



HAL
open science

Network effects in a bounded confidence model

Igor Douven, Rainer Hegselmann

► **To cite this version:**

Igor Douven, Rainer Hegselmann. Network effects in a bounded confidence model. *Studies in History and Philosophy of Science Part A*, 2022, 94, pp.56-71. 10.1016/j.shpsa.2022.05.002 . hal-03921576

HAL Id: hal-03921576

<https://cnrs.hal.science/hal-03921576>

Submitted on 4 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Network Effects in a Bounded Confidence Model*

Igor Douven

IHPST / CNRS / Panthéon–Sorbonne University
igor.douven@univ-paris1.fr

Rainer Hegselmann

Frankfurt School of Finance and Management
r.hegselmann@fs.de

Abstract

The bounded confidence model has become a popular tool for studying communities of epistemically interacting agents. The model makes the idealizing assumption that all agents always have access to all other agents' belief states. We draw on resources from network epistemology to do away with this assumption. In the model to be proposed, we impose an explicit communication network on a community, due to which each agent has access to the beliefs of only a selection of other agents. A much-discussed result from network epistemology shows that densely connected communication networks are not always preferable to sparser networks. The aim of this paper is to investigate whether there are any noteworthy network effects in a version of the bounded confidence model augmented with communication networks, and in particular whether the aforementioned result from network epistemology can be replicated in that version.

Keywords: agent-based modeling; bandit problems; belief change; bounded confidence; network epistemology; robustness.

1 Introduction

It is a core insight of social epistemology that we owe much of our success *qua* epistemic agents to our interactions with others (Goldmann, 1999). While most social epistemologists have focused on interactions between two agents (as in much of the work on testimony; see, e.g., Lackey, 2015, and Fricker, 2019), or among no more than a handful of agents (as in some work on peer disagreement; see, e.g., Christensen, 2007), there is a growing body of work that looks at larger communities of agents. In particular, building on Bala and Goyal's seminal (1998), Zollman (2007, 2010, 2011) has asked whether certain types of communication structures to be found in such communities, or that could be imposed on them, increase or rather decrease the chances of the community to achieve its epistemic goals.¹

*The simulations reported in this paper were written in Julia (Bezanson et al., 2017). The Supplementary Materials for this paper consist of an HTML file that contains the code for the simulations and that can be run either locally or in the cloud (see the file for detailed instructions). The file can be downloaded from this repository: https://osf.io/w42fh/?view_only=469d3a0afc7b4bd9a9f32a0539765d2b.

¹See also Kummerfeld and Zollman (2016) and O'Connor and Weatherall (2017, 2019).

“Communication structure” is a broad term, which is here taken to comprise not just the social connections among agents—which agents talk to each other at all—but also mechanisms for taking others’ viewpoints into account, for aggregating them, as well as for combining them with evidence we receive directly from the world, through our senses. Especially in the philosophy of science, most of the relevant work has gone into studying communities of broadly Bayesian agents forming different types of networks. Central questions asked by researchers working in this field have concerned correlations between network topologies and overall epistemic success.

There has been a parallel approach to studying epistemically interacting agents, to wit, the so-called Bounded Confidence (BC) model, developed by Hegselmann and Krause (2002, 2005, 2006, 2009, 2015, 2019; see also Krause, 2002, and Hegselmann, 2004, 2014, 2020). This has been much in the limelight in the social sciences as well as in computer science but has received little attention from philosophers of science so far. There are some marked differences between the BC model and the “network epistemology” that has been more dominant in the philosophy of science. What may appear the most salient one is that there has been little explicit discussion of networks or network topologies in the literature on the BC model (though see note 2 and Sect. 4). To some extent, however, that is a matter of presentation rather than of substance. The two approaches are more similar than they may at first appear to be, and it is in fact straightforward to incorporate elements from one into the other.

To see the connection, distinguish between two kinds of networks: *access networks*, which determine, for each agent, the agents to whose belief states the agent has access (i.e., is fully informed about), and *impact networks*, which determine, for each agent, the agents whose belief states impact the agent’s belief state. In the standard BC model, access networks are trivial in the sense that they are always complete, which is to say that, at any given time, every agent has access to all the other agents’ belief states. Impact networks, by contrast, are dynamic in that model, meaning that they can change from one time step to another. In the standard BC model, impact networks *strictly* depend on closeness relations among belief states: at any given time, an agent’s belief state is impacted by all and only agents whose belief states are (in some sense) close enough to its own.

