LUND UNIVERSITY

Norms of Assertion and Communication in Social Networks

Olsson, Erik J; Vallinder, Aron

Published in:
Synthese

DOI:
[10.1007/s11229-013-0313-1](https://doi.org/10.1007/s11229-013-0313-1)

2013

[Link to publication](#)

Citation for published version (APA):

Olsson, E. J., & Vallinder, A. (2013). Norms of Assertion and Communication in Social Networks. *Synthese*, 190(13), 2557-2571. <https://doi.org/10.1007/s11229-013-0313-1>

Total number of authors:
2

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

Norms of Assertion and Communication in Social Networks

Erik J. Olsson and Aron Vallinder

Abstract: Epistemologists can be divided into two camps: those who think that nothing short of certainty or (subjective) probability 1 can warrant assertion and those who disagree with this claim. This paper addressed this issue by inquiring into the problem of setting the probability threshold required for assertion in such a way that the social epistemic good is maximized, where the latter is taken to be the veritistic value in the sense of Goldman (1999). We provide a Bayesian model of a test case involving a community of inquirers in a social network engaged in group deliberation regarding the truth or falsity of a proposition p . Results obtained by means of computer simulation indicate that the certainty rule is optimal in the limit of inquiry and communication but that a lower threshold is preferable in less idealized cases.

1. Introduction

There are two camps of epistemologists: those who think that nothing short of subjective probability 1 can warrant assertion, and those who think this is untrue and that a person can be warranted in asserting a proposition even if her subjective probability in that proposition is less than 1. In the first camp we typically find those who subscribe to the knowledge rule of assertion. A number of distinguished epistemologists have evaluated that rule favorably, including Keith DeRose (2002), John Hawthorne (2004), Jason Stanley (2005) and Timothy Williamson (2000). As Williamson puts it, “[o]nly knowledge warrants assertion” (2000, p. 243). To the extent that an advocate of the knowledge rule also takes the standard position that knowledge requires (full) belief, he or she will be committed to the view that nothing short of probability 1 can warrant assertion.¹ Those advocating a truth or (full)

---

¹ For some proponents, such as Williamson (2000), the knowledge rule does not link any specific psychological state, such as belief, to assertion.

belief rule of assertion also belong to this group (e.g. Weiner, 2005). Members of the second camp typically promote a "rational credibility" or "reasonable belief" rule of assertion. Proponents include Igor Douven, Frank Jackson, Jennifer Lackey and David Lewis. According to David Lewis (1976, p. 297) for example, "[t]he truthful speaker is willing to assert only what he takes to be very probably true". In a similar vein, Jackson (1979, p. 565) says that a proposition "is assertable to the extent that it has high subjective probability for its assertor".

Douven (2006, p. 481) argues that the debate has led to a standoff between these two camps in terms of the number of prominent researchers belonging to each camp and in terms of how well the two camps deal with objections and problems. This may lead one to conclude that there is no fact of the matter, i.e. no objectively valid norm of assertion. This would be similar to how people have argued from there being longstanding and wide-spread disagreement in moral affairs to their being no objectively valid ethical standard. Alternatively, one might infer that the traditional philosophical resources of reflection and intuition-reliance may have been exhausted, and that it is time to look for new approaches and new sources of evidence.

The project to be carried out in this paper is more in line with the second option. First, rather than viewing the problem of setting the threshold of assertion as basically a problem of individualistic epistemology, we will conceive it as a problem in social epistemology. Our fundamental assumption is captured by the following Epistemic Principle of Assertion Rule Validity:

(EPARV) An assertion rule R is valid for a given community C to the extent that R , if universally followed by members of C , maximizes social epistemic good.

