Extending the Hegselmann-Krause Model II

Alexander Riegler and Igor Douven

ABSTRACT

Hegselmann and Krause have developed a computational model for studying the dynamics of belief formation in a population of epistemically interacting agents who try to determine, and also get evidence concerning, the value of some unspecified parameter. In a previous paper, we extended the Hegselmann-Krause (HK) model in various ways, using the extensions to investigate whether, in situations in which random noise affects the evidence the agents receive, certain forms of epistemic interaction can help the agents to approach the true value of the relevant parameter. This paper presents an arguably more radical extension of the HK model. Whereas in the original HK model each agent is solely characterized by its belief, in the model described in the current paper, the agents also have a location in a discrete two-dimensional space in which they are able to move and to meet with other agents; their epistemic interactions depend in part on who they happen to meet. We again focus on situations in which the evidence is noisy. The results obtained in the new model will be seen to agree qualitatively with the results obtained in our previous extensions of the HK model.

Modern science is largely a community enterprise; scientists working in relative isolation from their colleagues, in the way Kepler or Newton did, are exceptionally rare nowadays. Interactions between working scientists are multiple and multifarious, ranging from jointly carrying out experiments and coauthoring papers to discussing half-baked ideas during coffee breaks. As a result of this, or perhaps as a direct effect of a separate form of epistemic interaction, a scientist's belief on a given matter will often be influenced to at least some extent by his or her colleagues' beliefs on the same matter. But—one may ask—is this for better or for worse? It is not entirely inconceivable that in general scientists would do best, in terms of achieving whatever epistemic goals their research is meant to serve, by going purely by the data, and not allowing their beliefs to be affected by those of any of their colleagues. At a minimum, the said form of interaction might be inessential from a strictly scientific perspective.

In previous work, we have begun studying this matter from a truth-tracking perspective, that is, with an eye towards answering the question whether various forms of adjusting one's belief in response to learning the beliefs of others can help one to achieve the scientific goal of approximating the truth (see Douven, 2009 and Douven and Riegler, 2009). In that work, we could build on pioneering research carried out by Rainer Hegselmann and Ulrich Krause, who in a series of papers have developed a model for investigating the dynamics of belief formation in societies of truth-seeking agents.[1] While the merits of their work are beyond doubt, it has also been observed that their model has severe limitations. Especially if one wants to use it to study the aforementioned questions, many of the assumptions of the Hegselmann-Krause (HK) model cannot but strike one as being too idealized. In our own earlier work, we have sought to take some first steps towards "concretizing" the HK model (in the sense of Nowak 1980), that is, we have proposed extensions of the model that do away with some of the idealizations inherent in it. Most importantly perhaps, we gave up the assumption that the agents in the model receive "noise-free" experimental data. The resulting model enabled us to argue that particular ways of responding to learning the beliefs of certain colleagues are, from a truth-tracking viewpoint, clearly preferable as an epistemic strategy to ignoring those beliefs and going purely by the data, at least in environments in which the data are noisy—as they tend to be in reality.

Still, many more limitations of the HK model remain to be addressed if it is to inform us about the value (or otherwise) of the said form of epistemic interaction as it is to be observed in the practice of science. One of the major impediments in this respect is that, in the HK model, all agents are supposed to be privy to the beliefs of all the other agents at any time. Due to this, the model would seem to apply to only a very limited range of actual situations. Perhaps it is true that the members of a research group are mostly aware of each other's beliefs (insofar as these are relevant to their ongoing research). But epistemic interaction is not restricted to members of one and the same research group, and it would certainly be false to suppose

[1] See Hegselmann and Krause (2002, 2005, 2006).

that scientists are generally aware of the beliefs of all others—members of different research groups included—working in their field. To overcome this problem, we below present a further extension of the HK model in which agents are only aware of the beliefs of the subgroup of agents who they happen to meet at a given time. Specifically, and in contrast to the HK model, we endow agents with a location in a discrete two-dimensional space in which they are able to move about and to meet with other agents. These meetings determine, at least partly, with whom they epistemically interact. We examine various aspects of this new model. Again, we also focus on situations in which the agents receive noisy data. It will be interesting to see that the results obtained in the new model agree, at least qualitatively, with the results obtained in our previous extensions of the HK model.

