Etica & Politica / Ethics & Politics, XV, 2013, 2, pp. 25–41

Is Epistemic Trust of Veritistic Value?[*]

Gregor Betz
Karlsruhe Institute of Technology
gregor.betz@kit.edu

Michael Baurmann
Heinrich Heine Universitaet Duesseldorf
baurmann@hhu.de

Rainer Cramm
Alfried Krupp Wissenschaftskolleg Greifswald
rainer.cramm@wiko-greifswald.de

ABSTRACT
Epistemic trust figures prominently in our socio-cognitive practices. By assigning different (relative) degrees of competence to agents, we distinguish between experts and novices and determine the trustworthiness of testimony. This paper probes the claim that epistemic trust furthers our epistemic enterprise. More specifically, it assesses the veritistic value of competence attribution in an epistemic community, i.e., in a group of agents that collaboratively seek to track down the truth. The results, obtained by simulating opinion dynamics, tend to subvert the very idea that competence ascription is essential for the functioning of epistemic collaboration and hence veritistically valuable. On the contrary, we find that, in specific circumstances at least, epistemic trust may prevent a community from finding the truth effectively.

KEYWORDS
Epistemic trust, competence ascription, veritistic value, opinion dynamics, bounded confidence, simulation

1. *Introduction*

Epistemic trust figures prominently in our socio-cognitive practices. By assigning different (relative) degrees of competence to agents, we distinguish between experts and novices and determine the trustworthiness of testimony. Given the extensive division of cognitive labour in our societies in general

[*]We'd like to thank two anonymous reviewers for helpful suggestions. Our research has been supported by the Alfried Krupp Wissenschaftskolleg Greifswald, Germany.

and in science in particular[1], the practice of competence ascription seems to be essential for successful epistemic cooperation.[2] Studies in the history of science, too, confirm the import of competence attribution and epistemic trust.[3]

This paper probes the claim that epistemic trust furthers our epistemic enterprise. More specifically, it assesses the veritistic value of competence attribution in an epistemic community, i.e., in a group of agents that collaboratively seek to track down the truth. The results, obtained by simulating opinion dynamics, tend to subvert the very idea that competence ascription is essential for the functioning of epistemic collaboration and hence veritistically valuable. On the contrary, we find that, in specific circumstances at least, epistemic trust may prevent a community from finding the truth effectively. This is not to say, however, that epistemic trust is always veritistically detrimental. Our simulation results confirm that socio-cognitive practices which strongly couple competence ascription and objective expertise may improve the collective truth-tracking ability.

While this study falls well in the field of veritistic social epistemology [Goldman, 1999], it modifies and applies models of opinion dynamics to assess the practice of expertise ascription.[4] Simulations of opinion dynamics are suitable for studying the veritistic value of a social practice, such as establishing relations of epistemic trust, because they allow one to model (i) how the beliefs of agents change in time, (ii) to which degree agents approach the truth and (iii) to which extent the agents' doxastic trajectories depend on different social or institutional settings. A variety of different models of opinion dynamics have been developed in the last decade.[5] Our model can be considered a synthesis of the Hegselmann-Krause model [Hegselmann and Krause, 2002, 2006] and the Lehrer-Wagner model [Lehrer and Wagner, 1981], as explained below.[6] To our knowledge, models of opinion dynamics have, so far, not been directly employed to investigate the veritistic value of competence ascription and epistemic trust. Still, differentiating between agents according to ascribed competence can be considered, quite generally, as a way of reducing information exchange in an epistemic group: Rather than listening

[1] Kitcher [1990, 1993] pioneered the philosophical investigation of division of cognitive labour in science; for recent, simulation-based studies, see Weisberg and Muldoon [2009], Muldoon and Weisberg [2011].

[2] See also Goldman [2001].

[3] Rudwick's seminal study of the Great Devonian Controversy, for instance, makes plainly clear that the perceived degrees of expertise had a major impact on the debate evolution [see in particular Rudwick, 1985, pp. 418-426].

[4] We hence apply a simulation approach to social epistemology [cf. Olsson, 2011].

[5] See, e.g., Deffuant et al. [2002], Riegler and Douven [2009], Vallinder and Olsson [2012]; compare Douven and Kelb [2011] for a review.

[6] The model was originally developed in order to study the emergence and stability of fundamentalism [cf. Baurmann, 2007, Baurmann et al., 2013].

to everybody, agents mainly consider expert opinions and ignore novices. So, epistemic trust modifies the underlying network on which the opinion dynamics unfold. Now, simulations carried out by Zollman [2010, 2012] suggest that reducing information exchange by curtailing the network may be veritistically valuable (because it helps to maintain favorable diversity and may prevent individual agents from skewing the consensus formation process).[7] At first glance, Zollman's results seem to contrast with our findings (inasmuch as competence ascription may limit information flows while decreasing the overall truth-tracking ability). There is, however, no contradiction because the fact that a reduction of information exchange *may* enhance the truth tracking ability is consistent with the fact such a reduction *may* (in other circumstances) be veritistically detrimental.

The outline of this paper is as follows: We motivate and introduce the formal model of opinion dynamics, employed in this study, in Section 2. The model derives from the Hegselmann-Krause model and includes a representation of the agents' mutual competence attributions (second-order beliefs). Based on the model, we simulate opinion dynamics for three different set-ups (initial and boundary conditions), as described in Section 3. These scenarios are: (S1) attributed competence is independent of objective truth-tracking ability, (S2) some agents gradually adjust their competence attributions in line with objective truth-tracking ability, and (S3) initial competence attribution is correlated with objective truth-tracking ability. Section 4 finally reports the simulation results, compares the three scenarios with the original Hegselmann-Krause model and assesses the veritistic value of the three socio-cognitive set-ups.

2. *The Model*

As the model employed in this study derives from the Hegselmann-Krause model (henceforth "HK-model"), it seems appropriate to restate the template model in the first place before introducing our modifications.

