



# A diachronic perspective on peer disagreement in veritistic social epistemology

Erik J. Olsson<sup>1</sup> 

Received: 12 June 2015 / Accepted: 5 September 2018 / Published online: 14 September 2018  
© The Author(s) 2018

## Abstract

The main issue in the epistemology of peer disagreement is whether known disagreement among those who are in symmetrical epistemic positions undermines the rationality of their maintaining their respective views. Douven and Kelp have argued convincingly that this problem is best understood as being about how to respond to peer disagreement repeatedly over time, and that this diachronic issue can be best approached through computer simulation. However, Douven and Kelp's favored simulation framework cannot naturally handle Christensen's famous Mental Math example. As a remedy, I introduce an alternative (Bayesian) simulation framework, Laputa, inspired by Alvin Goldman's seminal work on veritistic social epistemology. I show that Christensen's conciliatory response, reasonably reconstructed and supplemented, gives rise to an increase in epistemic (veritistic) value only if the peers continue to recheck their mental math; else the peers might as well be steadfast. On a meta-level, the study illustrates the power of Goldman's approach when combined with simulation techniques for handling the computational issues involved.

**Keywords** Peer disagreement · Social epistemology · Truth · Computer simulation · Bayesian · Alvin I. Goldman

## 1 Introduction

Two individuals are epistemic peers with respect to some question if and only if they are equals with respect to (1) their familiarity with the evidence and arguments which bear on that question and (2) general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias (Kelly 2006, p. 175). Peer disagreement happens when two individuals who are peers with respect to the question at hand disagree about the

---

✉ Erik J. Olsson  
erik\_j.olsson@fil.lu.se

<sup>1</sup> Department of Philosophy, Lund University, Lund, Sweden

answer to the question. The disagreement becomes known when the peers disclose their diverging views to each other.

Plausible cases of known peer disagreement include examples from law and science, such as when the Supreme Court is divided in a difficult legal case or when distinguished professors in paleontology disagree about what killed the dinosaurs. In politics, issues about the legalization of drugs have been mentioned in this connection because “[i]ntelligent, well-informed, well-meaning, seemingly reasonable people have markedly different views on this topic” (Feldman 2006, p. 218), and of course philosophy is the area of known peer disagreement par excellence.

The main issue in the epistemology of peer disagreement is, in Thomas Kelly’s words, “whether known disagreement among those who are epistemic peers in this sense must inevitably undermine the rationality of their maintaining their respective views” (Kelly 2006, p. 175). There are, as one can imagine, two main answers to this question: yes, known peer disagreement does undermine the rationality of the respective views; no, it does not. Yes-people subscribe to some conciliatory response: suspending judgment, splitting the difference, downgrading one’s confidence etc. No-people subscribe to the steadfast response defending the epistemic right to hold on to one’s original view.

The standard argument for being conciliatory is essentially an appeal to symmetry. The only thing that would justify one in maintaining views that are rejected by one’s epistemic peers would be if one had some positive reason to prefer one’s own views over the views of those with whom one disagrees. But, by assumption, no such reason is available in such cases. The proponent of the opposite, steadfast, approach can argue, with Kelly, that to be steadfast “I need not assume that I was better qualified to pass judgment on the question than [my peers] were, or that they are likely to make similar mistakes in the future, or even more likely to make such mistakes than I am. All I need to assume is that *on this particular occasion* I have done a better job with respect to weighing the evidence and competing considerations than they have” (Kelly 2006, p. 180).

My point of departure in this article is the diachronic approach to peer disagreement presented in Douven (2010) and Douven and Kelp (2011). The diachronic approach focuses on strategies for how to react to peer disagreement over the long run rather than on mere one-shot cases. The latter turn out to be less important than is commonly assumed. Douven and Kelp argue that the diachronic case is best investigated through computer simulation, in contrast to the a priori methodology usually employed by epistemologists. However, their simulation framework cannot convincingly handle Christensen’s famous Mental Math example. As a remedy, I introduce an alternative simulation framework, Laputa, which relies on Alvin Goldman’s seminal work on veritistic value (Goldman 1999), and I show how the example can be diachronically addressed in this other setting in line with Christensen’s conciliatory recommendation (with some steadfast qualifications).

## 2 Douven and Kelp on diachronic peer disagreement

Douven (2010) and Douven and Kelp (2011) propose to investigate the issue of peer disagreement from a diachronic perspective. They explain (Douven and Kelp 2011, p. 273):

So far, this issue [of peer disagreement] has been mainly studied from a static or synchronic perspective: only single cases of disagreement have been considered, and authors have queried what the best response in those cases would be. We take instead a more dynamic or diachronic perspective on the matter by considering strategies for responding to disagreements and querying how these strategies fare, in terms of truth approximation, in the longer run, when they are applied over and over again to respond to arising cases of disagreement.

When investigating the diachronic case, Douven and Kelp rely on the influential simulation model due to Hegselmann and Krause (2002, 2006). The model allows agents to be influenced by their own inquiry as well as by the views of other agents (whose views are sufficiently close to their own view). Taking into account the views of other agents means taking the average of their and one's own view, which is one sense of “splitting the difference”.

For the case when the data coming from the world is noisy the simulations demonstrate the following (Douven 2010, p. 151):

- (A) “[O]n average members of the no-difference-splitting societies get within a moderate distance of the truth relatively quickly”.
- (B) “[O]n average, convergence towards the truth occurs at a slower pace for the members of the difference splitting societies, but in the somewhat longer run they are on average much closer to the truth than the members of the no-difference-splitting societies”.

These conclusions are illustrated in Figs. 1 and 2, respectively.