In the approach to social learning taken by Zollman and others, access and impact networks coincide and are, moreover, static: whether an agent’s belief state is impacted by another agent’s belief state depends on whether they are connected in a network that is given at the start of the collective learning process and stays in place until the very end. Zollman (2007, 2010) reported a surprising finding about learning in such networks, to wit, that sparser networks will sometimes be more conducive to improving the accuracy of the agents’ belief states than more densely connected networks. As Rosenstock, Bruner, and O’Connor (2017) discovered, however, that finding needs qualification, in that it holds only in a small area of the parameter space for Zollman’s model.

For now, we note that there is nothing in the BC model per se that prevents it from taking on board the assumption underlying the work of Zollman and others that our belief states do tend to be impacted by those with whom we stand in certain social relations. Specifically, unless we limit the community of agents in the BC model to some small group of people who are constantly in touch with each other—like, perhaps, the members of a laboratory, or a group of friends—the assumption of total access implicit in that model is unrealistic. Thus, it would make sense to abandon the assumption tacitly being made in the BC model that the access network is trivial and replace it by assuming instead still fixed but non-trivial access networks in the manner of the aforementioned alternative approach.

It is perfectly legitimate to study social learning from a purely descriptive perspective, as much work on the BC model has done, but like Zollman and others, we are mostly interested in the normative questions this type of learning raises, in particular, in whether social learning comes with any epistemic benefits or costs. If we follow the mainstream and see our epistemic goal as believing the truth and nothing but the truth (e.g., Lehrer, 1974; Bonjour, 1985; Foley, 1993), then the question of possible epistemic benefits can be formalized using known measures of accuracy, indicating how close our belief states are to the truth.

In this paper, we present the new, combined model of social learning, and start investigating it by applying it to the kind of problems—so-called bandit problems—that Zollman and others have studied using the fixed network model. We are interested in possible effects of varying network topology on the epistemic performance of agents in the network and also in whether our model gives rise to a network effect comparable to the one Zollmann observed in his model.

Section 2 summarizes the version of the BC model that we build on in this paper. Section 3 describes the kind of learning problems that the model will be applied to. Section 3 makes the aforementioned distinction between access networks and impact networks further precise and illustrates their relevance to social learning by means of computational modeling. Sections 5 and 6 describe the outcomes of simulations conducted with two different instances of our model aimed at answering the question of whether any network effects are to be found in the model and if so which. The former section focuses on group-level effects, the latter on agent-level effects.

2 The Bounded Confidence model and beyond

2.1 The standard model

In the first published version of the BC model (Krause, 2000; Hegselmann & Krause, 2002), agents update their opinions by “averaging” (in a specified way) over the opinions of those that are (in a specified sense) near to them. In a later version of the model (Hegselmann & Krause, 2006, 2009, 2015, 2019), updates also take into account evidence that the agents receive directly from the world. The later model is explicitly meant to capture, in a formalized and idealized way, a fundamental truth about human learning, to wit, that it is a “mixture” of learning from others and learning from the world. As said earlier, the BC model has been widely used to study descriptive questions, including questions concerning when the opinions of communities of interacting agents converge on a common opinion and when they diverge (see apart from Hegselmann and Krause’s writings already referenced, e.g., Deffuant et al., 2000; Lorenz, 2003, 2008; Chen & Lou, 2019; Hahn, Hansen, & Olsson, 2020). Much work has also gone into studying the amount of time it takes for a community to reach a fixed point, in that further updates leave the opinions of the agents unchanged (Chazelle, 2011; Kurz & Rambau, 2011; Kurz, 2015; Hegarty & Wedin, 2016). But there is also some work on the BC model that has looked at normative questions, for instance, concerning the practice of assertion (Olsson, 2008), the resolution of disagreement amongst peers (Douven, 2010; De Langhe, 2013), and efficient truth approximation (Douven & Kelp, 2011). In this paper, we also take a normative perspective and present a new extension of the BC model developed specifically with an eye toward addressing normative questions that have been raised in the philosophy of science literature. We start by describing the standard BC model (in the later version first presented in Hegselmann & Krause, 2006), on which we then build in the following.

In the standard BC model, a community of agents is trying to determine the value of some




Figure 1: Repeated BC updating of communities of 50 agents, for $\tau = 0.7$ with different values for α and ϵ .