Thus we will ask the somewhat Kantian sounding question: what choice of a threshold of assertion would, if implemented by everyone, lead to the best epistemic result for society at large? Second, rather than discussing assertion rules in the abstract, as philosophers usually do, we will focus on a concrete test case: setting the threshold for assertion for a community of inquirers in a social network engaged in group communication concerning the question whether a particular proposition p is the true. We will accordingly ask whether the best option is to set the threshold to 1 or whether a lower

value should rather be chosen. This way of posing the problem corresponds to the philosophical dispute we have just surveyed and any results we obtain will bear on that debate. We will depart from tradition, finally, in asking a question that is slightly more general than the usual one. What our investigations will focus on is identifying the optimal threshold in any given case. Hence, if we find that a threshold below 1 is sometimes to be preferred, we will not settle for this answer but we will continue asking what the optimal threshold value is.²

A rigorous treatment of these matters requires that we model group communication in social networks in formal (in our case probabilistic) terms. Thus, we need to specify, for instance, what we mean by saying that one assertion rule is socially optimal from an epistemic standpoint; and we need some way of determining the goodness of an assertion rule under certain specified assumptions, such as the rule "Assert something only if you confidence is above 0.9". Given the model we are about to present, this could in principle be done mathematically. In practice, however, this is far too tedious.

In section 2, we introduce the concept of veritistic value as a plausible more specific way of understanding the social epistemic goodness of a practice. In section 3, we present our Bayesian social network model. In section 4, we investigate the problem of finding the optimal threshold of assertion using a simulation program that has been devised for the purpose of making mathematical investigations into our favored Bayesian model tractable. We also test the robustness of our results by varying the circumstances under which they were obtained. The last section of the paper is devoted to the question of how best to explain the outcome of our simulation experiments.

2. Veritistic value

---

² We focus in this paper on identifying the isolated impact of specific rules of assertion. A problem in this regard concerns the background against which such an investigation takes place. In general, there may be other rules that are used by a given community and those other rules may affect the consequences of a new rule under consideration. These are broad issues that we cannot address extensively in this paper. Nevertheless, the stability tests that we perform in section 4 serve to address some of these concerns.

We will here give the essentials of Alvin I. Goldman's theory of veritistic value, which we will rely on as a plausible and independently supported way of measuring the goodness of a social practice, such as the practice resulting from the universal implementation of an assertion rule.

According to Goldman, 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

Goldman's basic concept is that of *a person's degree of belief (DB) having veritistic value (V-value) relative to a question Q*. The idea is simply that the V-value of person's DB in the true answer to a question *Q* equals 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 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

This takes care of the V-value of a person's degree of belief. Our primary concern, however, is of course to define the V-value of a social practice. So how do we get from individual V-values to the V-values of practices? 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, for instance, 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 ,

---

³ This section is based on Goldman (1999), pp. 87-100, and the exposition in Olsson (2011).

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.

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).

To fix ideas, 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 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. In that sense, the V-value of a practice is its *expected* V-value, but for simplicity we will often simply call it V-value.

We have now reached the point at which we can state what we will refer to as Goldman’s fundamental Principle of Veritistic Value:

(PVV) The social epistemic goodness of a practice is determined by the degree to which the practice maximizes veritistic value.

From EPARV and PVV we may infer a Veritistic Principle of Assertion Rule Validity:

(VPARV) An assertion rule R is valid for a given community C to the extent that R , if universally followed by members of C , maximizes veritistic value.

In other words, the rule of asserting something only if one’s degree of belief has reached a certain threshold, if implemented, is a social practice. This being the case, we may ask what veritistic value it has. To what degree can a particular threshold of assertion, if generally implemented, be expected to steer the average member in direction of adopting a higher credence in the true answer to the underlying question? The answer to this question will in accordance with VPARV be informative concerning the validity of that particular way of setting the threshold.

3. A Bayesian social network model

As we stated in the introduction, we wish to evaluate the epistemic, or veritistic value, of various implemented rules of assertion in a concrete test case. For that purpose we envisage a social network of inquirers engaged in joint inquiry into the question whether a given proposition, p , is the case. The inquirer can, at any point in time, engage in their own inquiries (ask an external source) and/or assert their current view to their peers in the network. Formally, we can take a social network S to be a set Γ , which we for reasons that will be made obvious shortly will call the set of *inquirers*, together with a binary relation R on Γ , which we call the *network structure*. This means that a social network is a directed graph.⁵

Following Bayesian tradition, the epistemic state of a person α at time t is assumed to be given by a *credence function* $C_\alpha^t : L \rightarrow [0,1]$ instead of just a set of sentences. L can be taken to be a classical propositional language, and C_α^t is assumed to fulfill the standard axioms of a probability measure. Due to its close connection to probability theory, Bayesianism is well suited for statistical models, and as the models that we will investigate in this paper are of this kind, we will adopt the Bayesian approach.

