values of belief states and (derivatively) of practices, I assumed that the beliefs have objective truth-values. This assumption does not imply, however, that those truth-values are *known* to the veritistic theorist, or that they are easy to ascertain. A practice's V-properties are what they are, whether or not they are known to the theorist ... If they are not known, though, of what use are they? Why bother with such abstract definitions if V-performance cannot be determined?

Serious as these worries may seem, Goldman thinks that they do not after all present a fundamental threat to his epistemological enterprise (ibid.):

My measures of V-value are intended to provide conceptual clarity, to specify what is sought in an intellectually good practice, even if it is difficult to determine which practices in fact score high on these measures. Conceptual clarity about desiderata is often a good thing, no matter what hurdles one confronts in determining when those desiderata are fulfilled. An analogous situation is encountered in creating and filling positions in a business or organization. Clearly specifying the desired qualifications of a job-holder is highly desirable, however tricky it may be to identify an applicant who best satisfies those qualifications. Similarly, we want clear specifications of what it means for a practice to be V-valuable, however difficult it may be to identify the practices that actually exemplify this virtue.

I would like to challenge Goldman's central claim that "[c]onceptual clarity about desiderata is often a good thing, no matter what hurdles one confronts in determining when those desiderata are fulfilled".

Goldman's claim applies, in particular, to cases in which the hurdles confronted are so severe that it is practically impossible to determine when the desiderata are fulfilled. Let us zoom in on that special class of cases (to which our veritistic theorist's predicament is assumed to belong). If Goldman is right, achieving conceptual clarity about the desiderata is still valuable and worthwhile. But his business analogy points in another direction. Suppose it turns out to be impossible to determine whether a candidate satisfies the specified job qualifications. Then, surely, the time consumed specifying those qualifications will be seen, with hindsight, as time ill spent. To put it in plain terms: if your company wants to hire someone, and you come up with a list of qualification which is such that it cannot be determined whether a given candidate satisfies it or not, then you will be hearing from your boss very soon. Analogously, if it turns out to be impossible to determine the V-value of a practice, then the efforts invested in clarifying the concept of V-value were largely wasted.

I conclude that Goldman has little consolation to offer the veritistic theorist worrying about the extent to which the V-value of a practice can actually be determined. His business analogy works rather in the opposite direction of adding urgency to those concerns. Therefore, it remains crucial for the veritistic epistemologist to show how the V-performance of a practice can actually be determined in concrete cases. How can the V-performance be assessed and, in particular, how is this possible in situations in which the veritistic theorist cannot be assumed to know the true answer to the question at hand?

Suppose for the sake of the argument that the truth objection could be convincingly tackled. A computational problem would still remain. Goldman summarizes the computational difficulties involved in ascertaining the V-value of a practice in the following sobering words (1999, p. 91, my italics):

Veritistic social epistemology seeks to assess not only the practices currently employed by people and communities, but to inquire whether there might be better practices to replace those presently in use. This means that practices must be evaluated that, so far, have no track record at all. To evaluate such hitherto undeployed practices, one must consider how they would perform in a range of possible applications. In other words, we must consider their veritistic “propensities”, not just their veritistic “frequencies”. In fact the same point holds of practices that do have a prior track record. Whatever that track record is, it may be partly due to various accidental features, which are not firm guides to the future performance of the practice. Needless to say, it is not easy to determine the prospective performance of a practice. It cannot be determined by direct empirical observation, only by theoretical considerations, typically conjoined with background empirical information. *This makes the task of veritistic epistemology extremely difficult.*

It is noteworthy that Goldman does not mention the possibility of determining the V-performance of a practice by means of computer simulation. In section 5, I will argue that computer simulation is a promising technique in this regard. Before doing that, however, I will address the truth objection.