The HK-model describes a group of n agents, $G = \{1, \dots, n\}$. Each agent $i \in G$ possesses a (first-order) belief, which is represented by a real number between 0 and 1 ($x_i \in [0, 1]$).[8] We assume that there is a correct belief, the truth ($T \in [0, 1]$). Whether and how an individual agent changes her

[7] For a review see Zollman [2013].

[8] Some beliefs can be naturally represented as a real number (e.g., beliefs about coastal water temperatures, or beliefs about a fair minimum wage). It's not straightforward to represent sets of binary, propositional beliefs as a single real number. But, instead of translating all propositional beliefs into real numbers, Riegler and Douven [2009] have shown how to extend the HK-model so as to cover complex propositional belief states.

opinion depends on two things: (i) the opinions of other trusted agents and (ii) the agent's ability to track the truth. Ad (i), an agent is assumed to trust and listen to all other agents whose opinions lie within her confidence interval, i.e., are not 'too far off'. The size of the confidence interval is given by the parameter $\epsilon \in \mathbb{R}^+$. We may define, for each agent i at time step t, the subgroup of agents trusted by i as,

$$B_i(t) := \{j \in G : |x_i(t) - x_j(t)| \leq \epsilon\}. \tag{1}$$

Without truth-tracking ability, the first-order belief of an agent at step $t + 1$ is simply the average of the beliefs of all agents she trusted at t. But, ad (ii), the agents may possess the ability to track the truth, which means that their beliefs are not only determined by opinions of trusted agents but by the truth itself, as well. In sum, this gives us the following dynamics (for all agents $i \in G$),

$$x_i(t+1) \;=\; \alpha_i\Big(T\Big) + (1-\alpha_i)\Big(\frac{1}{|B_i(t)|}\sum_{k\in B_i(t)} x_k(t)\Big), \tag{2}$$

The first summand represents the objective component, the second summand the social component of belief formation. The parameter $\alpha_i \in [0, 1]$ describes the agent-specific ability to track the truth and determines the relative weights of the objective and social component. The more proficient an agent is in terms of finding the truth, the less attention she pays to the opinions of her trusted peers. We say that an agent i is an expert if $\alpha_i > 0$, and a novice otherwise.

In the HK-model, the agents possess varying degrees of objective competence (α-values) but don't assign levels of competence or expertise to each other. As we want to investigate the role of competence ascriptions, this is exactly where we extend the model. We assume, accordingly, that each agent i possesses, besides her first-order belief (x_i), n further second-order beliefs that gauge, from i's perspective, the expert status of the group's agents.[9] Thus, the second-order belief $y_{i,j}$ is the level of competence agent i assigns to agent j. This entails that every agent assigns a degree of competence to herself. Like the first-order beliefs, the second-order beliefs are real numbers between 0 and 1 ($y_{i,j} \in [0, 1]$).

Having introduced second-order beliefs, we adjust, in a next step, the conditions under which agents trust each other. It's plausible to assume that an agent i won't trust an agent j in case she considers herself an expert and j

[9] These second-order beliefs correspond, of course, to the weights in the Lehrer-Wagner model [Lehrer and Wagner, 1981].

a novice. Along this line, we stipulate that it is necessary for agent i to trust agent j that i ascribes j a level of competence which is at least as great as the level of competence she assigns to herself. We may redefine, accordingly, for each agent i at time step t, the subgroup of agents trusted by i as,

$$B_i(t) := \Big\{ j \in G : (|x_i(t) - x_j(t)| \leq \epsilon) \wedge (y_{i,j} \geq y_{i,i}) \Big\}. \quad (3)$$

Note that, in contrast to the original HK-model, the modified model allows for asymmetric relations of epistemic trust; i.e., agent i may trust agent j while agent j does not trust agent i.

Like the first-order beliefs in the HK-model, the second-order beliefs in the modified model evolve in time as a function of (i) social and (ii) objective factors. The second-order dynamics can, in fact, be modeled in close analogy to the first-order dynamics. Hence, agent i changes her competence ascription to some agent j in the light of what (i) her trusted agents think about j's competence and (ii) whether j has some objective first-order competence, i.e., the objective ability to track the truth. Formally, the dynamics of our model can be specified as follows (for all agents $i \in G$),[10]

$$\begin{aligned} x_i(t+1) &= \alpha_i\Big(T\Big) + (1-\alpha_i)\Big(\frac{1}{|B_i(t)|}\sum_{k \in B_i(t)} x_k(t)\Big) \quad (4) \\ y_{i,j}(t+1) &= \beta_i\Big(\lceil \alpha_j \rceil\Big) + (1-\beta_i)\Big(\frac{1}{|B_i(t)|}\sum_{k \in B_i(t)} y_{k,j}(t)\Big), \quad (5) \end{aligned}$$

where in both equations the first summand represents the objective component and the second summand represents the social component of belief formation. We retain the first-order updating process of the original HK-model, hence (4) is identical with (2). While, as before, the agent-specific parameter α_i indicates the degree to which an agent is able to track the truth (in terms of first-order beliefs), the parameter $\beta_i \in [0, 1]$ describes the degree to which an agent adjusts her competence ascriptions to (other) agents in light of those agents' truth-tracking ability (i.e., α-value). In other words, α_i and β_i represent an agent's objective first-order and second-order competence, respectively. Note that these two parameters are independent. One may be able to track the truth without seeing which other agents possess this ability; and one may be highly competent in recognizing objective experts without being an expert oneself.

Finally, we may observe that the original HK-model (equations 1,2) is, in a strict sense, a special case of the modified model (equations 3-5). If, e.g.,

[10] Note that $\lceil \alpha_j \rceil = 1$ if $\alpha_j > 0$, and $\lceil \alpha_j \rceil = 0$ if $\alpha_j = 0$.

for all agents i, j (i) $\beta_i = 0$ and (ii) $y_{i,j}(t_0) = 0$, then the modified version reduces to the original one.