For the purposes of this paper, let us confine ourselves to the case where inquiry is aimed at discovering whether a single proposition p is true or false. We assume, conventionally, that p happens to be true, since this will simplify calculations further on. Every inquirer will then have a credence $C_\alpha^t(p)$ in p , which is a real number between 0 and 1, for every moment t .

In our model, there are two fundamentally different ways for the inquirers to receive new information: inquiry and communication. Inquiry can here be taken to include any kind of method of altering a credence function which does not base itself on information given by anyone else in the network. The paradigmatic cases of inquiry are observation, experiment, and perhaps taking advice from persons outside the social network S .

Not all participants' approaches to inquiry are the same, and they tend to vary both in their degree of activity and their effectiveness. Let $S_{i\alpha}^t p$ be the proposition “ α 's inquiry gives the result that p at time t ”, $S_{i\alpha}^t \neg p$ be the proposition “ α 's inquiry gives the result that not- p at t ”, and $S_{i\alpha}^t p \vee S_{i\alpha}^t \neg p$ the

---

⁵ This section is based on the account in Angere (forthcoming).

proposition that α 's inquiry gives *some* result at t . We represent the participants' properties *qua* inquirers by two probabilities: the chance $P(S_{i\alpha}^t)$ that, at any moment t , α receives a result from her inquiries, and the chance $P(S_{i\alpha}^t p | S_{i\alpha}^t \wedge p)$ that, when such a result is obtained, it is the right one. To simplify matters, we assume that the chance that inquiry gives a true result does not depend on whether p is true or false.

$P(S_{i\alpha}^t)$ will be referred to as α 's *activity*, and $P(S_{i\alpha}^t p | S_{i\alpha}^t \wedge p)$ as her *aptitude*. An inquirer without interest in p would generally have a low activity value, while one very interested in p , but engaging in inquiry using faulty methods would have a high activity value but an aptitude close to 0.5, or even below that. In the latter case, the results of her inquiry would actually be negatively correlated with the truth. As a simplification, we will assume α 's activity and aptitude to be constant over time, so we will generally write them without the time index t .

Just as inquiry represents the flow of information into the network, communication deals with how this information is disseminated. Analogously to the inquiry notation we define

$$S_{\beta\alpha}^t p =_{df} \beta \text{ says that } p \text{ to } \alpha \text{ at } t$$

$$S_{\beta\alpha}^t \neg p =_{df} \beta \text{ says that } \neg p \text{ to } \alpha \text{ at } t$$

$$S_{\beta\alpha}^t =_{df} \beta \text{ says that } p \text{ or that not-}p \text{ to } \alpha \text{ at } t$$

This strength of a link $\beta\alpha$ is then representable as a probability $P(S_{\beta\alpha})$, being the chance that β communicates that p or that not- p to α , at any given moment t .

Given that β communicates with α , what does she say? And what makes her say it? These questions are answered by a property of the link $\beta\alpha$ that we will call its *threshold of assertion* or just threshold for short: a value $T_{\beta\alpha}$ between 0 and 1, such that

If $T_{\beta\alpha} > 0.5$, β tells α that p only if $C_\beta(p) \geq T_{\beta\alpha}$, and that not- p only if $C_\beta(p) \leq 1 - T_{\beta\alpha}$;

If $T_{\beta\alpha} < 0.5$, β tells α that p only if $C_\beta(p) \leq T_{\beta\alpha}$, and that not- p only if $C_\beta(p) \geq 1 - T_{\beta\alpha}$; and

If $T_{\beta\alpha} = 0.5$, β can tell α that p or that not- p independently of what she believes, which is modeled by letting her pick what to say randomly.

So far we have described how the inquirers in a social network engage in inquiry and communication, but we have said nothing about how they react to the results of these practices. This is the purpose of the following considerations.

