

Extending the Hegselmann–Krause Model I

Igor Douven

Institute of Philosophy, University of Leuven
igor.douven@hiw.kuleuven.be

Alexander Riegler

Institute of Philosophy, University of Leuven
alexander.riegler@hiw.kuleuven.be

Abstract

Hegselmann and Krause have developed a simple yet powerful computational model for studying the opinion dynamics in societies of epistemically interacting truth-seeking agents. We present various extensions of this model and show their relevance to the investigation of socio-epistemic questions, with an emphasis on normative questions.

Agent-based computer simulations have been successfully employed in the study of broadly social processes for over thirty years now. Researchers have investigated by these means such diverse phenomena as the spread of wealth in a population, the emergence and evolution of friendship networks, integration and segregation of different racial groups, the transmission of cultural values, and the propagation of infectious diseases.¹ It is only in the past decade that agent-based simulations have come to be used for studying distinctively socio-epistemic (or socio-doxastic) questions, such as, most notably, questions concerning the roles various types of social interaction play in the acquisition and transmission of beliefs and knowledge. Pioneering work in this area has been carried out by Rainer Hegselmann and Ulrich Krause, who have developed, and systematically investigated, a simple yet already quite powerful model for simulating the opinion dynamics in societies of epistemically interacting truth-seeking agents.² By thus initiating what they have termed “Computer Aided Social Epistemology,” they have, in our eyes, made an important contribution to social epistemology. After all, social epistemology can be considered to be as much a branch of philosophy as of social science, and one may thus expect that it is able to profit from the deployment of simulation tools no less than the more traditional social sciences have been (and still are).

Nevertheless, it has been generally recognized that, in addition to considerable virtues, the Hegselmann–Krause (HK) model also has some important inherent limitations, these being largely due to the fact that the model is starkly idealized in several respects. The present paper proposes various extensions of the model that do away

¹See Epstein and Axtell [1996] and the references given there.

²See, e.g., Hegselmann and Krause [2002], [2005], and [2006]; see also Dittmer [2001], Fortunato [2004], and Lorenz [2007]. Similar work is to be found in Deffuant et al. [2000], Weisbuch et al. [2002], and Ramirez-Cano and Pitt [2006].

with some of the idealizations. While extending a given simulation model is typically not difficult, nor automatically significant, we hope to show that the extensions of the HK model to be presented below are valuable in that they are better suited than the original model is for the investigation of certain questions—especially normative ones—that are central to social epistemology. We note already here that, in this paper, the emphasis will be on the presentation of the various extensions and on highlighting their relevance to social epistemology; although we state some basic properties of the extensions, a systematic investigation of them must await another occasion.

1. The Hegselmann–Krause Model. Hegselmann and Krause use computer simulations to study the opinion dynamics of communities of agents who are individually trying to determine the value of a certain real-valued parameter. The agents are supposed to know that the true value of the parameter lies within a given interval, which for computational convenience is assumed to be the half-open interval $(0, 1]$. The simulations start by picking a given number of random values in that interval, which represent the initial opinions of as many agents. The agents then update their opinions repeatedly and simultaneously at discrete time-steps by taking a weighted average of, on the one hand, the (straight) average of the opinions that are “close enough” to their own and, on the other hand, of the true value of the parameter they are trying to determine. To make this precise, let $\tau \in (0, 1]$ be the true value of the parameter, $x_i(t)$ the opinion of agent x_i at time t , and $\alpha_i \in [0, 1]$ the weighting factor used by agent x_i . Further define a neighborhood function $X_i(t) := \{j : x_i(t) - \delta_i \leq x_j(t) \leq x_i(t) + \eta_i\}$, with $\delta_i, \eta_i \in [0, 1]$, and let $|X_i(t)|$ be the cardinality of $X_i(t)$. Then the opinion of x_i at $t + 1$ is given by

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

If $j \in X_i(t)$, we shall call x_j a *neighbor* of x_i at t . We shall say that an agent x_i *talks to its neighbors* (or sometimes even just that it *talks*) iff both (i) $\alpha_i > 0$ and (ii) either $\delta_i \neq 0$ or $\eta_i \neq 0$ (or both). Naturally, the situation might occur in which the agent does not have any genuine neighbors to talk to, that is, there might be times t such that the neighborhood $X_i(t) = \{x_i\}$.