3. *The Scenarios*

We simulate opinion dynamics for three different scenarios (initial and boundary conditions) and will, in Section 4, compare the results with those of the HK-model.

The three scenarios as well as the simulations with the HK-model share the following specifications: The epistemic community contains $n = 100$ agents, and the objective truth equals $T = 0.1$. The confidence interval has constant width, $\epsilon = 0.2$. The objective first-order competences (α-values) are determined randomly according to the following parameters: With probability 0.2 (0.8) an agent i displays a competence of $\alpha_i = 0.1$ ($\alpha_i = 0.0$). So, on average, 20% of all agents are truth-trackers. Moreover, the initial first-order beliefs (0.01, 0.02, 0.03, ..., 0.99, 1.0) are regularly distributed over the interval $[0, 1]$. The initial second-order beliefs and the agents' second-order competences (β-values) vary from scenario to scenario (see below).

The three set-ups we distinguish are:

Scenario S1) No objective second-order competence. The initial second-order beliefs are chosen randomly from the interval $[0, 1]$ (uniform distribution); the agents possess no objective second-order competence ($\beta_i = 0$ for all agents $i = 1 \dots n$). In this scenario, the competence ascriptions (second-order beliefs) to some agent are not directly linked to that agent's objective truth-tracking ability.

Scenario S2) Positive second-order competence. The initial second-order beliefs are chosen randomly from the interval $[0, 1]$ (uniform distribution); the agent specific β-values are determined in analogy to, but independently of, the α-values; i.e., with probability 0.2 (0.8) an agent i displays a second-order competence of $\beta_i = 0.1$ ($\beta_i = 0.0$). So, on average, 20% of the agents can track objective first-order expertise, but these agents with $\beta_i > 0$ are not necessarily identical with the agents that display $\alpha_i > 0$. In other words, there is no initial (prior) knowledge about the group members' objective expertise, but some agents gradually learn, in the course of the opinion exchange, which agents do and which don't possess objective first-order competence.

Scenario S3) Prior knowledge of truth-tracking ability. The initial competence ascriptions (second-order beliefs) to other agents are correlated with

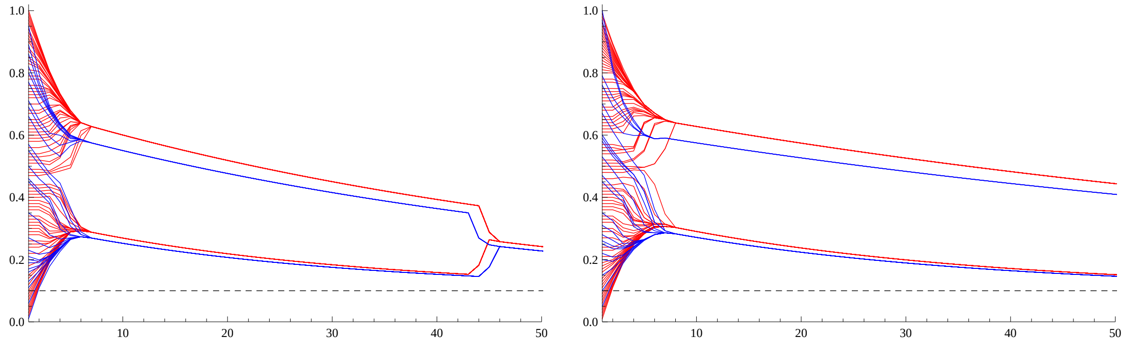

(a) (b)

Figure 1: Illustrative opinion dynamics of the original HK-model. First-order beliefs are plotted against time. The doxastic trajectories of truth-trackers ($\alpha_i > 0$) are colored blue, those of novices ($\alpha_i = 0$) red. The truth ($T = 0.1$) is marked by a dotted line.

> those agents' objective first-order competence (α-values).[11] But the agents possess no objective second-order competence ($\beta_i = 0$ for all agents $i = 1 \dots n$). So in this third scenario, it is at least approximately clear, to all agents and right from the start, who is a truth-tracker and who isn't—and this prior knowledge is reflected in the initial competence ascriptions.

4. *The Results*

Figure 1 displays two illustrative opinion dynamics of the original HK-model, i.e., without second-order beliefs. These dynamics are thoroughly studied in Hegselmann and Krause [2006] and serve as a point of reference for our investigation. The two plots nicely illustrate the key features of the HK-model (given the specific initial and boundary conditions we assume): Partial consensus formation predates the gradual truth rapprochement; the truth-trackers (blue) drag the novices (red) towards the truth; and the agents converge, in the long run, against the truth.

[11] That is, for every agent i, the initial value of $y_{i,j}$ is randomly drawn from the interval $[0.3, 1]$ if agent j is a truth-seeker ($\alpha_j > 0$), and from the interval $[0,0.7]$ otherwise. Consequently, truth-seekers are, on average, ascribed higher competence than non-truth-seekers. More precisely, consider a truth-seeker j and a non-truth-seeker k. Then, for every agent i, the probability that $y_{ij} > y_{ik}$ is,

$$\begin{aligned}
\Pr(y_{ij} > y_{ik}) &= \Pr(y_{ij} > y_{ik} | y_{ik} < 0.3)\Pr(y_{ik} < 0.3) + \Pr(y_{ij} > y_{ik} | y_{ik} \geq 0.3)\Pr(y_{ik} \geq 0.3) \\
&= 3/7 + 4/7\big[\Pr(y_{ij} > y_{ik} | y_{ik} \geq 0.3 \& y_{ij} > 0.7)\Pr(y_{ij} > 0.7) \\
&\quad +\Pr(y_{ij} > y_{ik} | y_{ik} \geq 0.3 \& y_{ij} \leq 0.7)\Pr(y_{ij} \leq 0.7)\big] \\
&= 3/7 + 4/7[3/7 + 1/2 \cdot 4/7] = 41/49 \approx 0.84.
\end{aligned}$$