We define α 's source σ 's *reliability* as

$$R_{\sigma\alpha} \text{--}df\ P(S_{\sigma\alpha}p|S_{\sigma\alpha} \wedge p) = P(S_{\sigma\alpha}\neg p|S_{\sigma\alpha} \wedge \neg p)$$

This definition presupposes that the probability that any source gives the answer p if p is the case is to be equal to the probability that it gives the answer not- p , if not- p is the case. This *source symmetry* simplifies our calculations, although it can be relaxed if we encounter cases where it does not provide a reasonable approximation.

It follows at once that the reliability of α 's reliability is identical to her aptitude. For other sources, it is an abstraction based on those sources' performances as indications of truth. In general, an inquirer has no direct access to this value, but this does not stop her from forming beliefs about it. Since the number of possible values for the chance $R_{\sigma\alpha}$ is infinite, we need to represent α 's credence as a density function instead of a regular probability distribution. Thus, for each inquirer α , each source σ , and each time t , we define a function $\tau_{\sigma\alpha}^t: [0,1] \rightarrow [0,1]$, called α 's *trust function for σ at t* , such that

$$C_{\alpha}^t(a \leq R_{\sigma\alpha} \leq b) = \int_a^b \tau_{\sigma\alpha}^t(\rho) d\rho$$

for a, b in $[0,1]$. $\tau_{\sigma\alpha}^t(\rho)$ then gives the credence density at ρ , and we can obtain the actual credence that α has in propositions about the reliability of her sources by integrating this function. We will also have use for the expression $1 - \tau_{\sigma\alpha}^t$ (which represents α 's credence density for propositions about σ *not* being reliable) and we will refer to this function as $\bar{\tau}_{\sigma\alpha}^t$.

Now, it is obvious that an inquirer's credences about chances should influence her credences about the outcomes of these chances. The way this should be done is generally known under a name Lewis gave to it: the principal principle (Lewis 1980). This says that if α knows that the chance that an event e will happen is ρ , then her credence in e should be exactly ρ . Applied to our case, this means that the following principle (PP) must hold:

$$C_{\alpha}^t(S_{\sigma\alpha}^t p | S_{\sigma\alpha}^t \wedge R_{\sigma\alpha} = \rho \wedge p) = \rho$$

$$C_{\alpha}^t(S_{\sigma\alpha}^t \neg p | S_{\sigma\alpha}^t \wedge R_{\sigma\alpha} = \rho \wedge \neg p) = \rho$$

for all t , i.e. α 's credence in σ giving the report p , given that the source gives any report at all, that σ 's reliability is ρ , and that p actually is the case, should be ρ .

We also have use for an independence postulate. While not strictly necessary, such a postulate will simplify calculations and modeling considerably. The independence assumption we use here will be referred to as *communication independence* (CI):

$$C_{\alpha}^t(p \wedge S_{\sigma\alpha}^t \wedge R_{\sigma\alpha} = \rho) = C_{\alpha}^t(p) C_{\alpha}^t(S_{\sigma\alpha}^t | R_{\sigma\alpha} = \rho)$$

Communication independence implies that whether σ says anything is independent of whether or not p actually is true, as well as of what reliability σ has.

Given (PP) and (CI) we can now define the crucial expressions $C_{\alpha}^t(p | S_{\sigma\alpha}^t p)$ and $C_{\alpha}^t(p | S_{\sigma\alpha}^t \neg p)$, the credence an agent should place in p at t given that the source σ says that p or not- p , respectively, as follows:

$$C_{\alpha}^t(p | S_{\sigma\alpha}^t p) = \frac{C_{\alpha}^t(p) \langle \tau_{\sigma\alpha}^t \rangle}{C_{\alpha}^t(p) \langle \tau_{\sigma\alpha}^t \rangle + C_{\alpha}^t(\neg p) \langle \bar{\tau}_{\sigma\alpha}^t \rangle}$$

$$C_{\alpha}^t(p|S_{\sigma\alpha}^t \neg p) = \frac{C_{\alpha}^t(p)\langle \bar{\tau}_{\sigma\alpha}^t \rangle}{C_{\alpha}^t(p)\langle \bar{\tau}_{\sigma\alpha}^t \rangle + C_{\alpha}^t(\neg p)\langle \tau_{\sigma\alpha}^t \rangle}$$

where $\langle \tau_{\sigma\alpha}^t \rangle$ is the expected value of the trust function $\tau_{\sigma\alpha}^t$. By the Bayesian requirement of conditionalization, we must have $C_{\alpha}^{t+1} = C_{\alpha}^t(p|S_{\sigma\alpha}^t p)$, whenever σ is the only source giving information to α at t . This means that these formulae completely determines how α should update her credence in such a case.