As the role of τ in (1.1) may be easily misunderstood, it merits remark that the idea of this equation is *not* that the agents know the true value of the parameter; if they did, then we might expect them to immediately adopt it as their new opinion. The idea, rather, is that an agent gets information at each time t , for instance by performing experiments, that points in the direction of τ , and which, together with the opinions of its neighbors at t , determines its opinion at $t + 1$ in a way captured by (1.1). The rule or rules the agents use to accommodate the data are left implicit by Hegselmann and Krause (see their [2006, Sect. 1] for more on this).

Among the questions they address are then the following: Under which circumstances do communities of agents who update by (1.1) reach a consensus? Under which circumstances do they polarize? How do the values of δ_i and η_i relate to the evolution of the opinions of the agents? Does it matter whether these are equal, or whether they are shared by all agents? How does the speed of convergence to the truth depend on the weight the agents assign to the opinions of others relative to the one they assign to the truth, and does it matter whether they all assign the same

weights? How does speed of convergence depend on the “location” of the truth in the $(0, 1]$ interval? And so forth.

These and similar questions receive detailed, often even exhaustive, answers in Hegselmann and Krause’s writings. Before giving a few snapshots of their findings, we should note that in many of their simulations Hegselmann and Krause assume that there is an ε such that $\delta_i = \varepsilon = \eta_i$ for all i , and also that there is an α such that $\alpha_i = \alpha$ for all i . These same assumptions hold for all simulations to be presented below. We further assume throughout that $\tau = .75$.³

Figure 1.1 shows the results of repeated updates of the opinions of twenty-five agents, where every agent updates by (1.1), with $\alpha = 1$ and with varying values of ε . (The x -axis represents time, the y -axis opinion.) Note that to assume that $\alpha = 1$ is




Figure 1.1: $\alpha = 1$; $\varepsilon = .05$ (left), $\varepsilon = .15$ (middle), and $\varepsilon = .25$ (right)

really to assume that the agents do not receive any data about τ but update their opinions strictly by talking to their neighbors. How strongly data-gathering can affect the evolution of the agents’ opinions becomes apparent when we compare the graphs of Figure 1.2 with those of Figure 1.1. Everything that holds for Figure 1.1 holds for




Figure 1.2: $\alpha = .9$; $\varepsilon = .05$ (left), $\varepsilon = .15$ (middle), and $\varepsilon = .25$ (right)

Figure 1.2 as well, the only exception being that the agents in the updates represented in the latter, while they still give much weight to the opinions of their neighbors— $\alpha = .9$ —do not exclusively go by those opinions. The effect of allowing the data to influence their opinions is quite dramatic, as one can see. Hegselmann and Krause’s further results show that even slight variations in some of the other parameters can have similarly dramatic effects on how the agents’ opinions evolve over time.

It will not have been missed that the questions stated above are all of a strictly descriptive nature. And indeed, Hegselmann and Krause do not address any normative

³We did run the simulations to be presented in Section 2 with different values of τ . However, this made so little difference that we refrained from varying τ for the other simulations. Nevertheless, a systematic investigation of the extensions to be proposed would require redoing the simulations for other values of τ as well.

questions in their papers. This is not by oversight: they explicitly opt for a non-normative approach (see their [2006:2]). Nor does it make their results irrelevant to social epistemology. First, descriptive questions of the kind they ask fall well within the scope of social epistemology, at least if this is broadly understood as “the study of the social dimensions of knowledge or information” (Goldman [2006]). Second, not only can their model also help to address normative socio-epistemic questions, it is so general and flexible that it can easily be extended to models which are even more helpful in this regard. That, at least, is what we hope to show in the following sections. But before embarking on this, let us be a bit clearer about the normative aspects of social epistemology, and why we think they motivate a search for extensions of the HK model.