4. Addressing the truth objection

The truth objection can be summarized as follows:

- (1) Determining the veritistic value of a practice means determining its tendency to promote or impede truth.
- (2) The veritistic theorist cannot normally be assumed to know where the truth is in the applications she is studying.
- (3) Therefore: the veritistic theorist cannot normally determine the V-performance of a practice.

As we saw, Goldman's response was largely ineffective in calming this worry.

It seems to me, though, that the problem is to a large extent only apparent. The reason is that even though both premises in the above argument are true, the inference to the conclusion is not valid. In other words, even if the veritistic theorist has no idea whatsoever how to answer the question whether *p*, it doesn't follow that he or she cannot determine the veritistic value of a practice that aims at answering that type of question.

How can this be? The key observation is that many practices are such that their V-values plausibly do not depend on what answer to the underlying question is actually true. Suppose for example that you wonder whether it will be raining tomorrow, and that you decide to find out by asking your meteorologist friend. Clearly, the intellectual goodness in so doing is normally the same whether or not it will in fact rain. The probability that the meteorologist will say that it will rain, given that it will, equals the probability that she will say that it won't rain, given that it won't. In other words, the

reliability of your friend does not depend on whether or not it will actually rain. The same is true *mutatis mutandis* for almost all practices that readily come to mind. They are all what I will call *truth invariant*. The fact that most practices one can think of are truth invariant has the important consequence that the determination of their V-performance does not presuppose, on the part of the verististic theorist, any knowledge of the true answer to the underlying issue.

It is, to be sure, possible to come up with practices that do not satisfy truth invariance. Suppose our meteorologist friend is unusual in the respect that she will always say that it will rain, regardless of whether it will. Then asking her will be an excellent thing to do, if it will rain. For in that case she is, in a trivial sense, completely reliable. However, it will not be such a good thing to do, if it won't rain, in which case our friend will give us false information. This practice is truth *variant* as opposed to truth *invariant*.

Nevertheless, almost all practices that have been seriously discussed by epistemologists are plausibly truth invariant: trusting other people, asking a friend, lying, telling the truth, conducting inquiry, thinking, reasoning. What is common to these practices is that they have a certain degree of generality which other practices lack: they don't refer to particular persons, for instance. Maybe this is what made them epistemologically interesting in the first place.

Although I have already given my basic response to the truth objection, it might be complained that I have underestimated the frequency of interesting practices that are truth variant. There are more realistic examples in which we apply a method which yields a higher probability of a correct answer if that answer is 'yes' than if it is 'no', or vice versa. For example, many tests of diseases have this property. A test may have high

sensitivity (high probability of a correct answer if that answer is ‘yes’, i.e., low probability of a type I error), but low specificity (relatively low probability of a correct answer if that answer is ‘no’; i.e. a relatively high probability of a type II error). Or the other way round: high specificity and low sensitivity. That such asymmetries are widely present in testing methods shows that the truth objection is more serious than I seem to suggest.⁵

While I admit that medical tests are often of the kind in question, these practices are not of the kind which has attracted most interest in the social epistemology community. Practices like “relying on testimony”, “telling the truth” and “lying” have, and they are also plausibly truth invariant. Even so, it would be useful to be able to determine the V-value of truth variant practices as well. Fortunately, there is a rather obvious proposal for how this could be achieved.

Let π be the practice under consideration, so that we want to determine the veritistic value of π . Let $V\text{-value}(\pi \mid C)$ stand for “the veritistic value of practice π given condition C ”. My suggestion is that we proceed as follows. We assess both the V-value of π on the

⁵ Another example of an asymmetric method is statistical hypothesis testing. The hypothesis to be tested is always negative. It is hypothesized that there is no connection between variables. If the hypothesis cannot be rejected (if there is no statistically significant connection in our data), then it is accepted. This means that statistical testing is asymmetric; it treats hypotheses about the absence of connections as true, as long as they have not been shown to be false. Consequently, the outcome of statistical testing seems more reliable, when it leads to the claims about the existence of connections, than when it leads to ‘no-connection’ claims. I am indebted to (* name removed for purposes of anonymous refereeing *) for challenging me on the issue of truth variance.