The calculations become slightly more complex when we take into account the possibility of receiving several messages at the same time. We refer to Angere (forthcoming) for a detailed description. Here we will only hint at the mechanisms involved. Perhaps the most important idealization is that of source independence: we take inquirers to treat their sources as independent, given the truth or falsity of p . This assumption greatly simplifies the updating of the individual credence functions and it is, we submit, psychologically plausible as a principle of default reasoning.⁶

Not only α 's credence in p should be updated, however. Equally important is for α to keep track of how much to trust her sources. A source that generally gives very unlikely reports is unlikely to be veridical, and an inquirer should adjust her trust function in light of this. It turns out that our model already determines how to do this but, once more, we will not go into the details here. Suffice it to mention the following consequence of our model: Even if an inquirer happens to be a perfect inquirer insofar as her inquiry always gives the right result, a fairly low stability of her faith in inquiry, together with her prior judgment that p is unlikely, may conspire to make her distrust her own inquiry. This, in turn, may give rise to a vicious circle in which she becomes more and more convinced that p is false, and that her inquiry is negatively correlated with the truth.

It is time to connect this account of our favored Bayesian model with the previous discussion of the veritistic value of a social practice. What are social practices in this model, and how can we measure their veritistic value? Let us, for a start, define the veritistic value V of a social *network state* $S^t = \langle \Gamma^t, R^t \rangle$:

---

⁶ The issue of source independence is discussed at length in Olsson (forthcoming).

$$V(S^t) = \frac{1}{|I^t|} \sum_{\alpha \in I^t} C_{\alpha}^t(p)$$

Let an *individual state* consist of the values of the epistemic variables of a single inquirer α . Let a *network state* be defined as an assignment of values to these variables for all inquirers in a network, as well as for the links, and a *network evolution* as a sequence $E = S^0, S^1, S^2, \dots$ such that state S^{t+1} is obtainable from the state S^t by having the participants in S^t conditionalize on new information according to the model laid out a few moments ago. Now a *social practice* π can, as a first approximation, be viewed as a constraint on such evolutions, or, equivalently, as a set of them – those evolutions that are compatible with the practice. But every evolution is determined (at least probabilistically) by its initial state S^0 , so we may as well say that the practice is a set of network states, which are to be taken as allowed initial states in applications of the practice.

There is reason to distinguish between two sorts of practices: the *temporary*, which are implemented so as to have a projected end, and the *continual*, which do not have such an end. An example of the first is an investigation made by a committee, while the second type includes policies such as free speech. The practices we focus on in this paper, namely practices of the form “Asserting something only if one’s credence in that thing is over a threshold t ” also belong to the second category. For the first, there is for any evolution a specific final state S^f , and the difference between this state and the initial state appears reasonable to use as a measure. This agrees with Goldman’s own suggestion for how to evaluate practices. For a continual practice, the problem is harder. The only reasonable thing to do may be to identify some finite period during which the practice is to be evaluated. This is the line we will take below.

4. Simulation results

Our problem now is to vary the threshold of assertion and somehow compute the V-value of various such choices. In that way, we could in principle find the optimum threshold for any given situation. However, the problem is how to carry out this proposal in practice given the complexity of the

problem: not only do we have to consider all the possible evolutions of a social network, but for a practice also all the possible initial states it may be applied to. How are we to accomplish this?

Fortunately, there is a simulation program available that we can rely on. Laputa (written by Staffan Angere) is a simulation environment for experiments in social epistemology. It is intended to be useful for investigating several kinds of models of society, but currently it implements the above model.

Laputa allows an experimenter to design social networks, to simulate their evolution, and to collect some useful statistics, including – as we shall see – veritistic values.