We take the key normative issue in social epistemology to be whether we ought to let our opinions be guided, even if only partly, by those of others. One should not expect a uniform answer to this question, but rather one of the form: “Under such and such circumstances, it is good to take into account the opinions of certain others,” or even of the more specific form: “Under such and such circumstances, it is good to take into account, to such and such an extent (or in such and such a way), the opinions of certain others.” For instance, it might be (and has been) defended that, when a person appears to be trustworthy and to be epistemically better placed with respect to a given issue than we are, we should adopt her testimony as regards that issue as our opinion on it.^{4,5}

Evaluation of such candidate answers presupposes some conception of goodness as pertaining to epistemic actions, such as the taking into account of one’s neighbors’ opinions. While one can think of several *prima facie* plausible explications of goodness in this regard, the most natural one seems to be in terms of truth. Truth, after all, is generally regarded as the overarching epistemic goal, both in traditional (individualistic) epistemology and in social epistemology. It might thus be said to be good to take into account other people’s opinions if (and perhaps only if) by doing so one will have a greater chance of hitting upon the truth, or one will get there more quickly than one otherwise would, or one will get closer to the truth, or some combination of these (for instance, one will have a greater chance of converging to the truth quickly). In the following, we shall assume this “truth-tracking” perspective in evaluating normative socio-epistemic claims.

As will be seen shortly, we can evaluate in this way answers of the forms mentioned above already in the HK model as is. Nevertheless, there are at least two good reasons to look for extensions of that model. The first is that the model is highly idealized, or at any rate that it could be said to model more or less realistically only a very limited class of actual epistemic situations. It should not be hard to grasp why extensions which could be said to be more realistic in several respects are desirable. Suppose that we decide to test a normative principle in a computer simulation model, and find that, in the simulations, the principle does well in the sense that those agents who obey the principle do better on average—for instance, hit upon the truth more often

⁴See, for instance, Fricker [2006] and Douven and Cuypers [2009].

⁵An anonymous referee made the interesting suggestion that one might distinguish the type of social epistemology supposed here, in which, roughly, one asks how an individual ought to adapt her opinions in view of information about the opinions of others, from a type of social epistemology that starts from empirically plausible assumptions about individual doxastic or epistemic behavior and then asks what, epistemically speaking, is the best arrangement of the group of individuals. The same referee suggested, rightly we now think, that some of the extensions to be presented below fit in naturally, and perhaps even more naturally, in this “social engineering” type of social epistemology.

or end up having beliefs which are closer to the truth or some such—than those who do not obey the principle. It seems that this will be the more significant the more realistic the model is, that is, the more the simulated agents are, in relevant respects, like real agents and the more the other features of the model are like the relevant properties of the world we inhabit. Equally, a negative result of a normative principle in simulations would seem to be more telling against that principle the more realistic the model or models used for the simulations are.⁶

The second reason has to do with the fact that there is always the risk that a simulation model creates artifacts. Hegselmann [1996, Sect. 3] nicely illustrates this by showing how, for a certain class of cellular automata models, an ostensibly insignificant change in their size can have a large and systematic impact on the outcome of the simulations performed in the models. Replication of some phenomenon in a model working on the basis of (partly) different principles can give at least *some* assurance that we are not dealing with an artifact created by some ill-understood features of the HK model. This is of particular importance when we are using simulations to assess normative claims. The last thing one would want to do is make an epistemic or doxastic recommendation on the basis of what is really such an artifact.

2. Noisy data. When considering a simulation in which all agents have equal access to the data (or whatever secures the truth attraction), a normative question that all but jumps to mind is whether these agents should want to talk to others at all. Why not go purely by the data? This is not meant as a rhetorical question. Indeed, it is not immediately clear whether in this situation talking to one’s neighbors might not still have some positive epistemic effects. However, a look at Figure 2.1 makes one wonder which effects these could be. The left graph of the figure represents the development




Figure 2.1: Talking (left) vs. no talking (right)

of the opinions of twenty-five agents who, next to taking into account the data, talk to their neighbors, where $\varepsilon = .1$; the right graph represents the development of the opinions of agents who, starting from the same initial opinions as the agents in the first

⁶Of course, a search for (more) realistic extensions is not only worthwhile from the viewpoint of normative social epistemology but also from Hegselmann and Krause’s descriptive viewpoint: we may expect simulations to be predictively more accurate the more realistic they are. But we will leave the matter of descriptive accuracy aside here. (It will be understood that by “realistic extensions” we mean “extensions that are realistic in relevant respects.” One can think of ways of making simulations more realistic—for instance, one might endow the agents with a digestive system—which would create additional complexity for no apparent benefit, given that they would be epistemically wholly irrelevant. Thanks to an anonymous referee here.)

group, do not talk to any of the others, that is, they update by (1.1) with $\varepsilon = 0$ ($\alpha = .5$ in both cases). The difference is minimal. A more systematic comparison of the two types of societies did not reveal any significant differences either. Simulations with 100 societies, each consisting of twenty-five agents, all with different initial opinions, and all of whom update by (1.1), confirm that the average distance from the truth of the agents' opinions after five, ten, twenty-five, and fifty updates differs at most negligibly when the agents take into account the opinions of others from when they do not do so (further updates were not considered, given that after ten updates the average distance from the truth was already approximately 0).⁷ This result turned out to be fairly robust for different combinations of values of α and ε .