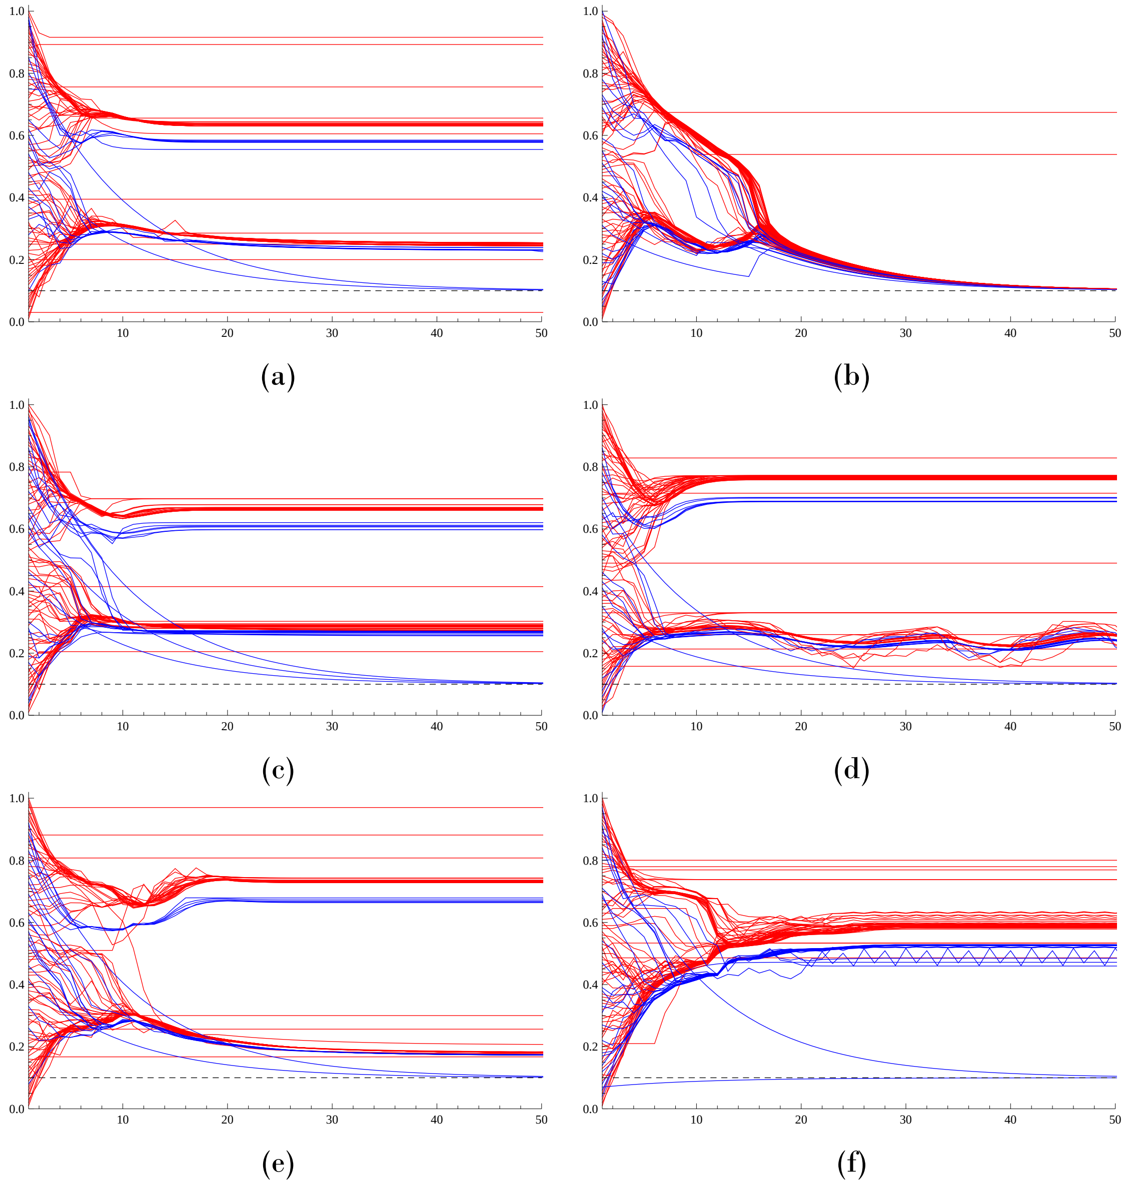

Figure 2: Illustrative opinion dynamics of the modified model in scenario S1). First-order beliefs are plotted against time. The doxastic trajectories of truth-trackers ($\alpha_i > 0$) are colored blue, those of novices ($\alpha_i = 0$) red. The truth ($T = 0.1$) is marked by a dotted line.

Figure 2 displays illustrative opinion dynamics of scenario S1). So, agents ascribe degrees of expertise to each other and only trust other agents that are deemed more competent than oneself. This modification, the plots make clear, heavily influences the opinion dynamics. We may, first of all, observe that, compared to the HK-model, the epistemic communities show a poor veritistic performance. Except in panel (b), where the entire group approaches the truth, only a few agents, and virtually no novices at all, succeed in tracking down the truth. Secondly, there are some red agents that don't alter their position at all (horizontal trajectories) as well as some blue agents which asymptotically approach the truth. These agents are not influenced by the opinions of other agents because they consider themselves as more competent (than all the other agents which fall within their confidence interval) and hence only trust themselves. Note that this doesn't prevent other agents from trusting them. Thirdly, we observe that truth-trackers (blue) which don't approach the truth typically stabilize somewhat below (but within the ϵ-interval of) a group of novices. These truth-trackers are 'caught' between the truth and a subgroup of novices they trust. Unlike in the HK-model, they don't succeed in dragging the novices towards the truth because, coincidentally, the novices do not trust the truth-trackers, that is ascribe lower competence levels to the truth-trackers than to themselves. So, it's because of asymmetric relations of epistemic trust that the agents don't reach the truth. Fourthly, although, in panel (f), the mainstream of the community has roughly agreed on a false opinion (spurious consensus), it seems that there is no full and broad consensus without having reached the truth (cf. panel (b)). In other words, the illustrative plots suggest that full agreement is an indicator of truth-proximity. Fifthly, the illustrative simulations (d) and (f) demonstrate that the modified dynamics do not necessarily stabilize on a fixed point but may exhibit oscillations, which is an interesting and potentially explanatorily powerful feature of the model.