(further unspecified) parameter $\tau \in (0, 1)$. The agents start with a random guess and then update their estimate of τ repeatedly, at discrete points in time. The core idea is that each agent updates on the basis of *both* the information about τ that it receives from the world *and* the estimates of all agents within its so-called bounded confidence interval (BCI), which are precisely those agents whose estimate of τ differs by at most ϵ (for some specific $\epsilon \in [0, 1]$) from the agent’s estimate. To be more exact, in the BC model agent x_i ’s opinion concerning τ after the $(u + 1)$ -st update is defined to equal

$$x_i(u+1) = \frac{1 - \alpha}{|X_i(u)|} \sum_{j \in X_i(u)} x_j(u) + \alpha \tau, \quad (\text{BC})$$

where $x_j(u)$ is agent x_j ’s opinion after update u ,

$$X_i(u) := \{j : |x_i(u) - x_j(u)| \leq \epsilon\}$$

are the agents in agent x_i ’s BCI after the same update, and $\alpha \in [0, 1]$ determines how the agent weights the “evidential” and the “social” parts of the updating relative to each other.

To illustrate the model, Figure 1 shows, for different settings of α and ϵ , how the opinions of 50 agents, all starting with a random guess of the value of τ , change over the course of 25 updates. Note that, although in their updates they give much greater weight to the estimates of the agents within their BCI than to the information they receive from the world (even where $\alpha = 0.3$), their opinions quickly converge to the true value of τ (which is 0.7 in this case). While that is so for all of the three parameter settings that were used, we see that exactly *how fast* the opinions converge does depend on the specific setting.

An often-mentioned advantage of the BC model is that it can be easily tweaked to whatever one’s specific research needs are. For instance, researchers have presented extensions explicitly meant to deal with situations in which agents receive *noisy* evidence from the world (Douven, 2010; Douven & Riegler, 2010; De Langhe, 2013), or in which they hold *vague* rather than sharp beliefs (e.g., vague rather than sharp estimates of τ) about the world (Crosscombe & Lawry, 2016). The new extension to be presented in this paper is more similar to extensions studied by Lorenz (2003, 2008), Jacobmeier (2004), and Pluchino, Latora, and Rapisarda (2006), in which agents’ belief states are characterized by more than one estimate. Also as in these extensions, in our extension all agents hold beliefs on issues that are *logically and statistically independent* of each other. That is exactly what we need for the kind of problems that we aim to tackle, but it is worth noting that there are also extensions of the BC model that feature agents holding multiple beliefs on logically interconnected issues (Riegler & Douven, 2009; Wenmackers, Vanpoucke, & Douven, 2012, 2014).

2.2 Two-dimensional bounded confidence updating

In Zollman’s bandit problems, agents have numerical beliefs about two hypotheses (to be interpreted as concerning the payoff of certain actions, or the reliability of methodologies; see below). Just extending the BC model to one in which agents’ belief states are characterized by two numbers instead of one is straightforward. The truth is now to be thought of as a tuple, $\tau = \langle \tau_1, \tau_2 \rangle$, with agents still basically updating via (BC), applying the update rule separately to τ_1 and τ_2 . Indeed, the generalization to any finite number of numerical beliefs is straightforward. The main difference is in the definition of the BCI. Where $x_i(u) = \langle x_1^{i,u}, \dots, x_n^{i,u} \rangle$ is a sequence of agent x_i ’s estimates of $\tau_1, \tau_2, \dots, \tau_n$, after the u -th update, we now have that the agents in x_i ’s BCI are precisely those in the set

$$X_i(u) := \{j: d(x_i(u), x_j(u)) \leq \epsilon\},$$

with d a distance metric. There are many metrics one could choose here, but the most plausible candidates are the Manhattan and the Euclidean metric, which are the instances with $p = 1$ and $p = 2$, respectively, of the following schema:

$$d(x_i(u), x_j(u)) = \sqrt[p]{\left(\sum_{k=1}^n |x_k^{i,u} - x_k^{j,u}|^p\right)}.$$

Figure 2 illustrates the extended BC model for communities whose agents hold numerical beliefs about two independent matters, for different parameter settings, and for both of the aforementioned metrics. The figure shows the trajectories of the belief states of 50 agents in the unit square, so with $t_1 \in (0, 1)$ as well as $t_2 \in (0, 1)$. While varying the parameters has a clearly visible effect, the same is not true for varying the metric. Naturally, that does not necessarily mean that there might not be an effect measurable in some other way, but in none of our computational experiments did the choice of metric make a qualitative difference. In the following, we only report the results based on use of the Euclidean metric.