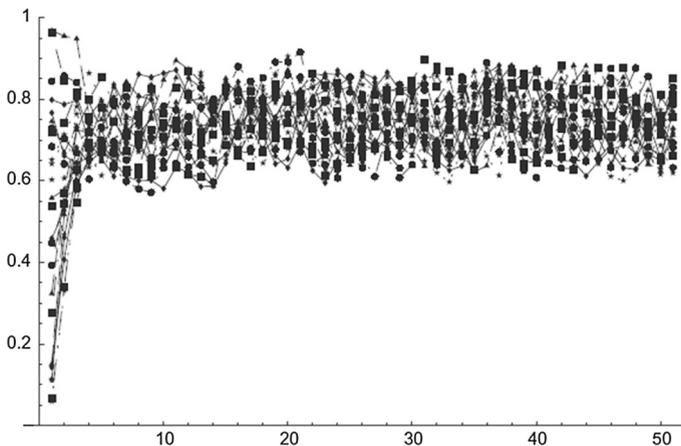
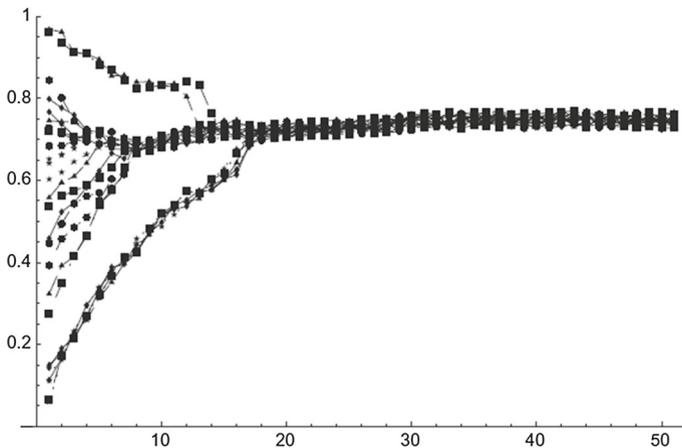


Fig. 1 Noisy data, no difference-splitting. Reproduced from Douven (2010, p. 151)



**Fig. 2** Noisy data, repeated difference-splitting. Reproduced from Douven (2010, p. 151)

Thus, if timing is important and approximating the truth is sufficient, then the steadfast approach is optimal, whereas if getting close to the truth is important and timing is not, then the conciliatory approach is better. Douven concludes that “the main lesson to be learned from the simulations we have looked at is that what is best to do in cases of disagreement with peers may depend on circumstances the obtaining of which is neither general nor a priori” (Douven 2010, p. 156). If this is true, then much of the armchair discussion of peer disagreement in epistemology can be questioned from a methodological standpoint. However, there are some severe problems for the diachronic approach as it is understood by Douven and Kelp.

### 3 Problems for the diachronic approach

In Christensen’s much discussed Mental Math or restaurant example, two peers at a restaurant come up with contrary answers when trying to determine the equal shares of the bill that each of the dining companions should pay. For instance, suppose you are at a restaurant with your friend. Tonight, you figure out that your share is \$43, and become quite confident of this. But then your friend announces that she is quite confident that your share is \$45. Neither of you has had more wine or coffee, and you do not feel (nor does your friend appear) especially tired or unfocused. How confident should you now be that your share is \$43? As Christensen (2007) has noted, the peers should, upon discovering the disagreement, lower their confidence in the correctness of their own belief and raise the confidence in the correctness of the peer’s belief. Christensen’s more specific proposal is that “you should become much less confident in \$43—indeed, you should be about as confident in \$45 as in \$43” (Christensen 2009, p. 757).

By contrast, the Douven–Kelp approach recommends adopting a new view corresponding to the average suggested amount to be paid, which in this case would mean that both individuals, upon discovering their disagreement, should be confident that their shares are \$44. Douven (2010) notes that this outcome is counterintuitive and that the underlying reason is that his model does not consider difference-splitting in

Christensen's sense. Thus, he concludes, "it might be said that, for all I have shown, in the relevant kind of situations, difference-splitting in the way of ... Christensen is always superior, from an epistemic perspective, to no-difference splitting" (Douven 2010, p. 154). If this were true it would undermine Douven's argument against a priori methodology in the study of peer disagreement; resolving such disagreement would after all be a simple and intuitive process amenable to armchair treatment.

However, Douven (*ibid.*, p. 155) proceeds to note a problem for Christensen in turn:

As to Christensen, it is not straightforward either to relate our simulations to his sense of difference-splitting for peers holding contrary beliefs, as there is no indication in Christensen's paper of how, according to him, the peers in the restaurant case, or peers in relevantly similar situations, are to proceed once they have lowered their confidence in the correctness of their own answer ... Christensen may deny this, or he may for some other reason think that the said considerations are irrelevant to his proposal. If so, then here too I am happy to leave it as a challenge to Christensen to be more forthcoming about what his view amounts to.

To illustrate the point, suppose that you were initially 99% confident that your share is \$43 and, similarly, that your friend was initially 99% confident that your share is \$45. Now the disagreement becomes known. According to Christensen this discovery should lead to significantly lower credences; you should, in Christensen's words, be "about as confident in \$45 as in \$43". This means that, unless you go all the way and think there is a 50–50 chance of \$43 or \$45, there is still some disagreement left. Let us say that you and your friend are now only 60% sure that your share is \$43 or \$45, respectively. Then you still disagree about how probable the alternatives are. You think that \$43 is the more likely alternative; your friend thinks that \$45 is more likely. Douven's point, as I understand it, is that Christensen has not explained how to proceed to eliminate this remaining element of disagreement. Christensen's proposal is, in this sense, diachronically incomplete.

It follows that there is no reason to think that the disagreement in the restaurant case can be completely resolved by means of a synchronic, one-shot approach if one adopts the confidence-lowering strategy. The result of one step of confidence-lowering may be a situation in which the peers are still in disagreement because they differ in how they place their confidence. Technically, mutually lowering confidence once may not lead to a fixpoint in the sense that further lowering would not yield a different result. What is needed is a diachronic approach to confidence lowering, which specifies not only what should happen in one step but in the succeeding steps as well.

Christensen's theory is incomplete in other respects, too. It does not specify exactly what degree of confidence the peers should have after the disagreement has become known. More seriously, Christensen does not address the problem how our trust in our peer should be affected by the disagreement. Should we not only lower our confidence in our belief but also lower our trust in the peer? Or should we rather raise this trust? Should trust remain unaffected? As far as I can see, Christensen is silent on this matter as well. In this sense, his proposal is even synchronically incomplete.