This might make one think that, from a normative viewpoint, the HK model has little of interest to offer; what in the previous paragraph we took to be the key normative question for social epistemology—whether we should let our opinions be influenced by those of (certain) others—would seem to receive a disappointingly dull answer on the present approach: it doesn't matter, do as you like! But this conclusion would be rash. As already intimated, it does not take much to tweak the model in ways which make it directly relevant to normative questions. One way to do this, which is already noted in passing in Hegselmann and Krause [2006:3], is by assuming that the truth attraction does not work quite so smoothly as is supposed in (1.1), and that instead of data pointing them precisely to the truth, the agents receive data that are somewhat noisy, pointing them to a value possibly only close to the truth. It should be clear that adding this assumption to the HK model only goes to make it more realistic, inasmuch as in everyday scientific practice noisy data are the rule rather than the exception.

To incorporate this assumption into the model, we slightly alter (1.1), as follows:

$$x_i(t+1) = \alpha \frac{1}{|X_i(t)|} \sum_{j \in X_i(t)} x_j(t) + (1-\alpha)(\tau + f(i,t)). \quad (2.1)$$

The function f returns a uniformly distributed random real number in the interval $[-\zeta, \zeta]$, for some $\zeta \in \mathbb{R}$, where the output of this function may be different for each agent and vary per update. That is to say, the noise is independent and identically distributed for each agent on each update.⁸

That talking may now make a big difference is already suggested by comparing the graphs in Figure 2.2. The graphs show the developments of the opinions of twenty-five agents who receive data that may be "off" by $\zeta = .2$. The difference is that the agents whose opinions are represented in the left graph do take into account the opinions of their neighbors, and those whose opinions are represented in the right graph do not. There is an unmistakable difference between the graphs: while in the right graph all agents seem to have an opinion more or less close to the truth from a very early stage on, they do not seem, on average, to come much closer to the truth after further updates. By contrast, in the left graph it seems to take longer before all agents are more or less close to the truth, but eventually they are, on average, much closer to the truth than the agents in the right graph ever get.

This is further confirmed by comparing the graphs in Figure 2.3. The dotted line represents the averages over one hundred simulations of the average distance from

⁷For all i and t , the distance of agent x_i 's opinion from the truth at time t is simply taken to equal $|\tau - x_i(t)|$.

⁸As a referee noted, this assumption may be quite crucial in that the results of the simulations to be reported below might have been very different had the errors in the data the agents in those simulations receive been highly correlated. See on the importance of the assumption also Douven [2009, Sect. 2].




Figure 2.2: Noisy data, talking (left; $\alpha = .9, \epsilon = .1$) vs. no talking (right; $\alpha = .5, \epsilon = 0$)

the truth for societies of agents who do not talk: $\alpha = .5, \epsilon = 0$; the dashed line does the same for one hundred simulations with societies of agents who do talk, and who in fact give a lot of weight to the opinions of their neighbors: $\alpha = .75, \epsilon = .1$; and the continuous line does this for societies of agents who talk even more (and more trustingly): $\alpha = .9, \epsilon = .2$. Clearly, if the data the agents receive are noisy, then talking




Figure 2.3: Talking (dashed and continuous line) vs. no talking (dotted line)

to neighbors makes a large difference. Under these circumstances, talking helps to get the agents on average closer to the truth, though convergence to an average value that is moderately close to truth takes longer.

These findings suggest that there is no uniform answer to the question whether or not it is good to talk to neighbors when everybody receives noisy data. After all, in some situations it will be important to get relatively close to the truth quickly but much less important to get very close to the truth, in others it will be important to get very close to the truth, even if that should take a while.⁹

3. Elite and non-elite scientists.

Some scientists are more skilled at performing experiments than others. Consider a society of twenty-five agents, five of which are elite

⁹See Douven [2009] for more on this, and for an application of basically the same results to the currently much debated question of how one ought to respond to the discovery that one disagrees with an epistemic peer on some topic.

scientists who are so skilled at performing experiments that they always receive non-noisy data, the other twenty being non-elite scientists who receive noisy data which may again be off by .2. Thus, the updates follow again Equation (2.1), albeit now for five i it holds that $f(i, t) = 0$ for all t .