2.3 Opening the black box

The BC model treats the updating on worldly evidence as a kind of black box: the value of τ impacts the agent’s opinion in a way that is left unspecified. This was a design decision of the developers of the model and was motivated by the fact that their main goal was to study the effects of social learning, however exactly agents updated on worldly evidence. Nevertheless, in the meantime researchers also became interested in the normative status of individual update rules in light of the fact that we are not *strictly* individual learners but also learn from others. Looking for a formal model that could facilitate studying individual update rules in a social setting, they turned to the BC model, although now equipping the agents populating the model with whichever individual update rule or rules the researchers were interested in.

In the context of comparing Bayes’ rule with probabilistic versions of the so-called Inference to the Best Explanation, Douven and Wenmackers (2017) and Douven (2019, 2022a) defined an extension of the BC model in which agents’ belief states were characterized as probability assignments to a finite set of self-consistent, mutually exclusive, and jointly exhaustive hypotheses, and in which each agent updated its probabilities on the worldly evidence it received either via Bayes’ rule or via a rule belonging to the aforementioned family of alternatives. These authors were then particularly interested in investigating which rule or rules offered the best trade-off between speed of convergence (how quickly the agents’ belief states approached the

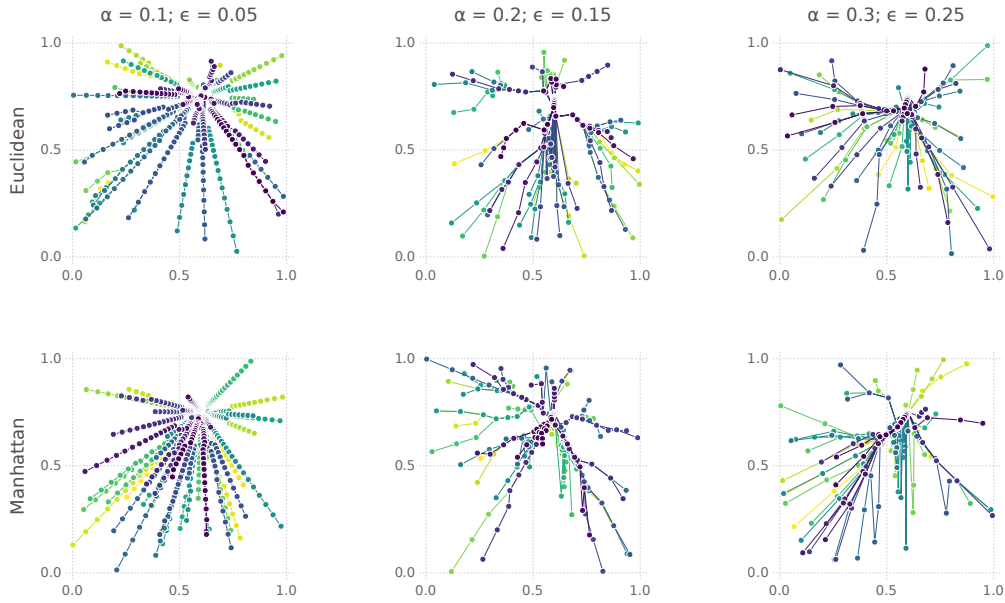


Figure 2: Repeated two-dimensional BC updating (20 updates) of communities of 50 agents, for $\tau = \langle 0.6, 0.75 \rangle$ with different values for α and ϵ per column, assuming the Euclidean metric (top row) and the Manhattan metric (bottom row).

truth) and accuracy (how closely their belief states approached the truth) and especially in what impact the social part of the updating process had on that trade-off, where the social part was as in the original BC model except that averaging of numerical beliefs was replaced by pooling of probabilities in the manner of Stone (1961).

In a variant of this extension, Douven (2022b) proposed a BC model populated with “Carnapian” agents, that is, agents equipped with update rules all belonging to the family of λ rules as first presented in Carnap (1952). To explain these rules, let $K = \{Q_1, \dots, Q_k\}$ be a set of elementary outcomes and $e_n \in K^n$ a sequence of n outcomes, such that, for all i , Q_i occurs with fixed probability p_i , and $\sum p_i = 1$. Let n_{Q_i} be the number of outcomes in e_n in which Q_i occurred. Then an agent is said to update its probabilities by means of a λ rule precisely if, for all j and n , and for some $\lambda \in \mathbb{R}_0^+ \cup \{\infty\}$,

$$\Pr(Q_j | e_n) = \frac{n_{Q_j} + \lambda/k}{n + \lambda}, \quad (\text{C})$$

with \Pr the agent’s probability function.