Figure 3 displays illustrative opinion dynamics of scenario S2). Some of the agents possess the ability to learn, in the course of time, which agents are truth-trackers, and which aren't (i.e., $\beta_i > 0$ for some i). In four out of six simulations—panels (a), (c), (d), (e)—, the epistemic community successfully tracks down the truth. So, firstly, second-order learning and expertise recognition (of some agents) seems to improve the collective veritistic performance. Secondly, if a group collectively approaches the truth, it's because the truth-trackers (blue), lying slightly below the novices (red), drag the novices with them; this is a mechanism which has already been observed in the HK-model. The effect of second-order learning is, thirdly, nicely illustrated by trajectories of some individual novices (red), e.g., of the agent 7 in panel (e) who, until step 38, holds a constant first-order belief of $x_7 = 0.07$

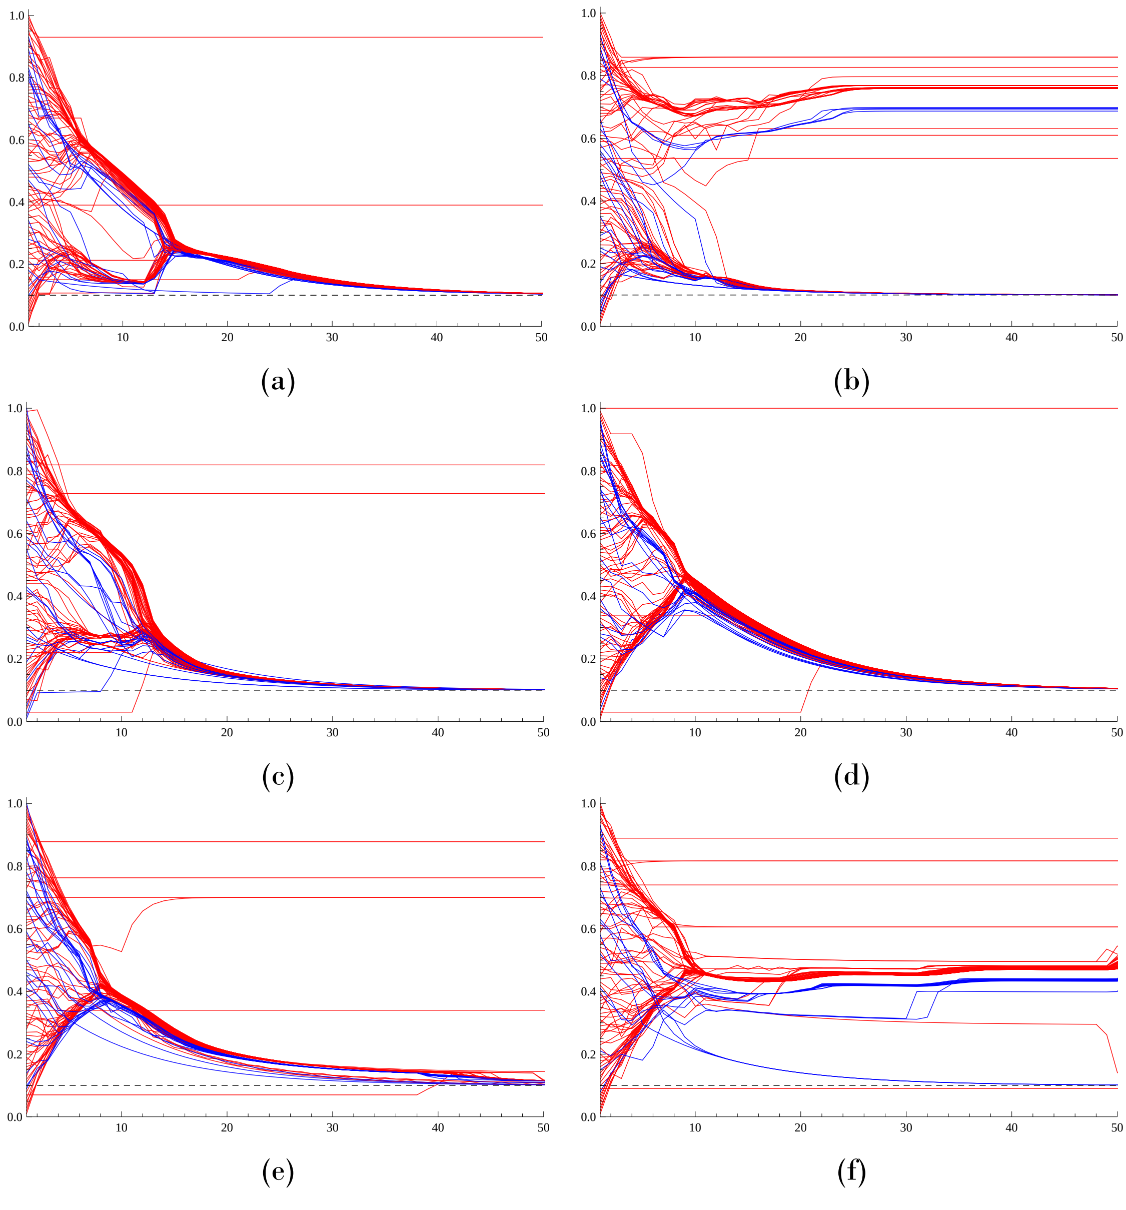

Figure 3: Illustrative opinion dynamics of the modified model in scenario S2). First-order beliefs are plotted against time. The doxastic trajectories of truth-trackers ($\alpha_i > 0$) are colored blue, those of novices ($\alpha_i = 0$) red. The truth ($T = 0.1$) is marked by a dotted line.