To summarize, the problem for Douven and Kelp is that the kind of difference-splitting or suspension which Christensen highlights cannot be naturally modeled in

their framework, and yet it makes good sense in the restaurant case and other similar situations.<sup>1</sup> The problem for Christensen is that his confidence-lowering approach is, in important respects, incomplete: there is lack of information concerning the exact procedure for lowering confidence and, in particular, lack of information as to how it should be repeatedly applied until, hopefully, a fixpoint is reached.

#### 4 Enter Laputa: an alternative simulation framework

The purpose of this section is to show how the problems for Douven and Kelp and for Christensen can be given systematic solutions in an alternative simulation framework called Laputa. Like the Hegselmann–Krause model, the Laputa framework has been studied and applied in a number of articles, and the reader is referred to these other studies for details about the mathematical model. Olsson (2011) uses Laputa to illustrate how the computational problem that arises for Goldman’s theory of social epistemology can be solved using computer simulation. Olsson (2013) and Master-ton and Olsson (2013) investigate Laputa as a theory of argumentation. Olsson and Vallinder (2013) applies Laputa to the debate in epistemology about the knowledge rule of assertion. Vallinder and Olsson (2013) uses Laputa to assess the strength of the so-called argument from disagreement in the ethics. Vallinder and Olsson (2014) contains technical background material on the underlying Bayesian model and investigates what happens when an agent’s trust in her own inquiry diverges from her actual reliability.<sup>2</sup>

The Laputa model differs from the Hegselmann–Krause model used by Douven and Kelp e.g. in keeping track of agents’ confidence in a proposition. These confidence values, or credences, are updated in accordance with Bayesian principles. These distinguishing features will be crucial when we turn, in the next section, to the representation of the Mental Maths example in Laputa.

In Laputa, the credence assigned to a proposition  $p$  by an inquirer  $S$  after communication in a social network depends (among other things) on:

- Reports from  $S$ ’s outside source, also termed “inquiry” in the following
- How many of  $S$ ’s network peers claim that  $p$ /not- $p$ .
- How often they do it.
- $S$ ’s trust in her network peers.

Among the notable assumptions and idealizations we find the following:

- At every step/round in deliberation, inquirers (outside sources) can communicate  $p$ , not- $p$  or be silent.
- Trust (= perceived reliability) is modelled as a second order probability: a credence in the reliability of the source.

<sup>1</sup> It may be possible in principle to model the restaurant case in the Hegselmann–Krause model by treating the numbers on the y-axis as credences. However, due to the way updating works in that model, by averaging rather than by conditionalization, the result wouldn’t correspond to a standard Bayesian approach. I assume throughout this paper that credences should be updated in the standard way prescribed by Bayesianism. I am grateful to an anonymous referee for demanding greater clarity on this point.

<sup>2</sup> Thus, Douven and Kelp’s remark about Laputa that “results of systematic simulations within this environment are still to be reported” (2011, p. 277, fn. 10) is no longer true.

- Reports coming from different sources at the same time are viewed by receiving inquirers as independent, in the sense that they have not, for example, colluded to report a certain result.
- Reports from outside sources are treated as independent.

Olsson (2013) argues that these assumptions are in line with Persuasive Argument Theory in social psychology and that they are justifiable from a dual process perspective found in cognitive science.

In the following I draw on Alvin I. Goldman seminal work in social epistemology on veristic value or V-value for short (Goldman 1999). The fundamental thought is that, ideally, an inquirer should have full belief in the truth, and that inquirers are accordingly better off the closer their degree of belief in the truth is to full belief.<sup>3</sup> Social arrangements (practices) are evaluated with regard to the average gain/loss in V-value that they induce. In particular, this goes for the social practice of reacting in a certain way to peer disagreement. Henceforth we assume that  $p$  is a proposition that is true and that not- $p$  false. Our Mental Maths example will involve two individuals, John and Mary, and we may think of  $p$  as the proposition that John’s share to pay for the dinner is \$43 and of not- $p$  as the proposition that he should pay \$45.

The Laputa model has been implemented in a simulation program. Once a network structure has been specified, the Laputa simulator can run tens of thousands of simulations holding that network fixed in “batch mode” while varying the background parameters, such as initial credence in  $p$ , reliability of outside sources etc. We will see how this works shortly in the case of the Mental Math example. Laputa outputs the average gain/loss in V-value and other interesting statistical information over all these simulations, making it possible to study the average collective “truth gain” resulting from responding to peer disagreement under various circumstances.

### 5 The synchronic and diachronic completeness of Laputa

We recall that our goal is to model the reaction to peer disagreement identified by Christensen which means lowering our confidence in our own view. Moreover, we want an account of this reaction that is diachronically complete in the sense that it can be used for studying the effects of this practice in the long run and not just in one-shot cases. In Laputa, both credences and trust values are updated dynamically by conditionalizing on what network peers and/or outside sources report to be the case. For instance, an agent  $\alpha$ ’s new credence in  $p$  after having listened at time  $t$  to source  $\sigma$  reporting that  $p$  or not- $p$ , respectively, is determined by the following equations referring to prior credence and (expected) trust:

$$C_{\alpha}^t(p|S_{\sigma\alpha}^t(p)) = \frac{C_{\alpha}^t(p) \times E[\tau_{\sigma\alpha}^t]}{C_{\alpha}^t(p) \times E[\tau_{\sigma\alpha}^t] + C_{\alpha}^t(\neg p) \times E[\overline{\tau_{\sigma\alpha}^t}]}$$

$$C_{\alpha}^t(p|S_{\sigma\alpha}^t(\neg p)) = \frac{C_{\alpha}^t(p) \times [\overline{\tau_{\sigma\alpha}^t}]}{C_{\alpha}^t(p) \times E[\overline{\tau_{\sigma\alpha}^t}] + C_{\alpha}^t(\neg p) \times E[\tau_{\sigma\alpha}^t]}$$

<sup>3</sup> See also Joyce (1998) for a similar account.

**Table 1** Derived updating rules in Laputa

	Message expected	Neither nor	Message unexpected
Source trusted	+ (↑)	↑ 0	– (↓)
Neither nor	0 (↑)	0 0	0 (↓)
Source distrusted	– (↑)	↓ 0	+ (↓)

Thus the new credence is determined by conditionalization on the report that  $p$  or not- $p$ , respectively, is the case. These two equations, together with the corresponding equations for updating trust values, give rise to a number of derived qualitative updating rules, as described in Table 1 (see Vallinder and Olsson 2014, for details and derivations).