1. The Hegselmann–Krause Model and Beyond

The HK model was devised to investigate the dynamics of belief formation in populations of epistemically interacting truth-seeking agents. In the model, a population consists of n agents, where each agent i ($1 \leqslant i \leqslant n$) holds a belief $x_i(t)$ at any time step t. The agents try to ascertain the value τ of a given parameter (which could be the mass of some particle, the viscosity of a fluid, the probability of life on Mars, etc.). It is assumed that $\tau \in [0,1]$, and that the agents know this antecedently; thus $x_i(t) \in [0,1]$ for all i and t. At discrete time steps, the agents simultaneously update their beliefs, where the following rule gives the belief of agent i at time step $t+1$ as a function of the agents' beliefs at t and the true value of the parameter:

$$x_i(t+1) \;=\; \alpha \frac{\sum_{j \in X_i(t)} x_j(t)}{|X_i(t)|} + (1-\alpha)\tau. \qquad (1.1)$$

Here, $|X_i(t)|$ is the cardinality of $X_i(t)$, which is defined to be the set of agents whose beliefs are "close enough" to the belief of agent i at t, or, in Hegselmann and Krause's terminology, that are within i's *bounded confidence interval* (BCI) at t. More exactly, $X_i(t) = \{\, j\colon |x_i(t) - x_j(t)| \leqslant \varepsilon \}$, for some real-valued ε. If $j \in X_i(t)$, we shall say that j is an *epistemic neighbor* of i at t; note that, trivially, every agent is its own epistemic neighbor at any time. Furthermore, $\alpha \in [0,1]$ is a weighting factor that determines the degree to which the agents' beliefs depend on those of their epistemic neighbors. In short, at each time step the new belief of an agent is calculated as the weighted average of the average of the beliefs of all its epistemic neighbors and the true value of the parameter. We shall say that an agent *talks to its epistemic neighbors* iff both $\alpha > 0$ and $\varepsilon > 0$. We might thus say that, if in addition $\alpha < 1$, an agent's new belief is the outcome of the combination of talking to its epistemic neighbors and making experiments, where the results of the experiments point in the direction of τ. In their papers, Hegselmann and Krause present some analytic results about this model, but for the most part they explore its properties by means of computer simulations.[2] In particular, they study the relation between the various parameters of the model and the convergence of the agents' beliefs to τ.

In our earlier-cited papers, we noted that the assumption, inherent in the HK model, that the agents receive precise evidence pointing in the direction of τ is not particularly realistic; in reality, scientists have to live with measurement errors and other factors that make their data noisy. For this reason, the assumption was dropped in the extension of the HK model presented in Douven (2009); this model explicitly allows data to be noisy. A further assumption of the original HK model that seemed questionable to us was that the agents' beliefs all weigh equally heavily in the updating process; the model is thereby unable to reflect the fact that in scientific practice the beliefs of some count for more than those of others. These considerations led us in Douven and Riegler (2009) to propose a further extension of the HK model that replaces (1.1) by

$$x_i(t+1) \;=\; \alpha \frac{\sum_{j \in X_i(t)} x_j(t) w_j}{\sum_{j \in X_i(t)} w_j} + (1-\alpha)\big(\tau + \mathrm{rnd}(\zeta)\big). \qquad (1.2)$$

In this equation, w_j denotes the fixed reputation of agent j, and $\mathrm{rnd}(\zeta)$ is a function returning a unique uniformly distributed random real number in the interval $[-\zeta, +\zeta]$ each time it is invoked ($\zeta \in \mathbb{R}$).

The main result of the computer simulations that we used to investigate this model can be put thus: in situations in which evidence obtained in experiments is noisy, populations of agents that attribute more weight to talking to each other (i.e., that have a higher value for α) converge to τ more accurately, albeit more slowly, than do populations that give less weight to talking and rely more on the evidence. See Douven (2009) and Douven and Riegler (2009) for the details. Reputation, on the other hand, appeared to have little to no effect on the outcome of the simulations. See Douven and Riegler (2009) for the details.

While we thus did away with some limitations of the HK model, it still holds for the models described in the aforementioned papers that all agents

[2] For the many virtues of this approach for investigating belief dynamics in multi-agent systems, see Humphreys (1991), Epstein and Axtell (1996), Hartmann (1996), Gaylord and D'Andria (1998), and Winsberg (1999), as well as the papers cited in the previous note.

are supposed to know, at any time, the beliefs of all the other agents. As intimated, this reduces the scope of the model considerably; the number of colleagues with whom scientists collaborate, or whom they meet at conferences, is generally just a fraction of the number of their epistemic neighbors. In the following, we describe an extension of the HK model that deviates more drastically from the original and that eliminates the said unrealistic assumption by letting agents move about in a two-dimensional environment in which they encounter others and interact only with the epistemic neighbors they happen to meet.