assumption that p is true and also the V-value of π on the assumption that not- p is true; and, finally, we take the average of those two V-values. In other words,

$$V\text{-value}(\pi) = \frac{1}{2} (V\text{-value}(\pi \mid p \text{ is true}) + V\text{-value}(\pi \mid \text{not-}p \text{ is true}))$$

This should do the job when p and not- p are equally probable. If p is the proposition that the person being tested has the disease, this means that that person is as likely to have as not to have the disease. The general situation is slightly more complex:

$$(EV) \quad V\text{-value}(\pi) = P(p) V\text{-value}(\pi \mid p \text{ is true}) + P(\text{not-}p) V\text{-value}(\pi \mid \text{not-}p \text{ is true})$$

In other words, absent knowledge about the truth value of p we need to focus on the *expected* veritistic value of the practice in question.

Plausible as it may seem, (EV) raises a question that we need to take seriously. In order to compute the V-value of a practice using equation (EV) we need to assess the probability of the proposition p under consideration. If that probability is not frequency-based but subjective, which would be the case if we do not know how frequent a given disease is in the population, then what we get is the *subjective* expected veritistic value, and it is not clear that this would be acceptable from Goldman's externalist perspective.

Interestingly, however, medical cases are not only cases in which we find many truth variant tests and practices; they are also cases in which there is a lot of reliable information available about frequencies. There are, for instance, good estimates of how many people are infected with HIV in various populations. Thus, according to a recent

estimate, 0.3 percent of the adult (15-49) population in West and Central Europe are infected.⁶ This means that the veritistic value of the practice of applying a new test for early detection of HIV could in principle be assessed using equation (V) in a manner that should be acceptable even to objectivists and externalists. To take another example, the frequency of Parkinson's disease among elderly is known to be approximately one percent, a fact that could be used for assessing the veritistic value of a Parkinson detection method.

5. Addressing the computation objection: the Laputa simulation framework

The computational problem has its root in the fact that assessing the veritistic value means collecting and processing tremendous amounts of information about various applications of the practice in question, including applications that have not actually been realized but are only possible. Processing large amounts of data is precisely what we use computers for, suggesting that the computational problems could be solved by means of computer simulation. In this section I hope to make likely that simulation greatly simplifies the computational task of veritistic social epistemology to the point of making that task relatively easy once the appropriate software has been developed. Developing the software is a non-trivial task from a practical perspective, but it can be done and indeed it has been done, at least on the prototype stage. I will now present the recently

⁶ Source: United Nations (<http://www.unaids.org/en/>).

developed Laputa simulation environment that is capable of computing (approximate) V-values of interesting practices.⁷

A basic notion in Laputa is that of a *social network* in which people can communicate with each other. Social networks are represented as graphs in which the nodes represent inquirers and the links represent communication channels. The links are directed, allowing for one-way communication (see Figure 1).

Figure 1: The social network of Sherlock Holmes represented in Laputa.

⁷ The name Laputa derives from Jonathan Swift's *Gulliver's travels*. Laputa is the result of joint work by Staffan Angere and the author. It is a work in progress that is still being improved and extended. There have been several other attempts to shed light on issues in social epistemology by means of computer simulation. For an influential example, see Hegselmann and Krause (2006) and for a discussion Olsson (2008). See also Zollman (2007) for a Bayesian simulation approach to issues in philosophy of science. Volume 6, issue 2, of the journal *Episteme* is devoted to simulation in social epistemology (guest editor: Igor Douven). Laputa seems to be the only model so far that computes veritistic values in Goldman's sense.