The gray lines in the left graph of Figure 3.1 show how the opinions of the non-elite scientists evolve over time when everyone in the society talks to their neighbors; the black lines show how, in that case, the opinions of the elite scientists evolve. The gray




Figure 3.1: Elite and non-elite scientists, talking (left; $\alpha = .9$, $\varepsilon = .1$) vs. no talking (right; $\alpha = .5$, $\varepsilon = 0$)

lines in the right graph show how the opinions of the same non-elite scientists evolve when no one in this society talks to anyone else; the black lines again show how the opinions of the elite scientists evolve.

These graphs already suggest that in a non-talking society the elite scientists converge rapidly to the precise value of τ and that in a talking society they do worse both in respect of speed of convergence and in respect of accuracy. They also suggest that, at least as regards long-run proximity to the truth, the non-elite scientists are worse off in a non-talking society than in a talking society. Runs with 100 different societies show that these results are in no way atypical. In particular, they confirm that on average talking brings the non-elite scientists as well as the society as a whole (i.e., elite and non-elite scientists together) eventually much closer to the truth than not talking. They further confirm, unsurprisingly, that by talking the elite scientists considered alone do worse than by not talking, in both of the above-mentioned respects.

Here, it might seem obvious that talking is good for the non-elite scientists and bad for the elite ones; that, to the extent you believe that you are an elite scientist, you should not want to take into consideration the opinions of any of your colleagues and, to the extent you believe that you are a non-elite scientist, you *should* want to talk at least to those colleagues whose opinions are close to your own. But the situation need not be this simple. *All* scientists, elite and non-elite alike, might benefit from the greater consensus in the community that is brought about by talking. For instance, suppose the community as a whole is involved in the search for a drug against a viral disease that threatens to decimate the world's population, and having a *relatively* accurate belief concerning the true value of τ greatly enhances, or at least is believed to greatly enhance, the chances of hitting upon such a drug, but having an even more accurate belief concerning that value does then not significantly increase those chances. Under these circumstances, even the elite scientists may prefer their society to be a talking society rather than a non-talking society. After all, they may well find it more

important that *someone* find the drug than that they themselves find it; in these circumstances, personal ambition may be less of a driving force than it normally is in science (if only because Nobel prizes are never awarded posthumously).

Of course, this is still very much idealized in that we are pretending that a society is either a talking or a non-talking society whereas in reality many intermediate cases are possible. For instance, scientists who believe themselves to be among the elite might want to consider making their beliefs public without themselves taking into consideration any of their colleagues' opinions. Or, more realistically still, people may be assumed to have reputations, where these reputations determine how much weight their opinions are given in discussions. We consider this possibility next.

4. Reputation. For better or worse, some scientists enjoy a higher status in their field of expertise than do others. No doubt one's reputation among colleagues has an effect on the weight these colleagues attach to one's opinions, at least on professional matters. An obvious question is whether taking into account scientists' reputations makes a difference to the results obtained so far. Another question, somewhat less obvious perhaps, is this: suppose the system of scientists' assigning reputations to one another works badly in that a scientist's reputation typically fails to reflect her scientific qualities (as cynics believe is true of certain fields of current science, or even of all of it); how much of a threat would this constitute to progress in science?

To address these questions, we again slightly tweak the HK model. This time the opinion of agent x_i at $t + 1$ is given by

$$x_i(t + 1) = \alpha \frac{1}{g(i, t)} \sum_{j \in X_i(t)} w_j x_j(t) + (1 - \alpha)(\tau + f(i, t)). \quad (4.1)$$

Here, $w_j \in \mathbb{R}^+$ stands for the reputation that is assigned to agent x_j , and the function g sums the reputations of the agents in $X_i(t)$, that is, $g(i, t) = \sum_{j \in X_i(t)} w_j$. Note that we assume that the agents have a fixed reputation, from the first till the last update, and that they have the same reputation in the eyes of each of their colleagues. One could of course consider relativizing reputations to updates and/or to agents,¹⁰ but we will not do so here.