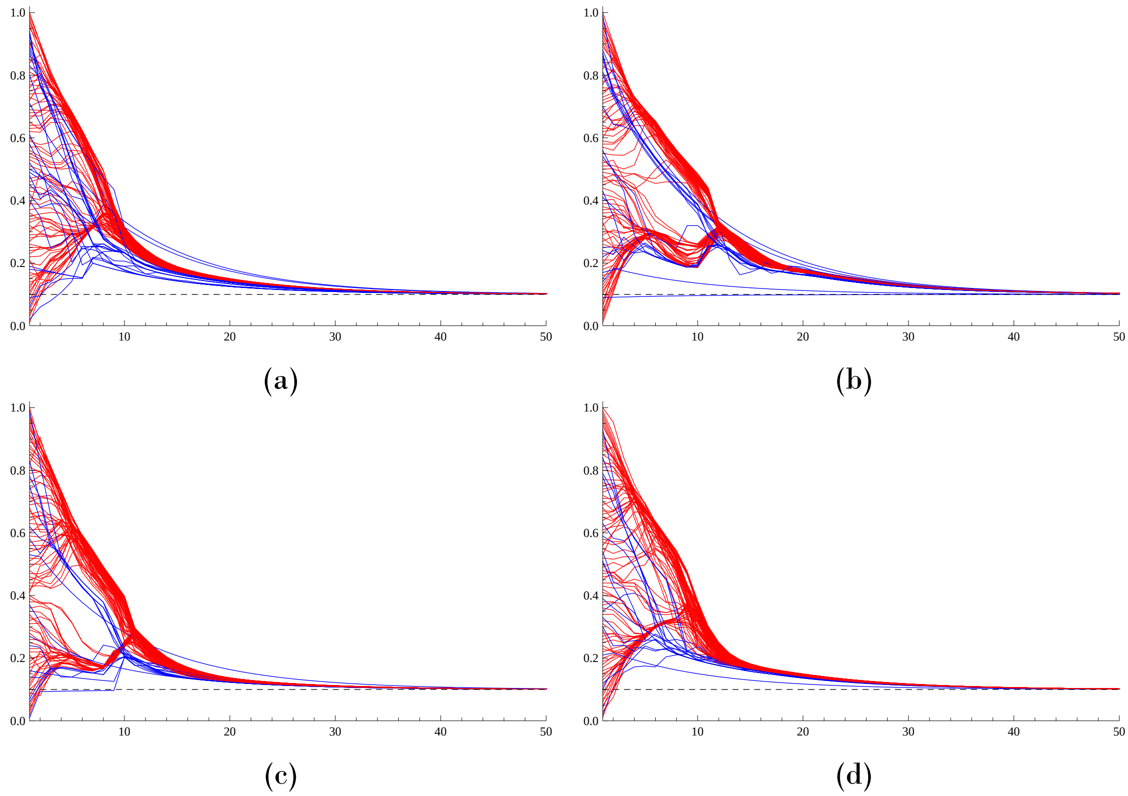

Figure 4: Illustrative opinion dynamics of the modified model in scenario S3). First-order beliefs are plotted against time. The doxastic trajectories of truth-trackers ($\alpha_i > 0$) are colored blue, those of novices ($\alpha_i = 0$) red. The truth ($T = 0.1$) is marked by a dotted line.

(near the bottom of the figure). Although the mainstream, including all the truth-trackers, is within 7's confidence interval at least since step 25, it's not before step 38 that agent 7 trusts some of those mainstream members. That can (only) be explained by the gradual adjustment of 7's second-order beliefs (she is downgrading her competence self-assignment while upgrading the competence attribution to truth-trackers) which, in step 38, leads agent 7 to consider some truth-trackers as more competent than herself, and hence to trust them. Fourthly, panel (f) shows that in scenario S2), as in scenario S1)—cf. Figure 2(f)—, the community's mainstream might agree on an approximate, albeit spurious consensus. That may happen because truth-trackers are not part of the mainstream or because they are not considered sufficiently competent by the mainstream's members. At best, only nearly complete consensus seems to be an indicator of verisimilitude.

Figure 4 displays illustrative opinion dynamics of scenario S3). All agents have imprecise prior knowledge about the group member's expertise and assign initial competence levels accordingly. First and foremost, we may observe that the four simulations exhibit quick consensus formation and (simultaneous) collective rapprochement towards the truth. All agents eventu-