For instance, the upper leftmost box in Table 1 describes the case in which an expected message comes from a trusted source, be it from inquiry or from another agent in the network. The + -sign means that the listener's current confidence is raised and the arrow in parenthesis means that the trust in the source goes up as well. To take a concrete example, John may initially be confident to degree 0.99 that his share is \$43. Now he receives the report that this is indeed true from Mary, whom he considers more trustworthy than not. The corresponding qualitative rule then specifies that both John's confidence in the proposition that he should pay \$43 and his trust in Mary should increase.

The more interesting case, for our purposes, is the upper rightmost box in Table 1 corresponding to the case in which the receiving agent obtains an unexpected message from a trusted source. This is the case which is at the center of the debate on peer disagreement. In the Mental Maths case this would be a scenario in which John is initially confident that his share is \$43, only to hear from Mary, whom he trusts, that his share is \$45. Of course Mary's situation is exactly symmetrical when she is confident that John's share is \$45, only to hear from John, whom she trusts, that his share is \$43.<sup>4</sup>

In a case like this, Laputa recommends that the peers lower their confidence in their current beliefs; this is what the – sign means in Table 1. This corresponds to Christensen's conciliatory view regarding what should happen in the restaurant case. We noted that Christensen's account is in a sense synchronically incomplete: it does not specify how the disagreement should affect the peers' trust in each other as information sources. Laputa gives a specific answer to this question which amounts to the recommendation that trust should go down as well; this is what the down-arrow means in the relevant box in Table 1. Thus Laputa gives an account of peer disagreement which is synchronically complete: it specified what should happen, in qualitative terms, to both confidence in the proposition in question and trust in the peer. The underly-

<sup>4</sup> There is a dispute regarding what the evidence is in cases of peer disagreement. From the Bayesian perspective adopted in the paper, this dispute translates into a dispute about how to describe the evidence that each peer conditionalizes upon when updating his or her credence in a case of peer disagreement. For example, the evidence could be viewed as "the fact that we disagree", "the fact that we disagree on a specific issue", "the fact that you say that not- $p$  when I believe that  $p$ " etc. In the paper, I investigate a specific and particularly simple way of spelling out the evidence as simply the report given by my peer, in a context in which that report conflicts with my own belief.

ing Bayesian machinery outputs specific numbers in all concrete updates, leaving no details for guesswork.

On the basis of these observations one could argue that Laputa reconciles the conciliatory and steadfast approaches to peer disagreement by giving some credit to both schools of thought, although I will not press this point here. The Laputa reaction is conciliatory in the sense that both peers lower their confidence in their belief, thus giving some weight to the peer's opinion. It is steadfast in the sense that the peers lower their trust in each other, thus signaling some doubt in the credibility of the other. It is like saying "I may be wrong, and so I lower my credence in my belief, but so may you, and so I lower my trust in your reliability". As we will see in simulations to come the reduction of trust that is occasioned by peer disagreement is very small—but these small changes may diachronically have significant effects.

As is also made clear in Table 1, Laputa gives an account of peer disagreement that is diachronically complete. After all, the table specifies what should happen for all combinations of trust/distrust and expectedness/unexpectedness of the message. Thus, whatever happens as the peers continue to interact after their initial agreement that activity will be covered by some Laputa rule.

*Prima facie*, then, Laputa systematically solves the problems that we identified for Douven and Kelp as well as for Christensen. It provides an account of peer disagreement along Christensen's lines and can, unlike Douven and Kelp's model, handle the restaurant example. Moreover, unlike Christensen, Laputa specifies how to proceed in a way that is synchronically and diachronically complete. Given the fact that the Laputa table for updating is derived from an underlying Bayesian framework there is strong systematic support for its particular way of handling these matters. In my view, this systematic support has been strengthened by the fact that the theory, as I noted, has been applied to a number of philosophical problems with, as I see it, interesting and enlightening results. Of course one may think of alternative ways of updating credence and trust but then the challenge would be to mount the same systematic or other support for the resulting theory as already exists for Laputa.

## 6 Simulating peer disagreement in Laputa

In order to simulate peer disagreement in Laputa we will draw a social network involving two agents who initially disagree. In Fig. 3, the nodes represent agents—John and Mary—and the links are communication links between them. A link from one agent to another indicates that the first agent can send messages ( $p$  or not- $p$ ) to the second agent. A green node means that the agent believes the truth. We stipulate that the truth in this case is that John's share is \$43. A red node means that the agent in question holds a false belief—that John's share is \$45. Thus the situation in Fig. 3 is that John is right and Mary is wrong about the amount to be paid by John. The color of the links go from green to red, whereby green represent the case in which the receiver fully trusts the sender, and red the case in which the receiver fully distrusts the sender, i.e. views the latter as a systematic liar. As for John and Mary, they both strongly trust each other.

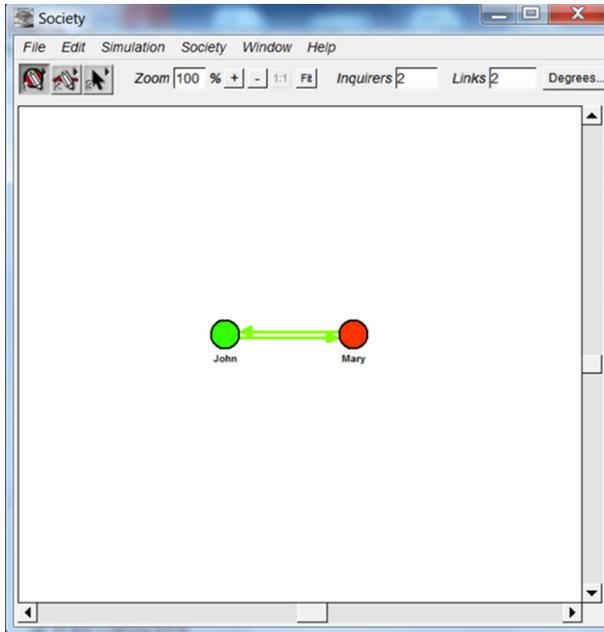
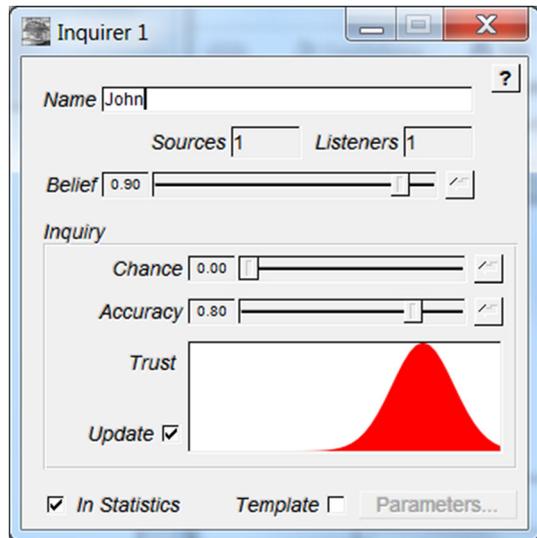