One might think that giving most respect to the best scientists is highly important already for reasons of fairness, but that it certainly is when the respect we pay someone determines the weight we give to that person's opinion. Surprisingly, however, when everyone forms his or her opinion by means of (4.1), then, at least in the kind of situation we have been considering (everyone is trying to determine the value of a parameter in the interval $[0, 1]$), it turns out to be of little or no importance, epistemically speaking, whether the scientists give most respect to their best colleagues, or whether everybody is equally well respected, or whether the not-so-good scientists are held in higher regard than the good ones. The array of graphs in Figure 4.1 already gives some indication of this. The left graph represents the development of the opinions of twenty-five scientists, five of who receive non-noisy data, the others receiving noisy data ($\zeta = .2$); all scientists assign a weight of 2 to the opinions of the colleagues who receive the non-noisy data and a weight of 1 to the opinions of the rest. All this holds for the other two graphs as well, with the exception that for the middle graph

¹⁰As is done in Lehrer and Wagner [1981] for instance; see in this vein also French [1956] and DeGroot [1974].




Figure 4.1: The role of reputation ($\alpha = .9$, $\varepsilon = .1$)

the opinions of *all* scientists are given weight 1, and for the right graph the opinions of the scientists receiving the *noisy* data are given weight 2 and the opinions of the rest weight 1. In other words, in the left graph we see what happens when the good scientists are most respected, in the middle graph we see what happens when the good and the not-so-good scientists are equally respected, and in the right graph we see what happens when the not-so-good scientists are most respected. The differences are hardly discernible, and may simply be due to random fluctuations of the noise in each of the runs. More importantly, systematic simulations with weight 2 assigned to the best scientists and weight 1 to the others did not reveal any differences either, nor did systematic simulations in which we assigned higher weights still to the best scientists while keeping the weight assigned to the rest constant at 1, nor, in fact, did it make a difference when we assigned instead still higher weights to the less good scientists while keeping the weight assigned to the best scientists constant at 1.

We merely register this apparent insignificance within our model of the allocation of reputations to the members of the societies as a (for us, at least) unexpected result, leaving a deeper investigation of it to another occasion. Needless to say, reputation may play a role in scientific practice in various ways that go unaccounted for in the present model. For instance, a scientist's reputation may determine how much grant money she is able to attract and thereby how much time and effort she can put into the pursuit of her research projects. This and similar factors may greatly affect the development of her scientific discipline. Thus, it would be rash to conclude from the results presented in this section that, from a purely epistemic viewpoint, it is immaterial whether the system of allocating reputations to scientists functions properly in that the better a scientist is, the more reputed she is.

5. Coin flipping. In the HK model, and also in the extensions of it considered so far, the agents update their opinions by simultaneously taking into account the opinions of their neighbors as well as the information securing truth attraction. This is not particularly realistic. In reality, a scientist might perform an experiment one day and update her opinion on the information she receives from that, and might go to a conference on the next day and update her opinion by taking into account the opinions of certain colleagues who she meets there. Thus, an obvious further extension of the HK model, one which allows us to simulate societies whose members are free at any given time to update their opinions either on the basis of experimental information or by averaging the opinions of their neighbors, is given by the following equation:

$$x_i(t+1) = \begin{cases} \alpha x_i(t) + (1-\alpha)(\tau + f(i,t)) & \text{if } h(i,t) \text{ is heads,} \\ \frac{1}{|X_i(t)|} \sum_{j \in X_i(t)} x_j(t) & \text{if } h(i,t) \text{ is tails.} \end{cases} \quad (5.1)$$

Here, h is a function that returns the outcome of the flip of a coin and that, as is suggested by the notation, depends on both i and t . Thus, at any given time the agents update their opinions in a way that is independent of what the other agents do at the same time, indeed, independent of what any other agent does at any time. The coin may have any possible bias, ranging from consistently yielding heads to consistently yielding tails. Although it would be interesting to experiment with taking differently biased coins for different agents or groups of agents—as scientists with a lot of confidence in their experimental skills may be assumed to have a greater tendency to do experiments, as opposed to discussing matters with their colleagues, than those who are less confident in this respect—in the simulations we performed the coin always had the same bias for all agents; it also had the same bias at each time step.