Figure 1: A sample social network in Laputa.

When the user hits the “run” button, Laputa runs through a series of steps, each step representing a chance for an inquirer:

- to conduct an inquiry
- to communicate 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 in accordance with the underlying probabilistic model previously described.

Laputa contains a highly useful facility, called the “batch window”, which allows the user to specify constraints on social networks and to instruct Laputa to generate and study a number of networks satisfying the constraints. Since a social practice can be thought of as a constraint on a network or network evolution, this facility allows us to study the effect of a given social practice. Laputa computes the V-value of imposing the constraints imposed as minus “average error delta”.

For instance, the user can instruct Laputa to generate and study only networks in which the inquirers, say, trust each other to a certain degree d . Trusting others to that degree may be understood as a social practice. The output V-value can be interpreted as the veritistic value of that practice in the context of the other constraints imposed. If those other constraints can be considered “normal” a case can be made that the output V-value measures the V-value of the target practice itself.

We will study a case of barely reliable, and well-calibrated, inquirers. They are barely reliable in the sense that they are generally 60 percent reliable in their own inquiries. They are well-calibrated in the sense that, generally, they initially believe (1) that they are 60 percent reliable in their own inquiries, and (2) that the same holds true for their peers. Technically, these parameters are normally distributed around 0.6 with a standard deviation of 0.15. All other parameters are (also) set to reasonable values as shown in Figure 2.

Figure 2: Batch window describing a case of barely reliable inquirers

Referring to Figure 2, the flat distribution for “Starting belief” indicates that Laputa, when selecting the initial degrees of belief for the inquirers in 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 0.6 means that Laputa, when selecting the inquiry accuracy for the inquirers in the generated networks, will have a preference for assigning an accuracy of 0.6 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.

The first simulation consisted in determining how the V-value as a function of the threshold value for a case of just a few – in this case 10 – simulation steps while keeping our background assumptions fixed. This means that each network that Laputa considered was allowed to evolve for 10 steps, each step representing, as we saw, an opportunity for the inquirers in the network to communicate and/or conduct their own inquiries. The results are shown in Figure 3.

Figure 3: The V-value as a function of the threshold of assertion for the case of 10 simulation steps.

As Figure 3 shows, the V-value takes on its maximum for a threshold of about 0.92. Lower V-values are obtained if the threshold is increased beyond that point.

This result raised the question whether the maximum would be obtained for the same threshold value regardless of the number of steps in the simulation. Further simulations showed this not to be the case. Figure 4 plots the V-optimal threshold as a function of the number of simulations steps while keeping everything else fixed.

Figure 4: The V-optimal threshold as a function of the number of simulation steps.

As shown in Figure 4, the V-optimal threshold increases with the number of simulations steps, approaching 1 in the limit.

The simulations suggested that the V-optimal threshold will lie roughly in the interval $0.9 - 1$, and that the exact position of the V-optimal threshold depends on the number of steps in the simulation. More precisely, the more time there is for inquiry and communication, the higher the V-optimal threshold, so that the V-optimal threshold approaches 1 in the limit of inquiry and communication.

To test how robust these results are, we varied other parameters in the model observing how these changes affect the V-optimal threshold of assertion after the simulation has gone through 100 steps. If the results are robust, we should, based on the results in Figure 4, expect the V-optimal threshold to lie somewhere around 0.99, regardless of the values of these other parameters. Since we are interested in the limit of inquiry and communication, we chose to run the simulation for 100 steps so as to give slow network evolutions sufficient time to converge to results close enough to those in the limit.

More specifically, we varied all relevant parameters (network size, starting belief, inquiry accuracy,

inquiry trust, communication trust, inquiry chance, communication chance and link chance), one at a time while holding the others fixed at the values specified in Figure 2. For each such parameter setting, we varied the threshold of assertion so as to find the V-optimal one.