Fig. 3 A social network of two disagreeing agents, John and Mary

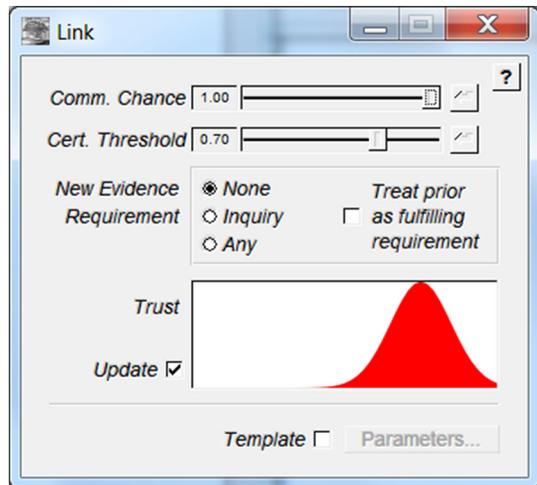
Fig. 4 John's properties



In the Laputa simulator there are a number of properties that need to be specified for both the nodes and the links. The node/agent properties chosen for John are depicted in Fig. 4.

The belief parameter in Fig. 4 signifies John's initial degree of belief in  $p$ , which is in this example 0.90. Inquiry chance is the chance that John, at a given step in

**Fig. 5** Link properties for the link from John to Mary



the simulation, will conduct inquiry, i.e. consult an outside source (a source not in the network). We will think of this as John checking his mental math. This chance is set to 0 in our first simulation experiment, meaning here that John will never redo his mental math. Therefore, the inquiry accuracy parameter does not matter. For the record, this parameter controls the reliability of the outside source, which will be relevant in our second experiment. The final parameter, trust, refers here to John's trust in his own inquiry, i.e. in his mental math ability. For reasons that we do not have to go into, this property is modeled as a probability distribution over all possible degrees of reliability: from complete unreliability (systematic lying) to complete reliability (truth-telling) rather than as a single value. In many cases, however, it suffices to consider the expected value of this distribution (expected trust) which is a single number. The situation depicted in Fig. 4 is one of rather high, but not, complete trust in the mental math. Again, since the first experiment does not involve inquiry this parameter will not play any role. The properties for Mary are exactly the same as those for John, thus making their epistemic situations completely symmetrical as required by peerhood.

Every link in Laputa has its own properties that can be specified. Figure 5 depicts the properties of the link from John to Mary in our experiment. The communication chance is the probability that John will report something ( $p$  or not- $p$ ) to Mary at a given step in the simulation, given that his confidence is over the certainty threshold. The latter is here set to 0.7 and the former to 1, meaning that if John's confidence in one of the proposition  $p$  or not- $p$  is over 0.7, then he is certain to report this view to Mary. The new evidence requirement does not play any role in this paper. Trust, as a link property, signifies the trust that the receiver puts in the sender, i.e. in this case Mary's trust in John. We set the parameters for the link from Mary to John to exactly the same values, thus securing their peerhood.

We start the experiment by pressing the go-button in single mode in Laputa. What happens then is that the peers start communicating in accordance with the conditions and probabilities set in the windows for inquirer and link properties.

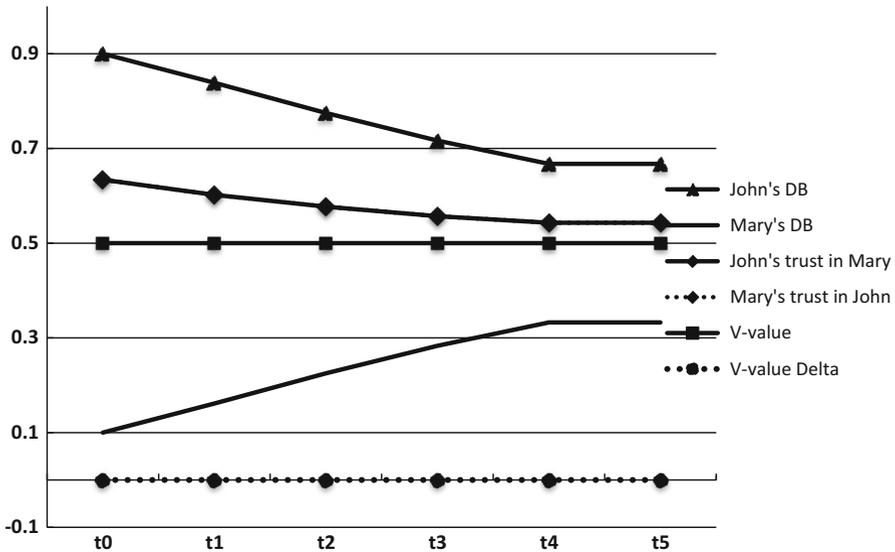


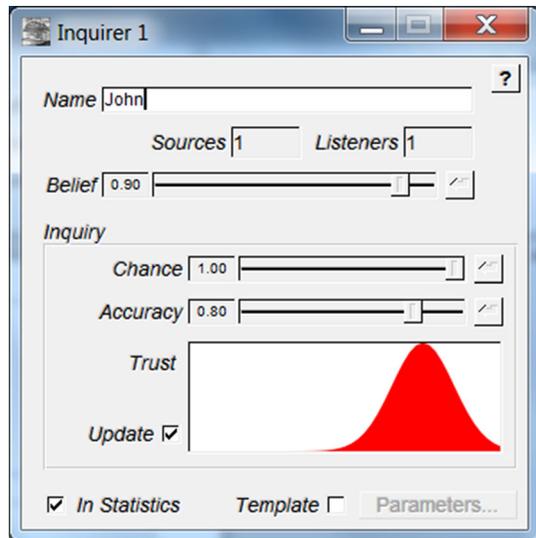
Fig. 6 The evolution of degree of belief (DB), trust and V-value in experiment 1