The λ parameter is to be interpreted as regulating the agent’s learning rate. If λ is set equal to 0, probabilities are always equal to observed relative frequencies, no matter the agent’s prior probabilities. The larger λ is set to be, the more the updater will stick to her prior probabilities, the limiting case being $\lambda = \infty$, which precludes any learning from data. Carnap saw λ as measuring how cautious the updater is and opined that there is no objectively correct value of λ . At the same time, Carnap (1952) does express a preference for $\lambda = 2$. In the kind of statistical model to be used in the following, the λ rule with that setting in effect amounts to Laplace’s rule of succession.

We are going to compare our model to the network model studied in Zollman (2007, 2010, 2011). Given that the agents in Zollman’s model are Bayesian updaters, it is worth clarifying how

they relate to agents updating strictly via the Carnapian schema (C). Zollman uses the same statistical model as is to be used in this paper. In this model, all agents receive binomial data (i.e., independent and identically distributed binary data), which come in one data point at a time. The typical instance of such a model is that in which the data consist of the outcomes of tosses with a coin having an unknown bias $b \in [0, 1]$ and we update our probability distribution over possible values for b again and again, as more and more data accrue. (The actual model used by Zollman and in the following is equivalent to one involving *two* coins, but that is immaterial for now.)

First note that, in this model, the parameter k in (C) equals 2, given that there are but two elementary outcomes; heads and tails, let us say. So, letting H designate the coin landing heads, for our model (C) can be rewritten as

$$\Pr(H | e_n) = \frac{n_H + \lambda/2}{n + \lambda}. \quad (\text{C}^*)$$

Furthermore, as Zollman notes, given that b (the bias of the coin, which is the objective probability with which H will occur) is a continuous parameter and the data follow a binomial distribution, the updating can be efficiently modeled via so-called beta distributions. The shape of a beta distribution is determined by two parameters, α and β . A nice property of the beta distribution is that it is a conjugate prior for the binomial distribution, which means that updating a beta distribution on data following a binomial distribution results again in a beta distribution. To be exact, suppose we start with a beta distribution with parameters α and β and observe n_H heads in n tosses. Then the updated beta distribution has parameters $n_H + \alpha$ and $n - n_H + \beta$. The final thing to notice now is that the mean of a beta distribution equals the first parameter over the sum of the two parameters, which for the present case yields

$$\frac{n_H + \alpha}{n_H + \alpha + n - n_H + \beta} = \frac{n_H + \alpha}{n + \alpha + \beta}.$$

Comparing this with (C*), we see immediately that the probability that a Carnapian assigns to heads, given that she observed sequence e_n , equals the mean of the probability distribution of a Bayesian agent who started out with a beta distribution with $\alpha = \beta$ and observed the same sequence of tosses. Asking a Bayesian for her probability for heads on the next toss, given the data observed so far, is asking her for a point estimate, and the mean of a distribution—also known as “Least Mean Squares estimator”—is generally considered to be the appropriate point estimate in the case of a continuous parameter, such as b (Bertsekas & Tsitsiklis, 2008, Ch. 8). Hence, a Bayesian agent starting with a beta distribution with parameters α and β will always post the same probabilities for heads as a Carnapian agent with $\lambda = 2\alpha = 2\beta$, supposing they have received the same data. In particular, a Carnapian with $\lambda = 2$ will always post the same probabilities for heads as a Bayesian starting from the flat beta distribution—which is the beta distribution with $\alpha = \beta = 1$ —again supposing they have seen the same data.