2. Adding Spatial Dimensions

In order to accurately model the aspect of encountering epistemic neighbors, we extended the HK model by introducing *spatial dimensions*. More concretely, each agent (exclusively) populates a site in a discrete two-dimensional toroidal grid, that is, a grid whose outside borders wrap around to the opposite side in order to avoid edge effects. (Such effects would, for example, reduce the number of possible adjacent neighbors for any agent at an edge, or in a corner, of the grid.) Our usage of discrete two-dimensional environments is neither original nor arbitrary, as they have long proven to be valuable tools in disciplines such as artificial life and sociology (see, e.g., Epstein and Axtell, 1996 and Gaylord and D'Andria, 1998).

To be more specific, the environment consists of a two-dimensional grid of 25×25 sites each of which can either be empty or occupied by an agent. Agents face one of the four cardinal points of the compass. They move about by leaping to the adjacent free site they are facing. As noted above, unlike in the original HK model and our earlier extensions thereof, in which an agent at each time step interacts with *all* its epistemic neighbors, in the present model it interacts only with those of its epistemic neighbors that are to be found in its *spatial* neighborhood. That is to say, for all i, j, and t, agent i epistemically interacts with agent j at time t iff:

1. j is in the *epistemic* neighborhood of i, that is, $|x_i(t) - x_j(t)| \leqslant \varepsilon$, and
2. j is in the *spatial* neighborhood of i.

The notion of spatial neighborhood can be—and in the literature has been—defined in various ways. In our simulations, we made use of the three neighborhood structures depicted in Figure 1, which are relatively common in the literature (see, e.g., Gaylord and D'Andria, 1998). Following a suggestion of Gaylord and D'Andria (1998), we also considered a variation of the von Neumann neighborhood in which only those agents are considered that face

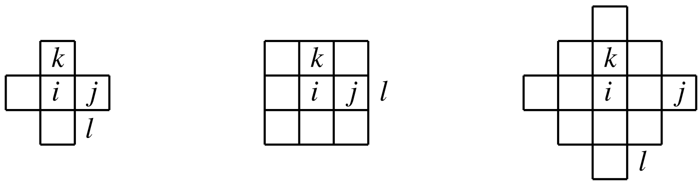

Figure 1: Three neighborhood structures: von Neumann (left); Moore (center); Gaylord-Nishidate (right). In each case, j and k are spatial neighbors of i, whereas l is not

one's own position (for the other two neighborhoods this extra "facing" condition makes no sense).

In accordance with the distinction between epistemic and spatial neighborhoods, the development of agents through time is characterized by both a belief-update rule and a migration rule. As in the HK model, the update rule takes care of the belief transitions an agent can undergo from one time step to another. The migration rule determines how agents move about in their environment. The former looks thus:

$$x_i(t+1) \;=\; \begin{cases} \alpha \frac{\sum_{j \in X_i(t)} x_j(t)}{|X_i(t)|} + (1-\alpha)\big(\tau + \mathrm{rnd}(\zeta)\big) & \text{if } |X_i(t)| > 1, \\ x_i(t) & \text{otherwise,} \end{cases} \tag{2.1}$$

where $X_i(t)$ now designates the set of i's epistemic neighbors at t that are also within its spatial neighborhood at that time. Note that the upper clause of equation (2.1) corresponds to equation (1.2) with reputation $w_j = 1$ for all j. This clause is invoked whenever there is at least one epistemic neighbor present in the spatial neighborhood besides the agent itself. If $|X_i(t)| = 1$, agent i's belief remains unchanged from t to $t+1$. The migration rule simply says that after an agent has updated its belief, it moves one step to the adjacent site it faces if that is free and not faced by at least one other agent, or else it randomly changes its orientation to any of the four cardinal directions (see Gaylord and D'Andria, 1998 for details).

3. Exploring the Model

In this section, we present the results of computer simulations we conducted in order to explore systematically the properties of our two-dimensional model. Like Hegselmann and Krause in the investigations of their model, we were particularly interested in questions concerning the relation between the values for the parameters α and ε and a population's ability to track the truth. First, however, we wanted to be clear about the role the type of the spatial neighborhood plays in this regard.