As Fig. 6 shows, the peers step-wise downgrade their confidence in their original belief and their trust in their peer. This happens until they reach a point where their confidence in either  $p$  or not- $p$  is below the certainty threshold. After that point, there is no more communication between them which means here that there is no further change in either confidence or trust. We also note that only the “peer disagreement box” (upper rightmost box in Table 1) is used. Hence, diachronic completeness is not really important if there is no inquiry going on. See “Appendix 1” for a print out of the output of Laputa in experiment 1.

The end result is that the peers end up with less confidence in their original belief. They also trust each other less than they did before, in the sense that mutual trust is slowly but surely eroding. (John’s trust in Mary and Mary’s trust in John stay the same throughout the experiment, whence the trust curves in Fig. 6 overlap.) The V-value is not affected in this process meaning that there is no progress (or backlash, for that matter) regarding the average distance to full belief in the truth. This is due to the fact that both peers approach a degree of belief 0.5 symmetrically from different directions, as it were. A consequence is that from a purely veritistic perspective the peers might as well have been (fully) steadfast in the sense of holding on to their original degree of belief.

Our next question is what happens if the peers (continue to) inquire, i.e. to check and recheck their mental math. This was studied in a second simulation experiment. The parameter settings for the new experiment were exactly the same as those for the first experiment, except that inquiry chance was now set to 1 for both John and Mary. Thus, both will not only communicate but also “ask their outside source”, i.e. consult

**Fig. 7** John with inquiry  
chance = 1



their mental math faculty, at each step of the simulation.<sup>5</sup> Figure 7 depicts John’s properties in the new experiment which correspond exactly to Mary’s properties.

In experiment 2, Mary’s initial error is gradually corrected and both peers approach a strong belief in the truth (Fig. 8). This particular run of the experiment involves the slightly unlikely event that Mary initially receives the wrong result (not- $p$ ) from inquiry. In the light of the actual reliability of her inquiry, this initial mistake can be attributed to an error “on this particular occasion”, in Kelly’s words, as quoted in the introduction. In the longer run, inquiry will deliver mostly correct result, which accounts for the gradual convergence of both peers on (close to) full belief in the truth. As a result there is a steady increase in V-value over time. As one would expect, the peers increasingly trust each other, as well as their own inquiry.

A printout of the output of Laputa in experiment 2 can be found in “Appendix 2”. A closer examination reveals that the updating makes use of several boxes in the Laputa updating table, not just the peer disagreement box. Thus, once inquiry is introduced in the diachronic process of negotiating peer disagreement, the diachronic completeness of the underlying model becomes crucial.

The actual increase in V-value is contingent on the string of results that come from inquiry. The average performance of this peer system can be studied automatically in Laputa’s batch window. The average result of 10,000 simulations, each of 30 steps, turns out to be V-value  $\Delta = 0.155$  which is much lower than the value obtain in experiment 2 where the string of results from inquiry turned out to be fortunate. In other words, the average increase in V-value if the peers engage in a repeated exchange of the type used in experiment 2 is 0.155. Thus, this practice is on average veritistically beneficial under the conditions specified. The social practice which these peers are

<sup>5</sup> A probability of 1 is interpreted in Laputa as a probability very close to 1. Literally assigning probability 1 to a proposition is alien to a Bayesian framework because it means that no revision is possible. Thus, when inquiry chance is set to 1, absence of inquiry is still possible, yet very unlikely.

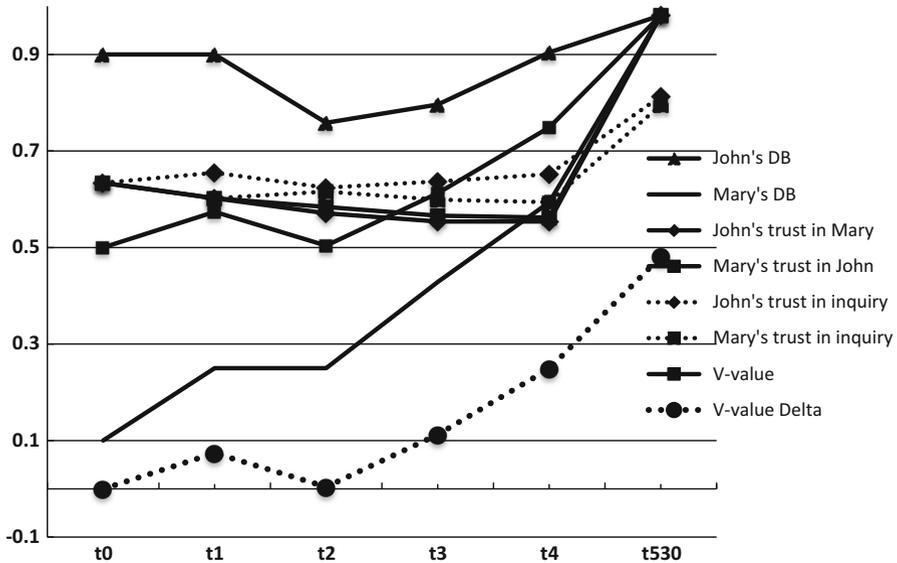


Fig. 8 The evolution of degrees of belief (DB), trust and V-value in experiment 2

engaged in is a good one, from the point of view of veritistic social epistemology à la Goldman, given that the circumstances are as specified by the parameter values.