Following Goldman, it is assumed that all inquirers focus on answering one and the same question: whether p or not- p . For example, p can be the proposition “Professor Moriarty committed the crime”, “The economic crises will soon be over” or “John suffers from Parkinson’s disease”. A number of parameters can be set for each inquirer. The *initial degree of belief* is an inquirer’s degree of belief in p from the start. *Inquiry accuracy* is the reliability of the inquirer’s own inquiries. The *inquiry chance* is the probability that the inquirer will conduct an inquiry. The *inquiry trust* is the inquirer’s degree of “self-trust”, i.e., her degree of trust in her own inquiries. Likewise, there are a number of parameters for each link. The *listen chance* is the probability that the recipient will listen to a message she receives. The *listen trust* is the recipients trust in the sender. The *threshold of assertion* is the degree of confidence in a proposition (“ p ” or “not- p ”) required for the sender to submit a corresponding message to the recipient(s). For instances, if the threshold is set at .90, this means that the sender needs to believe p (not- p) to a degree .90 in order for her to “assert” p (not- p) in the network.

The current version of Laputa computes the veritistic value of social practices that are truth invariant, so that the V-value of the practice does not depend on what proposition, p or not- p , is actually true. This means that we can take one of p or not- p to be true by convention. In Laputa it is stipulated that p is true.⁸

⁸ We plan to implement equation (EV) of section 4 in later version of Laputa allowing the system to compute the veritistic value of practises that are truth variant as well. This can actually be done already in the current version, although the process is somewhat cumbersome.

Running Laputa can mean to construct a network such as that in Figure 1, assign initial values to the inquirer and link parameters, and then click on a “run” button. What happens then is that Laputa runs through a series of steps, each step representing a chance for an inquirer to conduct an inquiry, to communicate (send, listen) to the other inquirers to which she is “hooked up”, or to do both. After each step, Laputa will update the whole network according to the information received by the inquirers. This is done in accordance with standard Bayesian techniques. Thus, a new degree of belief is computed for each inquirer based on the old degree of belief and the new information received through inquiry and/or listening to other inquirers. Laputa also updates the inquiry trust and listen trust parameters, which is once more accomplished in accordance with Bayesian principles. Figure 2 shows the output of running the network shown in Figure 1 two steps.

Time: 1

Inquirer 'Sherlock Holmes' heard that not-p from inquirer 'Mycroft Holmes', lowering his/her expected trust in the source from 0.189 to 0.188.

This raised his/her degree of belief in p from 0.52297 to 0.82427.

Inquirer 'Inspector Lestrade' received the result that not-p from inquiry, lowering his/her expected trust in it from 0.500 to 0.355.

This lowered his/her degree of belief in p from 0.93375 to 0.91253

Avg. error = 0.417, error dev. = 0.272, error delta = -0.060, avg. trust = 0.461, trust delta = -0.000.

Time: 2

Inquirer 'Mycroft Holmes' received the result that p from inquiry, lowering his/her expected trust in it from 0.396 to 0.316.

This lowered his/her degree of belief in p from 0.18691 to 0.13081

Inquirer 'Mrs Hudson' heard that p from inquirer 'Sherlock Holmes', lowering his/her expected trust in the source from 0.634 to 0.628.

This raised his/her degree of belief in p from 0.41000 to 0.54616.

Inquirer 'Dr Watson' received the result that not-p from inquiry, lowering his/her expected trust in it from 0.500 to 0.480.

This lowered his/her degree of belief in p from 0.56000 to 0.52256

Avg. error = 0.401, error dev. = 0.278, error delta = -0.076, avg. trust = 0.460, trust delta = -0.001.

Figure 2: Example of simulation output for the network in Figure 1.

By the *network structure* we shall mean the graph structure of the network, i.e. its nodes and links. A *network state* is a network structure together with values for all parameters for the inquirers and links. A *network evolution* is a series of network states resulting from running Laputa.