3.1. Different Spatial Neighborhoods

Given the different definitions of the notion of spatial neighborhood, one may ask which of these yields the best results in terms of accuracy and speed of convergence of the agents' beliefs to the value of τ. We put the four types described in the previous section to test by comparing two populations of $n = 100$ agents that attribute different weights to talking ($\alpha = .1$ and $\alpha = .9$, respectively). For both these populations, it holds that $\zeta = .1$ and $\varepsilon = .1$. We ran 100 simulations for each of the two populations and each of the four spatial neighborhoods, yielding eight trials of 100 simulations in total. In these, as in all other simulations we performed for the present paper, we set $\tau = .75$. For each trial, we calculated, at each time step, the average over the 100 simulations of the average distance from the truth of the beliefs of all the agents in the population, where the distance from the truth of agent i's belief at t was simply taken to be $|x_i(t) - \tau|$.

The results are shown in Figure 2. The most accurate and also fastest convergence is achieved by the Gaylord-Nishidate neighborhood (solid line), followed by Moore (dashed), von Neumann (dotted), and von Neumann facing (dash-dotted). This is so regardless of whether the agents give much weight to the evidence (low value for α) or rather to talking to others (high value for α), although in the former case the differences between the various neighborhoods are less pronounced. We opted for the Gaylord-Nishidate neighborhood in all further experiments. That the

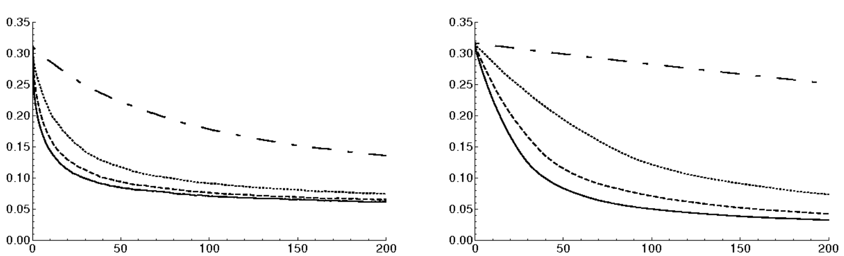

Figure 2: Convergence in different neighborhoods. The abscissa represents time steps, the ordinate is the average distance from τ. Left: $\alpha = .1$; right: $\alpha = .9$

Gaylord-Nishidate neighborhood, which encompasses more surrounding sites than any of the other neighborhoods, is superior in terms of convergence to the truth can already be considered as a first indication that, at least if the evidence the agents receive is noisy, the present model makes talking to others come out as a good epistemic strategy.

3.2. The Weight of Talking

The foregoing would clearly square with the earlier-cited conclusion of Douven and Riegler (2009) that, in situations in which evidence is noisy, giving much weight to talking to one's epistemic neighbors is an effective scientific strategy. But does this conclusion really carry over to the two-dimensional model? Given that in this model most of one's epistemic neighbors may be unavailable for talking at any given time step, the answer is not straightforward.

Talking is not a yes/no affair in the original HK model, in our previous extensions of that, or in the two-dimensional model. "How much" talking is going on in a population depends in all these models on the values of α and ε. To investigate the impact of varying these values in the two-dimensional model, we first systematically increased the value of ε between 0 and .3, in steps of .05, in two populations, one with $\alpha = .1$, the other with $\alpha = .9$. We ran 100 simulations with a population of $n = 100$ agents for each combination of α and ε. Interpolated results for the entire range of ε-values are shown in Figure 3.

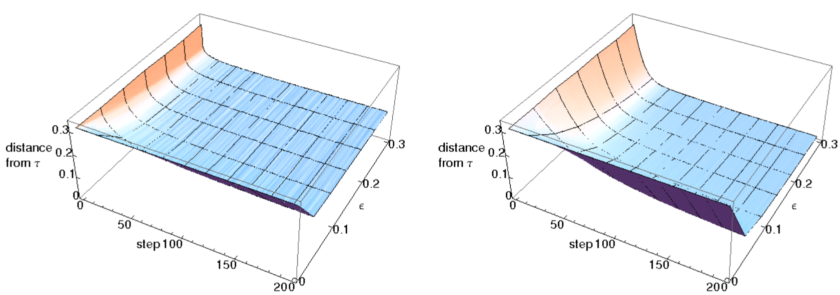

Figure 3: Varying ε for $\zeta = .3$. Left: $\alpha = .1$; right: $\alpha = .9$

The results suggest that, at least for the given level of noise, the larger ε the more accurate the convergence. Still, the value of α is seen to matter as well. For populations of agents that assign a value of only .1 to α, convergence to τ is far less accurate than for populations of agents that assign a value of .9 to α.