## 7 Conclusion

My point of departure was Douven and Kelp's observation that the problem of peer disagreement is best understood to be about how to respond to such disagreement over time, and that due to its context-dependence this diachronic issue can be best investigated through computer simulation. Alas, we saw that the Hegelmann-Krause model upon which Douven and Kelp's study relies cannot convincingly handle Christensen's Mental Math example. Moreover, we found Christensen's own recommendation for how to proceed lacking in detail, synchronically as well as diachronically.

As a remedy, I proposed the use of the Laputa simulation framework to handle the Mental Math case. Laputa relies on Alvin Goldman's seminal work on veritistic value for the computation of veritistic (epistemic) values. I showed how the Mental Math example can be diachronically addressed in line with Christensen's recommendation, supplemented with a minor steadfast element concerning the way in which trust was handled. I argued that there is strong systematic support for this supplement from the underlying Bayesian engine and the many applications that have emerged from Laputa. In this study it was crucial that Laputa is both synchronically and diachronically complete which makes it particularly suitable for a diachronic assessment of Christensen's example.

A critic may be dissatisfied with the claim that the present approach includes a steadfast element, on the following grounds. Unless we both lower our confidence

right down to the same degree, we will continue to disagree over how confident we should be in our contrary beliefs, as well as over how much trust we should place in the other's reliability. But this looks like an entirely conciliatory resolution of disagreement over both confidence and trust. In particular, there's nothing steadfast about lowering my trust in your reliability down to the same degree to which you are now confident (or trust) your own reliability; likewise, you are not displaying a firm an unwavering stance when lowering your trust in my reliability down to the same degree to which I'm now confident (or trust) my own reliability.<sup>6</sup> In response, it should be noted that the synchronic updating rules in Laputa contain a steadfast element in the sense that the downgrading of my trust in my peer is an expression of my self-confidence—a manifestation of my sentiment that you may be wrong and I may be right. To be sure, to many adherents of the traditional steadfast approach this will still look like a spineless response to peer disagreement. Yet, I am not hoping to capture steadfastness in the full sense of “a firm and unwavering stance” but rather in the weaker sense of self-confidence, a trust in one's own ability as an inquirer and epistemic agent that may be less than complete. The claim is that the present account involves an element of steadfastness or self-confidence, not the full package.

Another interesting aspect of the objection concerns the synchronic/diachronic distinction and its consequences. The long term effect of Laputa updating in a case of initial peer disagreement may be that the peers eventually reach perfect agreement not only regarding the proposition initially under dispute but also regarding trust and reliability, in the sense that one peer may end up trusting the other exactly at the level at which that peer trusts herself (or, to be exact, her own inquiry). What is steadfast or even self-confident about this result? It is certainly true that the result thus reached does not itself bear any particular mark of self-confidence. Even so, it is still the case that this result was reached through the repeated use of synchronic rules that are partly based on self-confidence at the level of peer trust-updating, as previously explained. Maybe the more general methodological insight here is that when we are working at the present level of detail, which surpasses that of most theorizing in the literature, we also need to sharpen the linguistic tools that we use to describe and interpret the updating mechanisms and the longer-term results thereby obtained. It seems to me that the terms “conciliatory” and “steadfast” may in the end be too crude to capture the nuances and distinctions that can be made using the resources provided by a more fine-grained and complete account.

What comes out of the simulation study conducted in this paper is that Christensen's approach, as reconstructed here, gives rise to an average gain in veritistic value only if the peers continue to recheck their results; else they might as well be (completely) steadfast. This complexity can be taken in support of Douven and Kelp's general methodological point that peer disagreement, properly understood in the diachronic sense, is a heavily context-dependent affair upon which a priori reasoning alone is unlikely to shed much reliable light. A priori reasoning still plays a role in assessing the plausibility of the models used, but when it comes to the actual behavior of various responses to peer disagreement over time we need more advanced tools. On a meta-

---

<sup>6</sup> I am indebted to an anonymous referee for raising this objection.

level, the present study illustrates, I hope, the power of the veritistic approach to social epistemology when combined with simulation techniques.

Which simulation framework is overall best for studying peer disagreement—Laputa or the Hegselmann–Krause model upon which Douven and Kelp’s work is based? The most attractive answer is probably to reject the presupposing of the question and advocate pluralism. While Laputa can handle examples like the Mental Math more convincingly than the H–K model, there are cases in which repeated averaging over peer opinions is a natural process, as evidenced by Douven and Kelp’s study. In such cases Hegselmann and Krause’s model would be the natural first choice.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

## Appendix 1: Laputa output in experiment 1

### Time: 1

Inquirer ‘John’ heard that not- $p$  from inquirer ‘Mary’, lowering his/her trust in the source from 0.634 to 0.602.

This lowered his/her degree of belief in  $p$  from 0.90000 to 0.83863.

Inquirer ‘Mary’ heard that  $p$  from inquirer ‘John’, lowering his/her trust in the source from 0.634 to 0.602.

This raised his/her degree of belief in  $p$  from 0.10000 to 0.16137.

V-value = 0.500, V-value  $\Delta$  = 0.

### Time: 2

Inquirer ‘John’ heard that not- $p$  from inquirer ‘Mary’, lowering his/her trust in the source from 0.602 to 0.577.

This lowered his/her degree of belief in  $p$  from 0.83863 to 0.77487.

Inquirer ‘Mary’ heard that  $p$  from inquirer ‘John’, lowering his/her trust in the source from 0.602 to 0.577.

This raised his/her degree of belief in  $p$  from 0.16137 to 0.22513.

V-value = 0.500, V-value  $\Delta$  = 0.

### Time: 3

Inquirer ‘John’ heard that not- $p$  from inquirer ‘Mary’, lowering his/her trust in the source from 0.577 to 0.557.

This lowered his/her degree of belief in  $p$  from 0.77487 to 0.71657.

Inquirer ‘Mary’ heard that  $p$  from inquirer ‘John’, lowering his/her trust in the source from 0.577 to 0.557.

This raised his/her degree of belief in  $p$  from 0.22513 to 0.28343.

V-value = 0.500, V-value  $\Delta = 0$ .

**Time: 4**

Inquirer 'John' heard that not-p from inquirer 'Mary', lowering his/her trust in the source from 0.557 to 0.543.

This lowered his/her degree of belief in p from 0.71657 to 0.66746.