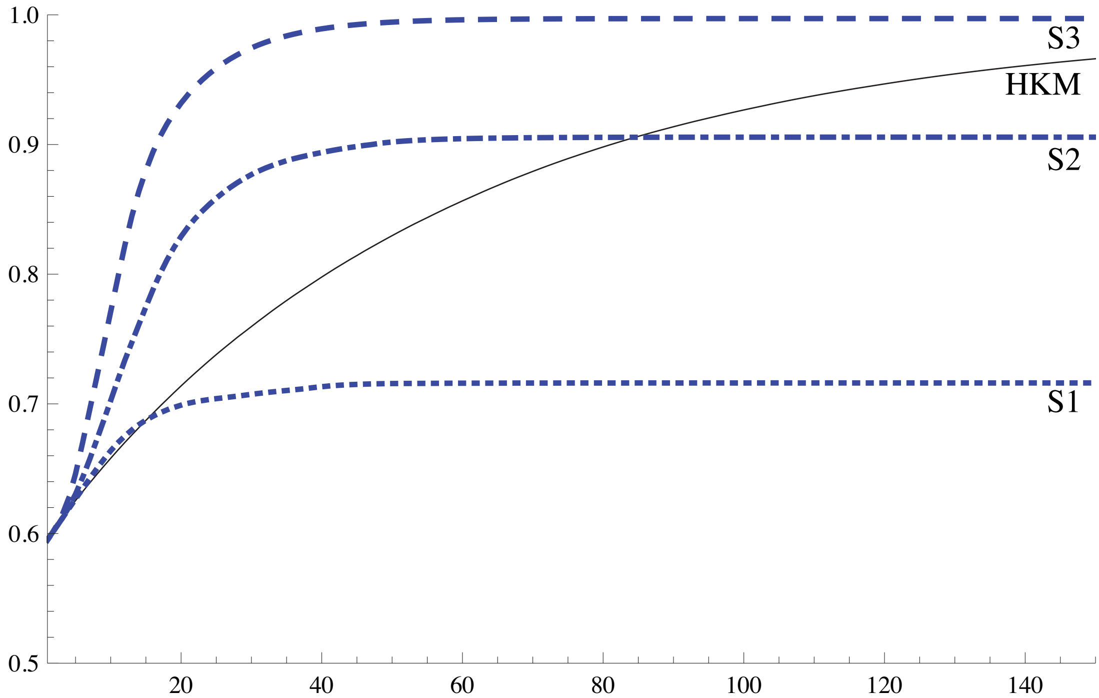

Figure 5: Ensemble-wide mean verisimilitude evolutions of the modified model in scenarios S1), S2), and S3); compared with the Hegselmann-Krause-Model. The mean verisimilitude, averaged over all agents in an ensemble of 100 simulations, is plotted against time.

ally track down the truth.[12] Secondly, the truth-trackers (blue) lie somewhat below the novices (red) and effectively pull them towards the truth; that's, again, the mechanism we know from the HK-model. Thirdly, there are no novices which stubbornly stick to their original position and fail to improve their beliefs. That is because there are no novices who consider themselves more competent than everybody else—on the contrary, novices typically acknowledge the greater expertise of truth-trackers. Some truth-trackers, however, are highly self-confident, only trust themselves and asymptotically approach the truth. Fourthly, experts are not led astray by novices. Again, the initial alignment of expertise ascription on the one side and objective expertise on the other side helps to explain this fact (as well as the complete truth rapprochement): Red agents consider blue agents as more competent right from the start and hence include them in their updating process; blue agents, on the opposite, tend to ignore red agents (for being less competent) and hence are not influenced by their first-order opinions.

Figure 5 visualizes the results of four ensemble simulations, namely of the three scenarios and the HK-model. Each ensemble consists of 100 individual simulations. Figure 5 displays the corresponding ensemble-wide mean

[12] We are only showing 4 plots because they are representative for all the opinion dynamics in scenario S3).

verisimilitude evolutions (where an agent's verisimilitude is defined as one minus her distance to the truth, $1 - |x_i - T|$). So how do the three scenarios and the HK-model compare in terms of verisimilitude increase?

In the short term ($t < 20$), agents approach the truth more rapidly in scenarios S2) and S3) than in the HK-model. In scenario S1), verisimilitude increases roughly at the same pace as in the HK-model. In sum, second-order dynamics (i.e., the practice of ascribing competence levels and forging relations of epistemic trust) are not detrimental in the short term; on the contrary, they may enable a community to approach the truth more rapidly.

In the long run, however, things look different. Consider the benchmark case (the HK-model), first. Here, agents track down the truth (in the long term) and reach a verisimilitude value of 1. Besides in the HK-model, this holds only in scenario S3), where verisimilitude approaches 1, too. In scenarios S1) and S2), in contrast, mean verisimilitude converges against clearly lower values (ca. 0.7 and 0.9, respectively) and the agents are veritistically worse off than in the benchmark case. As a consequence, second-order dynamics can hinder an epistemic community from tracking down the truth in the long run. Even with gradual second-order learning (cf. scenario S2), the collective truth-tracking performance may be lower than without any competence differentiation whatsoever (HK-model). In the different set-ups we have studied, the practice of competence assignment is veritistically valuable only if there is prior knowledge about the community members' objective expertise.

5. *Conclusion*

Before we draw any final conclusions from our investigation we better step back for a moment and reflect on a methodological issue that's looming large. We've used highly simplifying models with idealizing assumptions to study the veritistic value of competence ascription. It would be daring to claim that these models describe real opinion dynamics in an empirically adequate way or that they prescribe how opinions ought to be revised in a rational way. But then, what can we learn from such conceptual models anyway? We suggest to conceive of our model as a credible world in the sense of Sugden [2009]. Accordingly understood models are at least consistent with what we know about opinion dynamics in general and can hence be considered as possibly correct. Following Grüne-Yanoff [2009] and Betz [2010], they may show us that statements we had previously taken for granted might actually be false. We can learn from such models inasmuch as they expose possibilities we hadn't seen or acknowledged before. Models, understood as credible worlds, may also suggest novel hypotheses and hence fulfill a purely heuristic function.

This said, we suggest to draw, from the simulation exercise, a possibilistic conclusion only—albeit a relevant and interesting one. Against deeply entrenched intuitions, the socio-cognitive practice of differentiating amongst agents according to assigned expertise is not necessarily veritistically valuable. In other words, we have shown that epistemic communities might, in some situations, be better off without second-order belief dynamics; not distinguishing between experts and novices might foster such a group's collective ability to track down the truth.