With the following important exceptions, the V-optimal threshold was found to lie around 0.99:

- When all agents already have a sufficiently high initial credence in p (e.g., $C_\alpha(p) = 0.7$), a lower threshold will actually be veritistically superior. This is explained by the fact that under such circumstances, agents are much more likely to communicate p than *not-p* anyway. Therefore, raising the threshold will mostly lead to fewer true messages being communicated. However, for even higher initial credences ($C_\alpha(p) = 0.9$), there were not any statistically significant differences between the tested threshold values. In this case, because of the high initial credence, agents will be almost equally likely to communicate p regardless of threshold value.
- When inquiry is *negatively* correlated with truth ($R_{ia} < 0.5$), a lower threshold value is again more beneficial. Since $R_{ia} < 0.5$, most inquirers will end up moving closer to *not-p*, and a lower threshold value will therefore allow for more p -messages to be communicated. However, regardless of threshold value, $R_{ia} < 0.5$ yields a *negative* V-value.
- Similarly for inquiry trust: when the expected value of the inquiry trust function, $\langle \tau_{ia} \rangle$, is less than 0.5, a lower value is better. Again, this will lead to most inquirers moving closer to *not-p*, so that a lower threshold value will allow for more p -messages to be communicated.
- When the expected value of the communication trust function, $\langle \tau_{\beta\alpha} \rangle$, is lower than 0.5, the V-optimal threshold is 1, whereas for other values of communication trust, it lies around 0.99. In this case, since inquirers are actually (but barely) reliable, having a communication trust lower than 0.5 will on average lead agents to update in the wrong direction, which explains why the V-optimal threshold is 1: this threshold removes the misleading information.

For all other parameter settings we tested, the V-optimal threshold is around 0.99, thereby indicating

that the obtained results are robust to these changes. However, it is possible that there could be synergies between two or more parameters. We plan to address these possible concerns in future work.

5. Discussion and conclusion

Why does the V -value depend on the threshold of assertion in the way displayed in Figure 3 (for 10 simulation steps)? In other words, why does fixing the number of steps lead to an optimum somewhere between 0.9 and 1? We confess that we don't have a clear cut explanation of this fact. But we do think there is a reasonable hypothesis available: A higher threshold means that less communication will take place, and that the inquirer will rely more on her own inquiries. In general this is a good thing: inquiry provides the inquirer with a direct, if somewhat unreliable, "reality check", whereas reports from an informer depend not only on reality, via the informer's inquiries, but also on the credence the informer assigns to the communicated proposition. If the informer's prior credence is assigned randomly according to a flat distribution, as in our experiment, this should make "the word of others" less reliable than inquiry. However, if there are only a few simulation steps, as in this case, the inquirer may receive little input from inquiry, and so she may have little choice but to lend an ear to her peers even if listening to them is not as reliable a method of belief updating as inquiry itself. Therefore, the optimum threshold will be below 1 if the number of simulation steps is not too large. At the same time, that optimum will be close to 1: a small threshold will make the inquirer too influenced by communication with the increased risk of error which that brings with it. Moreover, the fewer the simulation steps, the scarcer the opportunities to inquire, and the more attractive the option of listening to others. In other words, fewer steps will lead to a lower optimum threshold, as witnessed by Figure 4.

There is an alternative, and possibly complementary, explanation of the observation (in Figure 4) of there being a positive dependence between the number of simulation steps and the V -optimal threshold. It is a general fact about Laputa that a society converges, normally in a matter of 15-30 steps, towards an "opinionated" state, i.e. one in which every inquirer either believes fully that p or fully that not- p . Therefore, fewer simulation steps means that fewer inquirers will have time to reach

the threshold of assertion, given that they started out, as in our experiments, with a prior credence sampled from a flat distribution. If the threshold of assertion is held fixed, this means that there will be less communication in the network. Since people are generally reliable in their inquiries, and hence in their communications, people, when they communicate, are more likely to say p than not- p , given that p is assumed true. Hence communication, everything else equal, has a positive general influence on the V-value. Hence, fewer simulation steps will lead to a reduced V-value. Now lowering the threshold compensates for this negative effect by allowing for more communication to take place. Therefore, fewer simulation steps will lead to the optimal threshold being lower, as displayed in Figure 4.