As seen from Figure 2, Laputa outputs not just what happens to the individual inquirers during simulation, but also collects some statistical data. For our purposes, “error delta” is of special interest. Error delta is the difference between the initial and final average degrees of belief in the true proposition p. Given error delta, we can

compute the veritistic value for a network evolution according to the following simple rule: $V\text{-value} = -\text{error delta}$. This means that an error delta of -0.076 equals a V-value of 0.076.

However, the veritistic value of a *network evolution* was obviously not what we were looking for. We wanted to assess the veritistic value of a *practice*. So how do we get from V-values of network evolutions to V-values of practices? The first thing to note is that what we have learned about Laputa so far allows us to study the V-value of a particular *application* of a practice. Consider for instance the practice of trusting other people. Before we run the network we can adjust the listen trust parameter for all the links so that this condition is satisfied. Now we run the network as previously described, preferably until the network stabilizes and relatively fixed degrees of belief have been obtained. What we get as a result is the V-value of the practice of trusting other people as applied to the particular network at hand and its initial state (e.g. the Sherlock Holmes network of Figure 1).

We still want to know, however, how to get from the V-value of a practice as applied to a particular network to the V-value of the practice itself. Laputa solves this problem by allowing its user to specify various features or “desiderata” of networks at an abstract level. The program can then randomly generate a large number of networks, of different sizes, having those features, letting them evolve, collecting the corresponding V-values and, finally, outputting the average V-value of all the network evolutions it has examined. This allows Laputa to compute the V-value of a large number of interesting practices. For instance, Laputa can be told, at the abstract level, to study 10 000 randomly generated networks in which inquirers trust each other to a certain degree. The resulting V-value is

a measure of the V-value of the practice of trusting other people itself, independently of any particular network. All this is done in Laputa's "batch window" (Figure 3).

Figure 3: The batch window in Laputa.

In the batch window, various probability distributions can be selected for the several inquirer and link parameters. For instance, the flat distribution for "Starting belief" indicates that Laputa, when selecting the initial degrees of belief for a generated network, will treat all possible degrees of belief as being equally likely to be realized. The selection of a normal distribution for "Inquiry accuracy", centered around 0.75 means that Laputa, when selecting the inquiry accuracy for the inquirers in the generated networks, will have a preference for assigning an accuracy of 0.75 and surrounding values. The population feature allows the specification of the lower and upper sizes of the networks to be examined. In this case, Laputa is instructed to generate and study

networks having 2 to 20 inquirers. “Link chance” specifies the “density” of the networks to be studied. A link chance of 0.25 indicates a 25 percent chance that two inquirers will be connected by a communication link. In Figure 3, the number of trials have been set to 1,000, meaning that Laputa will generate and study 1,000 networks in accordance with the statistical criteria specified in the batch window. Finally, the number of steps per trial has been set to 100, indicating that the focus is on the long-term effects of implementing the practice.

Apart from allowing the veritistic value of practices to be determined, the development of Laputa had two conceptual side-effects that are worth mentioning. One has already been alluded to: the possibility in Laputa of differentiating between the short run and long run V-performance of a practice. Suppose for example that we want to know how beneficial truth telling is in the long run. This problem could be studied by setting the number of steps per trial to, say, 100. If we are more interested in short term uses, we could instead set the number of steps to a smaller number, say, 5 or 10.

Secondly, Laputa can also help us to get clearer on what a social practice is. What do intuitively interesting social practices (like blind trust in others, free speech, telling the truth, and so on) have in common? From the point of view of Laputa, the answer is that they are all constraints on network states. This suggests *identifying a social practice with a network constraint*. Any such constraint which can be imposed in Laputa’s batch window (and there are a great many of those) can be studied from the point of view of veritistic value. Given the proper directive, Laputa will generate a great number of networks (“societies”) satisfying the constraints, allow them to evolve and, finally, output the corresponding veritistic value. This includes constraints that would perhaps not

normally be described as social practices, e.g., “being reliable in one’s inquiries to degree .75”. Nevertheless, identifying a social practice with a network constraint may still be a fruitful explication, in the sense of Carnap (1950), of the concept of a social practice.⁹

6. Putting Laputa to the test

My main point is that the method of computer simulation can be used to overcome the computational obstacle for veritistic social epistemology which Goldman has identified. It is possible to design and implement a computer program that generates and checks a wide range of applications of practices so as to determine the V-value of the practice as a socially aggregated average over those applications. In support of this contention, I have described a running simulation program, Laputa, which does precisely this. This shows that there is no problem *in principle* of computing V-values. This is the main issue here, and I take what I already said as sufficient evidence to regard it as largely settled.