More specifically, we have seen that even the consecutive improvement of competence ascriptions and the gradual alignment of second-order beliefs on the one side and objective expertise on the other side does not necessarily improve the collective truth-tracking ability—compared to the case where one treats all agents as equally proficient. This has further ramifications. The result suggests that prior knowledge about the agents' expertise is needed for the practice of competence ascription to be veritistically valuable. Question is, of course: How does one obtain such knowledge? How do we recognize experts—prior to the actual epistemic engagement? Goldman, who appreciates the importance of this question for veritistic social epistemology, tries to defuse the problem by giving it a temporal dimension [cf. Goldman, 1999, pp. 267-271]: He claims that one can gradually learn who is an expert by comparing the alleged expert's opinion-trajectory with one's own. But our simulations, then, did suggest that such gradual revelation of true expertise might not be enough to make competence ascription veritistically valuable; what we need, or so it seems, is prior expert recognition, and Goldman's proposal is insufficient.

This paper calls for further, more detailed research and raises a couple of significant questions. First of all, we have only considered very specific initial and boundary conditions in this study. But the model promises to display an interesting dynamic (cf. the oscillations) and deserves a systematic study, including, e.g., the exploration of the entire parameter space, the combination of second-order learning and prior expert recognition (scenarios 2 and 3), the introduction of dynamic α- and β-values, etc. A second follow-up question concerns the relation between nearly complete consensus and truth proximity. It is significant to understand whether and, if so, when a consensus reliably indicates verisimilitude; and the model might be used to study this question. A third project would look at the model from an optimization perspective: the task is, then, to identify optimal initial conditions (competence ascriptions) that maximize speed and extent of verisimilitude increase.[13] Fourthly, the agents' evolving second-order beliefs induce an endogenous, dynamic network of relations of trust, on which the opinion dynamics unfold. Given the current research on network epistemology [cf. Zollman, 2013], it seems to be promising to scrutinize the network structures that are generated by the model and to attempt to explain the opinion dynamics in terms of the network dynamics. Fifthly, since the model represents opinion dynamics in a highly aggregated way, it might be insightful to compare its findings with more detailed models of scientific debate.[14] Finally, the results of such simulation studies need to be underpinned by case-studies in the history of science. It seems to us that the well-documented and thoroughly analyzed Great Devonian Controversy [cf. Rudwick, 1985] might represent a worthwhile starting point. An equally promising case is the systematic manipulation of competence ascriptions and its detrimental veritistic effect in tobacco and climate change debates, as documented by Oreskes and Conway [2010].

References

Michael Baurmann. Rational fundamentalism? An explanatory model of fundamentalist beliefs. *Episteme*, 4(2):150–166, 2007.

Michael Baurmann, Gregor Betz, and Rainer Cramm. Charismatische Führer in fundamentalistischen Gruppen. Erklärungshypothesen auf der Basis von Simulationsmodellen. *In preparation*, 2013.

Gregor Betz. What's the worst case? The methodology of possibilistic prediction. *Analyse und Kritik*, 32(1):87–106, 2010.

Gregor Betz. *Debate Dynamics: How Controversy Improves Our Beliefs*. Synthese Library. Springer, Dordrecht, 2012.

Guillaume Deffuant, Frédéric Amblard, Gérard Weisbuch, and Thierry Faure. How can extremism prevail? A study based on the relative agreement interaction model. *Jasss-the Journal of Artificial Societies and Social Simulation*, 5(4):–, 2002.

Igor Douven and Christoph Kelb. Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis*, 75(2):271–283, 2011.

[13] Hartmann and Sprenger [2010] have studied an analogous problem for the Lehrer-Wagner model.
[14] Such as, e.g., Betz [2012].

Alvin I. Goldman. *Knowledge in a Social World*. Oxford University Press, Oxford, New York, 1999.

Alvin I. Goldman. Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63(1):85–110, 2001.

Till Grüne-Yanoff. Learning from minimal economic models. *Erkenntnis*, 70 (1):81–99, 2009.

Stephan Hartmann and Jan Sprenger. The weight of competence under a realistic loss function. *Logic Journal of the IGPL*, 18(2):346–352, 2010.

Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence: Models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3):–, 2002.

Rainer Hegselmann and Ulrich Krause. Truth and cognitive division of labour first steps towards a computer aided social epistemology. *Jasss-the Journal of Artificial Societies and Social Simulation*, 9(3):–, 2006.

Philip Kitcher. The division of cognitive labor. *The Journal of Philosophy*, 87(1):5–22, 1990.

Philip Kitcher. *The Advancement of Science : Science without Legend, Objectivity without Illusions*. Oxford University Press, New York, 1993.

Keith Lehrer and Carl Wagner. *Rational Consensus in Science and Society*. Philosophical studies series in philosophy v. 24. D. Reidel, Dordrecht, 1981.

Ryan Muldoon and Michael Weisberg. Robustness and idealization in models of cognitive labor. *Synthese*, 183(2):161–174, 2011.

Erik J. Olsson. A simulation approach to veritistic social epistemology. *Episteme*, 8(2):127–143, 2011.

Naomi Oreskes and Erik M. Conway. *Merchants of Doubt*. Bloomsbury Press, New York, 2010.

Alexander Riegler and Igor Douven. Extending the Hegselmann-Krause Model III: From single beliefs to complex belief states. *Episteme*, 6(2): 145–163, 2009.

Martin J. S. Rudwick. *The Great Devonian Controversy : The Shaping of Scientific Knowledge Among Gentlemanly Specialists*. Science and its conceptual foundations. University of Chicago Press, Chicago, 1985.

Robert Sugden. Credible worlds, capacities and mechanisms. *Erkenntnis*, 70 (1):3–27, 2009.

Aron Vallinder and Erik J. Olsson. Do computer simulations support the argument from disagreement? *Synthese*, Online First (May 2012), 2012.

Michael Weisberg and Ryan Muldoon. Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2):225–252, 2009.

Kevin J.S. Zollman. The epistemic benefit of transient diversity. *Erkenntnis*, 2010(1):17–35, 2010.

Kevin J.S. Zollman. Social network structure and the achievement of consensus. *Politics Philosophy Economics*, 11(1):26–44, 2012.

Kevin J.S. Zollman. Network epistemology: Communication in epistemic communities. *Philosophy Compass*, 8(1):15–27, 2013.