Agents strictly updating via (C) will occur in our model as a limiting case. Our model will be populated by *social* Carnapians, that is, agents that also take into account the opinions of some or all of their community members. We here are building here on Douven’s (2022b) proposal to cash out what is a black box in the original version of the BC model—viz., the procedure of updating on worldly evidence—in terms of Carnap’s λ rules. We state here the most general version of the resulting model. Let e_u^i designate the sequence of outcomes that agent i has seen after the u -th update, and let H again designate heads (or whatever counts as success). Furthermore, let $\Pr^i(H | e_u^i)$ be agent i ’s estimate of b , and so the probability of obtaining heads

on a toss, after having seen data e_u^i , with u_H^i heads. Then for all agents i and updates u ,

$$\Pr^i(H | e_{u+1}^i) = \alpha_i \frac{1}{|X_i(u)|} \sum_{j \in X_i(u)} \Pr^j(H | e_u^j) + (1 - \alpha_i) \frac{u_H^i + \frac{\lambda}{2}}{u + \lambda}, \quad (\text{BCC})$$

with the set of agents within agent i 's BCI after the u -th update now being defined as follows:

$$X_i(u) := \{j: |\Pr^j(H | e_u^j) - \Pr^i(H | e_u^i)| \leq \epsilon_i\}.$$

This is the most *general* version because it allows agents to have their own α and ϵ values. We look at that version in Section 6, but until then we focus on the model in which $\alpha_i = \alpha_j$ and $\epsilon_i = \epsilon_j$ for all agents i and j .

In this paper, we are going to apply this model to agents estimating the probabilities of two independent parameters. For a first illustration, suppose each agent has two coins with possibly different biases for heads, where these biases are the same for all agents in the community. Thus, if *one* agent has a coin with bias 0.6 for heads and another coin with bias 0.75 for heads, then *all* agents have one coin with bias 0.6 for heads and another with bias 0.75. The worldly input each agent receives at any time step comes from flipping each of the coins at the time step and observing whether it lands heads or tails. The individual part of the updating consists of applying the Carnapian rule with $\lambda = 2$ to update the estimates of the two biases. To combine this with the social part of the updating requires a slight generalization of the definition of the BCI, exactly in the way we saw above (so we again consider both the Euclidean and the Manhattan distance metric), but otherwise does not call for any changes. Once an agent has determined which other agents are in its BCI, it calculates, for each coin, the mean of the estimates of the coin's bias held by those in the agent's BCI and then makes its own new estimate of the bias a convex combination of that average and of the estimate obtained via the individual part of the updating process, where α again determines how the outcomes of the individual and social parts are weighted against each other, and ϵ again determines how "liberal" the agent is in counting others as being in its BCI. Figure 3 shows the trajectories of the belief states of 50 agents updating via (BCC), for different parameter settings, but always with $\tau = \langle 0.6, 0.75 \rangle$, which can be thought of as indicating that the first coin has a bias for heads of 0.6 and the other coin has a bias for heads of 0.75.

3 Bandit problems

In the two-dimensional BC model with social Carnapian agents that was illustrated in Figure 3, each agent flipped both coins in its possession at each time step. But now suppose it were allowed to flip only one coin each time step, where the choice which coin to flip at a given time step were up to the agent. Then, given that the agent would still like to determine the bias of *both* coins, it could decide to flip the coins interchangeably, first one coin, then at the next time step the other, then the first again, and so on. We add a further twist: each time a coin lands heads, the agent receives one dollar, and nothing if the coin lands tails. To maximize its expected returns, the agent might now decide to always pick the coin it estimates to have the highest bias for heads. While that appears reasonable, the agent realizes that its estimates of the biases may be wrong: it might be maximizing *expected* returns by always picking one coin, even though it would actually make more money in the long haul were it to pick the other coin instead. The agent could increase the accuracy of its estimate of the bias of the coin it now




Figure 3: Repeated two-dimensional BC updating (20 updates) of communities of 50 social Carnapian agents, for $\tau = \langle 0.6, 0.75 \rangle$ with different values for α and ϵ per column, assuming the Euclidean metric (top row) and the Manhattan metric (bottom row).

believes to be less profitable by also from time to time flipping that coin. But how to balance the expected gain in accuracy and the expected loss of income?

In the typical presentation of this problem, it is not about coins but about two-armed bandits. Regardless, the problem is not, at bottom, about games of chance but manifests itself in many ordinary situations. Doctors might have to choose between two different therapies. It would seem unethical not to choose the one with the highest chances of curing the patient. But what if they are mistaken about which therapy is the more efficacious one? And supposing that, in reality, the “other” therapy is more efficacious, how will they ever find out if it never gets applied? Similarly, scientists might have to choose between two research methods. They consistently choose the one they believe to have the highest success rate. But if the alternative has *in fact* a higher success rate, how are they ever to find out? As a final example we mention the problem of choosing a website a company will use to persuade visitors to purchase its products. Will a website with a red buy button (“Click here to buy!”) lead to more revenue than the same website but now with the buy button appearing in yellow? The company might have more confidence in the red button, but to find out whether that confidence is justified, they would have to make sure that at least some visitors see the web page with the yellow button—which, however, in the meantime could lead to a loss of revenue.