Figure 5.1 shows that the people who experiment more than talk—that is, people who base their decision of whether to experiment or talk on the outcome of the flip of a coin with a .25 bias for talking—on average converge fastest to the truth; the people who use a fair coin for the same purpose on average converge more slowly but eventually get closer to the truth; and the people who talk most often—they use a coin with bias .75 for talking—on average converge even more slowly but end up being closer to the truth than the other groups.




Figure 5.1: Bias for talking: dotted line .25; dashed line .5; continuous line .75 ($\alpha = .9$, $\epsilon = .1$)

6. Concluding remarks. We have sought to provide extensions of the HK model that are meaningful from the perspective of the social epistemologist, in that they might help elucidate or even answer normative questions central to social epistemology. Needless to say, the extensions we presented are no more than a sampling of extensions of the HK model that can be employed for purposes germane to social epistemology. One obvious, but still interesting, avenue for further research would be to try to combine several of the above extensions. More radical departures from the original HK model would seem to be necessary to overcome two limitations of that model which have not been addressed here, to wit, the fact that the agents are assumed to know, at all times, the opinions of all the other agents in the model, and the fact that the agents have only one belief, rather than, as real people have, a set of (typically logically interconnected) beliefs. In one sequel to this paper (Riegler and Douven [2009a]), we extend the HK model to a model in which the epistemic agents

are able to move in a discrete two-dimensional space, and are privy at any given time only to the opinions of those of their colleagues who they then happen to meet (if any). In another sequel (Riegler and Douven [2009b]), we extend the aforementioned model in turn by populating it with agents that have much richer belief states than the agents in the HK model.¹¹

References

- Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G. [2000] "Mixing Beliefs among Interacting Agents," *Advances in Complex Systems* 3:87–98.
- DeGroot, M. [1974] "Reaching a Consensus," *Journal of the American Statistical Society* 69:118–121.
- Dittmer, J. C. [2001] "Consensus Formation under Bounded Confidence," *Nonlinear Analysis* 7:4615–4621.
- Douven, I. [2009] "Simulating Peer Disagreements," manuscript.
- Douven, I. and Cuypers, S. [2009] "Fricker on Testimonial Justification," *Studies in History and Philosophy of Science*, in press.
- Epstein, J. and Axtell, R. [1996] *Growing Artificial Societies*, Cambridge MA: MIT Press.
- Fortunato, S. [2004] "The Krause–Hegselmann Consensus Model with Discrete Opinions," *International Journal of Modern Physics C* 15:1021–1029.
- French, J. R. [1956] "A Formal Theory of Social Power," *Psychological Review* 63:181–194.
- Fricker, E. [2006] "Testimony and Epistemic Autonomy," in J. Lackey and E. Sosa (eds.), *The Epistemology of Testimony*, Oxford: Oxford University Press, pp. 225–250.
- Gaylord, R. J. and D’Andria, L. J. [1998] *Simulating Society*, New York: Springer.
- Goldman, A. [2006] "Social Epistemology," in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), available at <http://plato.stanford.edu/entries/epistemology-social/>.
- 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*, Dordrecht: Kluwer, pp. 209–233.
- 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 <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 <http://jasss.soc.surrey.ac.uk/9/3/10.html>.
- Lehrer, K. and Wagner, C. [1981] *Rational Consensus in Science and Society*, Dordrecht: Reidel.

¹¹We are greatly indebted to Christopher von Bülow and two anonymous referees for very helpful comments on a previous version of this paper.

- Lorenz, J. [2007] "Continuous Opinion Dynamics under Bounded Confidence: A Survey," *International Journal of Modern Physics C*18:1819–1838.
- Ramirez-Cano, D. and Pitt, J. [2006] "Follow the Leader: Profiling Agents in an Opinion Formation Model of Dynamic Confidence and Individual Mind-Sets," *Proceedings of the 2006 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pp. 660–667.
- Riegler, A. and Douven, I. [2009a] "Extending the Hegselmann–Krause Model II," in K. Placek (ed.), *Proceedings of ECAP 6*, London: College Publications, in press.
- Riegler, A. and Douven, I. [2009b] "Extending the Hegselmann–Krause Model III: From Single Beliefs to Theories," *Episteme*, in press.
- Weisbuch, G., Deffuant, G., Amblard, F., and Nadal, J. P. [2002] "Meet, Discuss and Segregate!" *Complexity* 7:55–63.