For the record, there is still the question of whether our particular program, Laputa, computes V-values *correctly*. This is of course an entirely different ball game and

⁹ It could be objected that practices should by their very nature be something that can be voluntarily implemented, at least to some degree, and that this would allow “putting trust in others” but disqualify “being reliable in one’s inquiries”. However, even in the latter case we can exert some degree of voluntary control. By trying harder and being more attentive we can become more reliable as inquirers. An alternative to the present proposal would be to identify a practice with a constraint not on a network state but on a network evolution. These two alternatives need not be mutually exclusive. Perhaps some practices are most naturally thought of as constraints on a network, while others are better conceived of as constraints on how a network is allowed to evolve.

depends on many things that are not directly relevant to the philosophical issue raised by Goldman. In computing the veritistic value of practices, Laputa relies on a particular way of updating degrees of beliefs and trust in a network. As we saw, the basic updating mechanism is standard Bayesian conditionalization, a technique for belief updating that has a reasonably firm standing in the philosophical and scientific communities. But this does not mean, of course, that one couldn't imagine other formal frameworks for updating degrees of belief. Moreover, when actually implementing a design like Laputa, a lot of minor decisions have to be made which can potentially affect the outcome. (As always, the devil is in the details.) There could also be programming errors compromising the output. The bottom line is that the actual implementation of a design like Laputa can itself be a source of error and controversy.

This is not the place to attempt a large-scale validation of Laputa as a tool for computing veritistic value. However, I will make a small-scale, casual attempt to confirm Laputa in the context of a set of test cases in which we have reasonably clear antecedent expectations concerning what the result should be.

Test case 1: "Nothing comes from nothing". We would expect that, unless some inquirer is reliable, no practice can have non-zero V-value. If all people in the social network are completely unreliable (in the sense of "randomizing"), it doesn't matter if people communicate, trust each other, and so on. Nothing will come out of it.

Confirmation: We can test this prediction in Laputa by setting “Inquiry accuracy” to 0.5 in the batch window and letting the other parameters vary randomly. The result is, as expected, zero V-value.¹⁰

Test case 2: “Nothing comes from *perceived* nothing”. This is a variation on the previous test case. Unless some inquirer not only *is* a reliable inquirer but also *treats* her own inquiries as reliable, no practice can have a positive V-value.

Confirmation: Set inquiry accuracy to some value over 0.5, to secure inquiry accuracy, but inquiry trust to 0.5. The effect, once again, is zero V-value regardless of what other settings are made.¹¹

Test case 3: “Other things being equal, more reliability is always a good thing”. We would expect a higher degree of inquiry accuracy to be beneficial in a *ceteris paribus* sense. If people are generally more reliable in their own inquiries, that should benefit society at large.

Confirmation: We can test for this property by gradually increasing reliability accuracy in Laputa’s batch window, while keeping everything else the same. As noted before we

¹⁰ This principle should be valid in normal circumstances without being universally valid. The sociological *law of group polarization* states that “members of a deliberating group predictably move toward a more extreme point in the direction indicated by the members’ preliberation tendencies” (Sunstein, 2002, p.176). This is so even if no additional inquiry takes place during deliberation (and hence even if the inquirers are entirely unreliable in their own investigations). Hence, if the members’ preliberation tendencies is to think that p is more likely than not, they will move toward believing fully that p is the case. If p is true this means that the deliberation process had positive V-value. It can be verified that Laputa covers this exception as well.