While intuitively compelling, this second potential explanation rests on an assumption that is strictly speaking not true: that communication in general has a positive general influence on the V-value. We tested this claim by varying the link chance as well as the communication chance. As for the former parameter, we found that for sufficiently high thresholds (e.g. 0.99) the maximum V-value is reached when link chance = 1. However, for low thresholds (0.6 or 0.75) a higher link chance leads to a lower V-value. For some values in-between (e.g. 0.9) the optimum V-value is attained for a link chance strictly between 0 and 1. As for the latter parameter, we performed preliminary tests for threshold values of 0.6 and 0.99. For the former, the highest V-value was obtained for communication chance = 0.1. For the latter, the highest V-value was obtained for communication change = 0.4. Given these findings, it is obviously not correct to say that the V-value always increases when the link chance or communication chance increases. More communication is not always a good thing.⁷ What remains true, however, is the fact that there is a positive connection between more communication and higher V-value, although this connection is weaker than we would have expected: given a reasonable threshold, there is always some positive level of communication that is better than no communication at all.

---

⁷ See also Zollman (2007) for a defense of the less-is-more thesis regarding communication. Zollman also works in a Bayesian setting although at the level of details his model is quite different from ours.

The second explanation relies on a second claim that was also found to be easily testable in our framework: the claim that fewer simulation steps will lead to a reduced V-value (*ceteris paribus*). This hypothesis was corroborated in our experiments. We run simulations with 5 to 50 simulation steps for three different thresholds (0.6, 0.75 and 0.99). In all three cases the V-value increased with the number of simulation steps.

In summary, we suggested that norms of assertion should be evaluated based on their epistemic results for society at large, which we identified with Goldman's notion of veritistic value. The simulation study indicates that the certainty rule is veritistically optimal in the limit of inquiry and communication. In finite cases, a threshold less than one will be optimal, depending on the time available for inquiry and communication, and the more time is available, the higher is the optimal threshold. Our robustness test indicated that these results hold so long as the initial credence is not too high, and so long as inquiry reliability, as well as the expected value of inquiry trust and of communication trust, are all greater than 0.5. Future work could investigate whether this also holds when we co-vary other parameters.

We started out by noting that there are two groups of epistemologists: those insisting that nothing short of certainty can warrant assertion and those insisting that this is not so. Which group is right? Our answer is: both. Those who favor the certainty rule are right if their claim is that the certainty rule is optimal "in the limit". Those who reject the certainty rule are also right, if their claim is taken to be one about an investigative and communicative process extending over a finite period of time.⁸

References

Angere, Staffan (forthcoming) "Knowledge in a Social Network", *Synthese*.

DeRose, Keith. (2002) "Assertion, Knowledge, and Context." *The Philosophical Review* 111: 167–203.

---

⁸ Aron Vallinder contributed to this paper by performing and summarizing the results of the robustness tests described in section 4. He also wrote parts of section 5. We thank two anonymous referees for their suggestions.

- Douven, Igor. (2006) "Assertion, Knowledge, and Rational Credibility" *Philosophical Review* 115 (4): 449-485.
- Goldman, A. I. (1999) *Knowledge in a Social World*. Clarendon Press, Oxford.
- Hawthorne, John. (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Jackson, Frank. (1979) "On Assertion and Indicative Conditionals." *Philosophical Review* 88: 565–89.
- Lackey, Jennifer (2007) "Norms of Assertion", *Noûs* 41 (4): 594–626.
- Lewis, David. (1976) "Probabilities of Conditionals and Conditional Probabilities." *Philosophical Review* 85: 297–315.
- Lewis, David. (1980) "A Subjectivist's Guide to Objective Chance", in Richard C. Jeffrey (ed.) *Studies in Inductive Logic and Probability*, vol.2. University of California Press.
- Olsson, Erik J. (2011) "A Simulation Approach to Veritistic Social Epistemology", *Episteme* 8 (2): 127-143.
- Olsson, Erik J. (forthcoming), "A Bayesian Simulation Model of Group Deliberation".
- Stanley, Jason. (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Weiner, Matthew. (2005) "Must We Know What We Say?" *The Philosophical Review* 114: 227–51.
- Williamson, Timothy. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Zollman, K. J. (2007), "The Communication Structure of Epistemic Communities", *Philosophy of Science* 74 (5): 574-587.