For additional clarity, Figure 4 shows the convergence to τ for only four selected combinations of values for α and ε. It indicates clearly the advantage of assigning higher values to both α and ε for approaching τ. (We

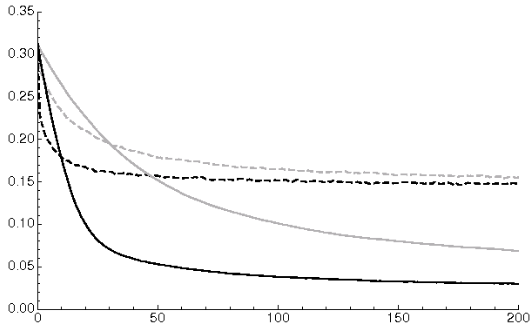

Figure 4: Dashed lines: $\alpha = .1$; solid lines: $\alpha = .9$. Gray lines: $\varepsilon = 0$; black lines: $\varepsilon = .3$

also ran simulations for $\zeta = .1$ and $\zeta = .2$, but this did not lead to a qualitative difference in the outcome. Quantitatively, the differences between the different parameter settings are a bit more pronounced with $\zeta = .3$, which is why we used this value for producing the graphs.)

In further computer experiments, we systematically increased the value of α between 0 and 1, again in steps of .05, keeping ε fixed at .2.[3] In these experiments, it appeared that the closest approximation to the truth was reached if α was given a value close to 1, as can be gleaned from Figure 5. While a high value for α delays the convergence, convergence is ultimately more accurate, that is, it leads to a smaller average distance of the agents' beliefs from τ. Based on these results, it would seem that scientists interested in quick approximate solutions should give less weight to talking than to the evidence they obtain, whereas scientists who want to approach τ as closely as possible—which from a purely scientific perspective would seem to be the more desirable goal—are well advised to do the opposite and rather give less weight to the evidence and more to talking to others.

[3] Again we used various values for ζ, and again this did not significantly influence the results. We also again ran 100 simulations for each relevant combination of the parameters. All experiments to be reported below have the same set-up.

This is fully in line with the above-mentioned conclusion from Douven and Riegler (2009) about our simpler extensions of the HK model.

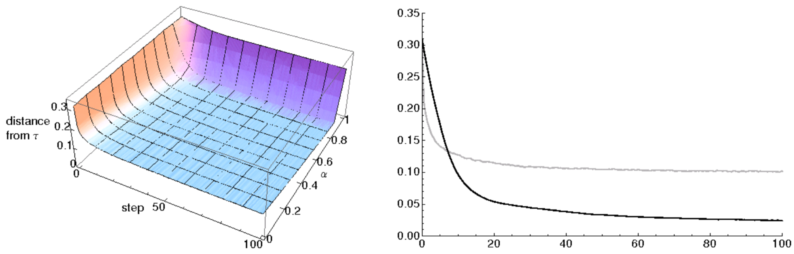

Figure 5: Varying α for $\zeta = .2$. Right: comparison of $\alpha = .1$ (gray) and $\alpha = .9$ (black)

3.3. Elite Versus Normal Scientists

Following Douven and Riegler (2009), we also in the two-dimensional model performed experiments in which we tried to model a distinction between elite scientists, who are very skilled experimenters, able to obtain noise-free data, and normal scientists, who are not equally skilled and obtain noisy data. To that end, we introduced two classes of agents that differ with respect to their noise parameter: of the 100 agents in each population considered in the simulations, there are 75 normal agents with a noise parameter $\zeta_N = .2$ and 25 elite agents with $\zeta_E = 0$. For ε we chose a value of .1 and for α we experimented with values between 0 and 1. Figure 6 shows the results for selected values of $\alpha = .1$ and $\alpha = .9$.