¹¹ Keith Lehrer should be credited for emphasizing the epistemological importance of “self-trust”. See, for instance, Lehrer (1997). For criticism of Lehrer’s theory and Lehrer’s responses, see Olsson (2003).

must keep inquiry trust positive in order to register any (positive) effect at all. As expected a higher V-value is obtained for higher values of reliability accuracy.

Test case 4: “Everything else equal, truth telling is better than lying”. It is better for society that people tell the truth than that they lie. Thus, the practice of telling the truth should receive a higher V-value than the practice of lying, at least in a *ceteris paribus* sense. This is an issue that is slightly more intricate than the others and it therefore requires a somewhat more extended treatment.

First we need to get clearer on what “truth telling” and “lying” mean. A truth teller can be someone who tells what is actually the truth. But this not what is usually meant by truth telling. Rather, a person is a truth-teller if she says what she *takes* to be the truth. Similarly, a person is a liar if she says what she takes to be false. At least, these are arguably the epistemologically more interesting senses of “truth telling” and “lying”. To fix ideas, we will think of truth tellers and liars in the following terms:

(A) A truth teller is someone who says that p (not-p) just in case her degree of belief in p (not-p) exceeds 0.9.

(B) A liar is someone who says that p (not-p) just in case her degree of belief in p (not-p) falls below 0.1.

A threshold of assertion to a value α below 0.5 is interpreted by Laputa as a “liar threshold”, i.e., the inquirer will say that p (not-p) when the degree of belief for p (not-p) is below α .¹²

¹² The most obvious choice would be to think of a truth teller as someone who says that p (not-p) just in case her degree of belief in p (not-p) equals 1; and of a liar as someone who says that p (not-p) just in case

We are now in a position to confirm Laputa by checking that truth telling indeed is a V-better practice than lying. We will check this against the background of some normalcy assumptions. As usual we need to have in place positive inquiry accuracy as well as a positive inquiry trust. We will also assume an initial positive listening trust, i.e., that people are initially somewhat inclined to rely on “the word of others”. To be specific, we will assume these factors to be normally distributed around 0.75. Laputa now gives the following output in the long run (100 steps per trial, 95 % confidence):

$$\text{V-value of truth telling} = 0.292 \pm 0.016$$

$$\text{V-value of lying} = 0.098 \pm 0.010$$

As expected, the V-value of truth telling is higher, indeed much higher, than the V-value of lying. The V-value of truth telling is actually rather impressive considering the fact that the maximum V-value that can be attained is 0.5 (as the reader can easily verify). What is somewhat surprising is that the V-value of lying is still positive: on the average, the practice of lying did not cause damage to society, although it would have been better had people been telling the truth instead. The explanation turns out to be straightforward. Remember that Laputa not only updates the inquirers’ degrees of belief, it also updates their trust in the word of their peers. What will happen is that liars tend to become, as it were, identified as such in the long run, i.e., the trust in them will become “negative”, meaning that the listener will tend to take a message to the effect that p (not-p) is true as

her degree of belief in p (not-p) equals 0. However, this is alien to the Bayesian approach to belief updating according to which 0 and 1 are special degrees of belief that should not be offhandedly assigned.

evidence not for p (not- p) but for not- p (p).¹³ There is something to be learned even from a liar, provided you have figured out that she is lying systematically.¹⁴ Still, telling the truth from the start is a more efficient policy because it does not involve the potentially tedious process of gradually downgrading one's initial trust.

7. Conclusions

In his important 1999 book, Goldman raises a number of serious worries for his account of veritistic value. Two of them concern the possibility of determining the veritistic value of a practice in a concrete case because we often don't know what beliefs are actually true, and even if we did there would still remain a computational problem to address. These objections challenge just about any account of veritistic social epistemology and not just Goldman's specific theory.