Inquirer 'Mary' heard that p from inquirer 'John', lowering his/her trust in the source from 0.557 to 0.543.

This raised his/her degree of belief in p from 0.28343 to 0.33254.

V-value = 0.500, V-value  $\Delta = 0$ .

**Time: 5**

(Silence)

V-value = 0.500, V-value  $\Delta = 0$ .

**Time: 6**

(Silence)

V-value = 0.500, V-value  $\Delta = 0$ .

**Appendix 2: Laputa output in experiment 2****Time: 1**

Inquirer 'John' received the result that p from inquiry, raising his/her expected trust in it from 0.634 to 0.655.

Inquirer 'John' heard that not-p from inquirer 'Mary', lowering his/her expected trust in the source from 0.634 to 0.602.

This did not affect his/her degree of belief in p.

Inquirer 'Mary' received the result that p from inquiry, lowering his/her expected trust in it from 0.634 to 0.602.

Inquirer 'Mary' heard that p from inquirer 'John', lowering his/her expected trust in the source from 0.634 to 0.602.

This raised his/her degree of belief in p from 0.10000 to 0.24993.

V-value = 0.574, V-value  $\Delta = 0.074$ .

**Time: 2**

Inquirer 'John' received the result that not-p from inquiry, lowering his/her expected trust in it from 0.655 to 0.624.

Inquirer 'John' heard that not- $p$  from inquirer 'Mary', lowering his/her expected trust in the source from 0.602 to 0.571.

This lowered his/her degree of belief in  $p$  from 0.90000 to 0.75854.

Inquirer 'Mary' received the result that not- $p$  from inquiry, raising his/her expected trust in it from 0.602 to 0.616.

Inquirer 'Mary' heard that  $p$  from inquirer 'John', lowering his/her expected trust in the source from 0.602 to 0.584.

This lowered his/her degree of belief in  $p$  from 0.24993 to 0.24993.

V-value = 0.504, V-value  $\Delta$  = 0.004.

### Time: 3

Inquirer 'John' received the result that  $p$  from inquiry, raising his/her expected trust in it from 0.624 to 0.637.

Inquirer 'John' heard that not- $p$  from inquirer 'Mary', lowering his/her expected trust in the source from 0.571 to 0.554.

This raised his/her degree of belief in  $p$  from 0.75854 to 0.79622.

Inquirer 'Mary' received the result that  $p$  from inquiry, lowering his/her expected trust in it from 0.616 to 0.599.

Inquirer 'Mary' heard that  $p$  from inquirer 'John', lowering his/her expected trust in the source from 0.584 to 0.566.

This raised his/her degree of belief in  $p$  from 0.24993 to 0.42859.

V-value = 0.612405, V-value  $\Delta$  = 0.1124.

### Time: 4

Inquirer 'John' received the result that  $p$  from inquiry, raising his/her expected trust in it from 0.637 to 0.652.

This raised his/her degree of belief in  $p$  from 0.79622 to 0.90426.

Inquirer 'Mary' received the result that  $p$  from inquiry, lowering his/her expected trust in it from 0.599 to 0.594.

Inquirer 'Mary' heard that  $p$  from inquirer 'John', lowering his/her expected trust in the source from 0.566 to 0.562.

This raised his/her degree of belief in  $p$  from 0.42859 to 0.59362.

V-value = 0.74894, V-value  $\Delta$  = 0.24894.

### Time: 530

Inquirer 'John' received the result that  $p$  from inquiry, raising his/her expected trust in it from 0.812 to 0.813.

Inquirer 'John' heard that  $p$  from inquirer 'Mary', raising his/her expected trust in the source from 0.981 to 0.981.

This did not affect his/her degree of belief in  $p$ .

Inquirer 'Mary' received the result that  $p$  from inquiry, raising his/her expected trust in it from 0.794 to 0.795.

Inquirer ‘Mary’ heard that  $p$  from inquirer ‘John’, raising his/her expected trust in the source from 0.981 to 0.981.

This did not affect his/her degree of belief in  $p$ .

V-value = 0.981, V-value  $\Delta$  = 0.481.

## References

- Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, *116*, 187–217.
- Christensen, D. (2009). Disagreement as evidence: The epistemology of controversy. *Philosophy Compass*, *4*(5), 756–767.
- Douven, I. (2010). Simulating peer disagreements. *Studies in History and Philosophy of Science*, *41*, 148–157.
- Douven, I., & Kelp, C. (2011). Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis*, *75*, 271–283.
- Feldman, R. (2006). Epistemological puzzles about disagreement. In S. Hetherington (Ed.), *Epistemology futures* (pp. 216–236). Oxford: Clarendon Press.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis, and simulations. *Journal of Artificial Societies and Social Simulation* 5. Available at <http://jasss.soc.surrey.ac.uk/5/3/2.html>. Accessed 1 Jan 2010.
- Hegselmann, R., & Krause, U. (2006). Truth and cognitive division of labor: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation* 9. Available at <http://jasss.soc.surrey.ac.uk/9/3/10.html>. Accessed 12 Feb 2011.
- Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, *65*(4), 575–603.
- Kelly, T. (2006). The epistemic significance of disagreement. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 1, pp. 167–196). Oxford: Oxford University Press.
- Masterton, G., & Olsson, E. J. (2013). Argumentation and belief updating in social networks: A Bayesian model. In E. Fermé, D. Gabbay, & G. Simari (Eds.), *Trends in belief revision and argumentation dynamics*. London: College Publications.
- Olsson, E. J. (2011). A simulation approach to veritistic social epistemology. *Episteme*, *8*(2), 127–143.
- Olsson, E. J. (2013). A Bayesian simulation model of group deliberation and polarization. In F. Zenker (Ed.), *Bayesian argumentation, synthese library* (pp. 113–134). New York: Springer.
- Olsson, E. J., & Vallinder, A. (2013). Norms of assertion and communication in social networks. *Synthese*, *190*, 1437–1454.
- Vallinder, A., & Olsson, E. J. (2013). Do computer simulations support the argument from disagreement? *Synthese*, *190*(8), 1437–1454.
- Vallinder, A., & Olsson, E. J. (2014). Trust and the value of overconfidence: A Bayesian perspective on social network communication. *Synthese*, *191*, 1991–2007.