Finding an optimal balance, in this kind of situation, between exploration (trying to improve the reliability of one’s estimates) and exploitation (maximizing profits) is an old problem with no easy solution (see, e.g., Berry & Fristedt, 1985). As Zollman (2010) notes, however, it may help to approach the problem from a social perspective. After all, in a social setting, instead of exploring the alternative they expect to lead to lower returns themselves, agents may be able to learn about it from others who actually put more confidence in that alternative and therefore choose it over the other option.

Zollman addresses bandit problems by placing Bayesian updaters in a network. Each agent always plays the arm it believes to have the highest success rate. However, it also has access to the results obtained by its neighbors in the network. Zollman (2007, 2010) is mostly interested in the impact of network connectivity on a certain practical outcome, viz., that all network members choose the better arm. It is assumed that the members of the network always play the arm they consider to be the better arm. Using computer simulations, he finds that, whereas greater connectivity tends to speed up convergence to a consensual view on what is considered to be the better arm, it tends to *decrease* the chance that the arm is indeed the better one. As Zollman explains, that is because, in a densely connected network, runs of outcomes deviating from the limiting frequency of success of an arm can more easily impact the network as a whole, quickly pushing all agents off track, while in sparser networks, there is a greater chance that at least some of the agents stick to the better belief, which they may then later disseminate to the rest of the community, also helping to correct agents who had been misled by unlucky subsequences of outcomes.

While interesting, Zollman’s results are limited in scope, as was discovered by Rosenstock, Bruner, and O’Connor (2017). Specifically, these authors found that the network effect documented in Zollman’s work only manifests itself in a small region of parameter space. In all of Zollman’s simulations, one arm has an associated success rate of 0.5 while the other one has an associated success rate of 0.501. Extending Zollman’s simulations, Rosenstock and colleagues show that the network effect becomes virtually negligible if the success rate of the second arm is just slightly increased, to 0.51, and that the effect vanishes completely as soon as its success rate is 0.53 or above.

Among other findings, this has raised concerns about the robustness of results from agent-based modeling studies; see, besides Rosenstock and colleagues’ paper, Frey and Šešelja (2018, 2020), Borg et al. (2019), and Šešelja (2022). In the case of Zollman’s studies, the problem is not only that the main finding fails to hold beyond a small fraction of the relevant parameter space; the problem is also that, where it holds, it is of the least interest. To be sure, if one therapy has a 0.5 probability of saving a patient’s life while a second has a 0.501 probability of saving that same patient’s life, then, all else being equal, we should prefer the second over the first any time. And for Amazon, it will make big difference to their revenue if, by changing the color of the buy button, they can increase the probability of a visitor of their website making a purchase from 0.5 to 0.501. But for many practical applications, an increase in success rate of that order of magnitude will make no real difference. We take these concerns about robustness seriously and address them using a method first proposed in Douven and Hegselmann (2021).

In this paper, we study bandit problems in the context of the extended BC model stated in the previous section. We populate the model with agents relying on Zollman’s decision rule (“play the best”) to always play the arm they believe to have the highest chance of success, and we also look at a variant rule, proposed by Kummerfeld and Zollman (2016), which lets agents play the arm they believe to be the most successful one with some fixed high probability p and play the other arm with probability $1 - p$. (Kummerfeld and Zollman call this rule “ ϵ -greedy”; because we already have a parameter ϵ , we call it “ δ -greedy.”)