I observed that that the first problem pertains only to practices that are truth variant so that the V-value depends on what is actually true. Fortunately, most practices that have received epistemological attention (e.g. relying on testimony) do not belong to that class but are rather truth *invariant*. However, the class of truth variant practices cannot be ignored entirely, and I suggested a way of assessing the V-value of such practices by

¹³ This process of revising the trust can be studied by examining the log reports that Laputa is capable of producing. See Figure 2 for an example.

¹⁴ This is so if, as in the basic scenario studied by Goldman, the underlying question is of a yes-no kind. In cases in which the inquirers face a question with more than two possible answers lying is less informative because there are then many different ways to lie.

determining their objective expected V-value. Paradigm cases of truth variant practices from medical science often allow for such determination.

The main part of the paper was devoted to the computational problem which, I argued, can be addressed in promising terms by means of computer simulation. In concrete support of this contention, I described the Laputa simulation environment which allows a large number of social networks structures to be generated and studied from a veritistic point of view. By specifying constraints on those networks, the user can let Laputa compute the veritistic value of various interesting practices. I concluded that the computational issue raised by Goldman does not in fact pose a fundamental threat to the project of veritistic social epistemology.

Another matter entirely is to what extent the output of Laputa, or any other simulation program, is reasonable. There is clearly many ways in which the details of such a program could be specified and implemented. Laputa represents only one possibility in this regard, although one which should be acceptable, at least in principle, to many researchers working in the influential Bayesian tradition. In the penultimate section, I made an attempt to test Laputa against the background of our intuitive expectations in some simple cases, which led to some rather encouraging, if elementary, results.

Acknowledgements

I am greatly indebted to Alvin Goldman, Klemens Kappel and the members of Klemens's Copenhagen-based epistemology group, for their advice and – not the least – patience during the development of the ideas that I am putting forward in this paper. I am

particularly grateful to Staffan Angere for co-developing the probabilistic model underlying Laputa with me and for implementing it in computer code.

References

- Carnap, R. (1950), *Logical Foundations of Probability*, University of Chicago Press, Chicago.
- Frankfurt, H. G. (2005), *On Bullshit*, Princeton University Press, Princeton.
- Goldman, A. I. (1986), *Epistemology and Cognition*, Harvard University Press, Cambridge Mass.
- Goldman, A. I. (1999), *Knowledge in a Social World*, Clarendon Press, Oxford.
- Goldman, A. I. (2000), “Replies to Reviews of *Knowledge in a Social World*”, *Social Epistemology* 14 (4): 317-333.
- Goldman, A. I., and Olsson, E. J. (2009), “Reliabilism and the Value of Knowledge”, in *Epistemic Value*, Haddock, A. et al (Eds.), Oxford University Press.
- Hegselmann, R., and Krause, U. (2006), “Truth and Cognitive Division of Labour: First Steps Towards a Computer-Aided Social Epistemology”, *Journal of Artificial Societies and Social Simulation* 9 (3).
- Lehrer, K. (1997), *Self-Trust*, Clarendon Press, Oxford.
- Maffie, J. (2000), “Alternative Epistemologies and the Value of Truth”, *Social Epistemology* 14 (4): 247-257.
- Olsson, E. J. (2003), *The Epistemology of Keith Lehrer*, Philosophical Studies Series 95, Kluwer Academic Publishers, Dordrecht.

-
- Olsson, E. J. (2008), “Knowledge, Truth, and Bullshit: Reflections on Frankfurt”,
Midwest Studies in Philosophy XXXII: 94-110.
- Schmitt, F. F. (2000), “Veritistic Value”, *Social Epistemology* 14 (4): 259-280.
- Sunstein, C. R. (2002), “The Law of Group Polarization”, *The Journal of Political Philosophy* 10 (2): 175-195.
- Zollman, K. J. (2007), “The Communication Structure of Epistemic Communities”,
Philosophy of Science 74 (5): 574-587.