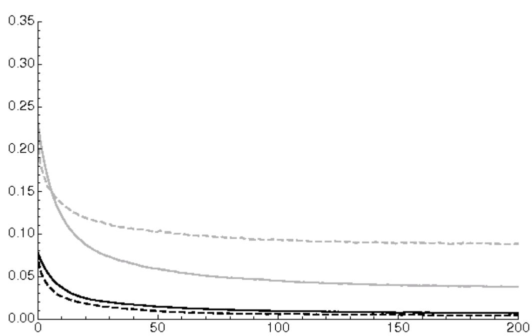

Figure 6: Black lines: elite agents; gray lines: normal agents. Dashed lines: $\alpha = .1$; solid lines: $\alpha = .9$

These results exhibit that normal agents (gray lines) benefit greatly if all agents—elite and normal ones alike—give much weight to talking, as they (the normal ones) then approach τ much more closely than they do if all agents give only little weight to talking. On the other hand, elite agents (black lines) are better off in the latter case. Nevertheless, since they do not do *much* worse in that case, a kind of epistemic utilitarianism would suggest that, provided the choice is between just these two options—either all agents assign a high value to α or all agents assign a low value to α—the former is preferable.

3.4. The Role of Reputation

Not only are some scientists better than others, some scientists are also more reputed than others. In many situations, the views of the highly reputed scientists weigh more heavily than the views of others. But *should* they do so? Might this aspect of scientific practice not bear *negatively* on the prospects for making progress? Perhaps scientists are generally bad at distinguishing the good from the not-so-good scientists, and tend to assign a higher reputation to the latter than to the former. This question can be investigated, to some extent, in our two-dimensional model by replacing the update rule (2.1) by

$$x_i(t+1) = \begin{cases} \alpha \frac{\sum_{j \in X_i(t)} x_j(t) w_j}{\sum_{j \in X_i(t)} w_j} + (1-\alpha)\big(\tau + \mathrm{rnd}(\zeta)\big) & \text{if } |X_i(t)| > 1, \\ x_i(t) & \text{otherwise,} \end{cases} \tag{3.1}$$

where, as in equation (2.1), $X_i(t)$ designates the set of agents that are both i's epistemic and its spatial neighbors at t, and where, as in equation (1.2), w_j represents the reputation of agent j.

The specific question we wanted to address was whether, in order to improve the accuracy in truth approximation of a population in its entirety, elite scientists should be given a higher reputation than normal scientists, where the difference between these groups is understood in the way defined above. To investigate this question, we set up three scenarios. In the first, the agent's status, *qua* elite ($\zeta_E = 0$) or normal ($\zeta_N = .2$), corresponds to its reputation: elite agents are given a reputation w_j of 2, normal agents a reputation of 1. In the second scenario, this is reversed: elite agents are given a reputation of 1, normal agents a reputation of 2. And in the third scenario, both the elite and the normal agents have a reputation of 1.

In line with the findings of Douven and Riegler (2009), the results of various computer experiments (using various combinations for the values of the other parameters) suggest that there is no difference among the three settings, neither in terms of accuracy of convergence, nor in terms of speed of convergence. As also explained in the aforementioned paper, this should not be taken to support the cynical conclusion that, from a purely scientific perspective, it is immaterial whether we assign a high reputation to the best scientists or rather to the crackpots: reputation plays a role in science in many ways, most of which go unaccounted for in the present model. Nevertheless, it is noteworthy that in at least one way in which one would have guessed reputation to matter, it does *not* matter.

3.5. Choosing Between Talking and Experimenting

In the computer experiments described above, we have been assuming that whenever agents talk to other agents they also take into account evidence pointing in the direction of τ. Clearly, this is an idealization. Human scientists may perform an experiment one day and share their results only the next day. One can model this to some extent by letting each agent at each time decide which one of the following it wants to do: (i) gather evidence pointing in the direction of τ by performing an experiment (or making observations, or some such); (ii) talk to those of its epistemic neighbors that are in its spatial neighborhood at the given time (if any are). For this purpose, we once more changed the update rule, as follows:

$$x_i(t+1) = \begin{cases} \alpha x_i(t) + (1-\alpha)\big(\tau + \mathrm{rnd}(\zeta)\big) & \text{if } |X_i(t)| > 1 \text{ and } \mathrm{rnd}(1) > \theta, \\ \dfrac{\sum_{j \in X_i(t)} x_j(t)}{|X_i(t)|} & \text{if } |X_i(t)| > 1 \text{ and } \mathrm{rnd}(1) \leqslant \theta, \\ x_i(t) & \text{if } |X_i(t)| = 1. \end{cases} \tag{3.2}$$

Here, $|X_i(t)|$ again designates the cardinality of the set of agents that are both within i's epistemic and within its spatial neighborhood at t; the function $\mathrm{rnd}(1)$ returns a uniformly distributed random real number in the interval $[0,1]$; and θ is a threshold influencing the agent's decision: a low value of this threshold corresponds to a greater likelihood of updating one's belief on the evidence and one's current belief only, a high value to a greater likelihood of updating by talking to others only.

Our results show a clear, even if not very significant, difference in the ability to approach τ for different values of θ. As can be seen in Figure 7, convergence is more accurate, though slower, for higher values of the threshold, that is, for a greater likelihood of talking (the graph shows the results for $\varepsilon = .1$ and $\alpha = .1$; the results proved relatively robust for variations in these values).

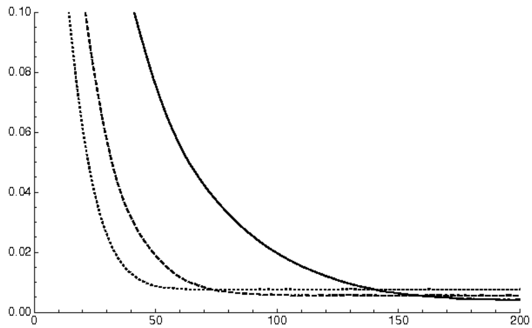

Figure 7: Dotted line: $\theta = .25$; dashed line: $\theta = .5$; solid line: $\theta = .75$

4. Conclusion

Hegselmann and Krause's model must be considered an important tool for studying the effects of certain types of epistemic interaction in populations of truth-seeking agents on these populations' abilities to track the truth. While their own model is highly idealized, it can easily be extended in ways that allow for more realistic interpretations. Above, we have presented an extension that introduces spatial neighborhoods, whereby one can drop the assumption of the original HK model that each agent knows at all times the beliefs of all other agents. The results about this extended model, which were obtained by means of computer simulations, reconfirm the main conclusion of Douven and Riegler (2009): assuming the evidence to be noisy, as in real life it tends to be, talking to others is an epistemically good strategy for a population of truth seeking agents in that it helps the agents to get, on average, closer to the truth, closer, at any rate, than if they disregard the beliefs of others and purely go by the data and their own current beliefs.

In closing, we note that the model presented above, while already much less idealized than the original HK model, is still rather limited in scope, as in it all agents are supposed to hold single beliefs only. Real scientists invariably have a great many beliefs that, moreover, tend to be logically interconnected. A model that equips the agents with propositional theories rather than a single numerical belief is described, and systematically explored, in Douven and Riegler (2009).[4]

[4] We are grateful to Christopher von Bülow for valuable comments on a draft version of this paper.

References

Douven, I. (2009). Simulating peer disagreements. (Manuscript).

Douven, I. and Riegler, A. (2009). Extending the Hegselmann-Krause model i. (Manuscript).

Epstein, J. and Axtell, R. (1996). *Growing Artificial Societies*. MIT Press, Cambridge, MA.

Gaylord, R. J. and D'Andria, L. J. (1998). *Simulating Society*. Springer, New York.

Hartmann, S. (1996). The world as a process: Simulations in the natural and social sciences. In R. Hegselmann, U. Mueller, and K. G. Troitzsch (Eds.), *Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View* (77–100). Kluwer Academic Publisher, Dordrecht.

Hegselmann, R. (1996). Cellular automata in the social sciences. In R. Hegselmann, U. Mueller, and K. G. Troitzsch (Eds.), *Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View* (209–233). Kluwer Academic Publisher, Dordrecht.

Hegselmann, R. and Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis, and simulations. *Journal of Artificial Societies and Social Simulation*, 5. available at URL = http://jasss.soc.surrey.ac.uk/5/3/2.html.

Hegselmann, R. and Krause, U. (2005). Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25, 381–405.

Hegselmann, R. and 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 URL = http://jasss.soc.surrey.ac.uk/9/3/10.html.

Humphreys, P. (1991). Computer simulations. In A. Fine, M. Forbes, and L. Wessels (Eds.), *Proceedings of the PSA 1990*, volume 2 (497–505). East Lansing MI. Philosophy of Science Association.

Nowak, L. (1980). *The Structure of Idealization*. Reidel, Dordrecht.

Riegler, A. and Douven, I. (2009). Extending the Hegselmann-Krause model iii: From single beliefs to theories. *Episteme*. (In press).

Winsberg, E. (1999). Sanctioning models: The epistemology of simulation. *Science in Context*, 12, 275–292.