For the most part, however, we will be interested in a generalization of Kummerfeld and Zollman’s proposal, to wit, that at each time step an agent chooses an arm with a probability *proportional* to the extent by which it believes the arm’s success rate to exceed that of the other arm. Formally, where a_1^t and a_2^t are a given agent’s estimates at time t of the success rates of the first and, respectively, second arm, the agent plays arm 1 at time t with probability

$$p_1^t = 0.5 + 0.5(a_1^t - a_2^t),$$




Figure 4: Visualization of proportional choice rule. (See the text for explanation.)

and it plays arm 2 at t with probability $1 - p_1^t$. One can think of this approach as a kind of equivalent to portfolio diversification. If an investor thinks one asset class will appreciate more than another, she may make the former a bigger part of her portfolio. But because price actions can never be predicted with certainty, she may hedge by also taking a smaller position in the other asset class. It is the same with the rule we are proposing: the agent will mostly play the arm (flip the coin, use the method, and so on) it believes to have a greater chance of leading to success, but it will also allocate *some* of its time to exploring the other option, and in fact it will allocate *more* time to that the closer it believes that other option's success rate to be to that of its competitor. Figure 4 visualizes this “proportional” choice rule, with color indicating the probability with which an agent will chose arm 1, as a function of its estimates of the success rates of the two arms.

Like Zollman, we are interested in the contribution network structure makes to how the agents in a community perform. The crucial difference with Zollman's approach that we want to focus on is that, where Zollman's networks are static in that they do not undergo any changes during the learning process, networks implicitly associated with the BC model are dynamic: they change, or at least can change, as the learning process unfolds. The idea is that, at each time step, the network is determined by the relation of an agent's being in another agent's BCI. That means the communication network can be constantly changing, given that different agents can be in an agent's BCI at different times. To help readers visualize the connection between the extended BC model and communication networks, the upper left panel of Figure 5 shows the trajectories of the belief states of 10 agents updating two independent estimates via (BCC), all agents having an α value of 0.3 and an ϵ value of 0.25, and using the proportional rule for selecting which arm to play; the other panels show, for the first seven time steps, the communication network structures corresponding to the BCI relations that exist at the given time step. (We are using a community of only 10 agents to avoid cluttering the graphs. Also note that, because all agents have the same ϵ value, BCI relations are symmetric and so we can leave off arrows from the edges. We further do not represent the reflexivity of BCI relations, given that an agent is always in its own BCI.) We see that, whereas at the beginning relatively few agents are connected, after six updates the agents form a fully connected communication network, given that every agent is in every other agent's BCI.

Note that nothing prevents us from combining a static communication network approach




Figure 5: Top left panel showing 20 updates of a community of 10 social Carnapian agents trying to estimate $\tau = \langle 0.6, 0.75 \rangle$ using (BCC) and the proportional choice rule. The other panels show the networks formed by the BCI relations at the first six time steps.

of the kind explored in Zollman’s work with the BC model, in which communication networks are dynamic. The basic idea would be that one agent’s belief state impacts another agent’s belief state at a given time t if, and only if, the first agent is connected to the second in the static network imposed on the community of agents *and* the first agent is within the second agent’s BCI at t . Such an amalgamation would make both Zollman’s model and our extension of the BC model more realistic. After all, we are not typically always informed about the belief states of all agents in our community—contrary to what the BC model and its extensions assume—nor however are we always influenced by what those whose belief states we *are* familiar with believe: even if our neighbor does not tire of sharing her views with us, we may never have thought much of those and have long decided to ignore them.

In the following, we will be interested in network effects, like Zollman was, although for us the central question is whether imposing static networks on a community of agents updating via (BCC) can benefit those agents, epistemically speaking, and if so, under which circumstances.²

4 Adding static networks

Figure 6 shows three standard types of network: a fully connected one, a wheel, and a cycle, all for ten agents. (We always assume that any agent is connected to itself, so strictly speaking there should, for every node, be an edge connecting that node to itself.) Say that an edge between two nodes in any of these networks means that the agents represented by the nodes have access to each other’s belief states at any time, and so that the agents *can* stand in a BCI relation to each other (whether or not they do). Then we can say that the BC model (any version of it) implicitly assumes a fully connected access network. But, as mentioned above, there is no impediment to imposing a more restricted access network on a community of agents. We also mentioned that, if we do so, we can say that one agent is in another agent’s BCI if the belief state of the former is “sufficiently close” (made precise as before) to that of the latter *and* there is an edge in the access network connecting the former to the latter. Where the sufficient closeness relation defines a network that can change from one update to another—as illustrated in Figure 5—the access network is fixed from the beginning of the learning process until the end.

Zollman (2010) found for his model that sparser networks—access networks in our sense—tended to improve the performance of the agents in the network in that agents in sparser networks had overall more accurate beliefs than agents in more densely connected networks. A first question to ask is whether there is an analogue of this effect for our extension of the BC model with different access networks assumed.

Figures 7 and 8 give an impression of what we might find. The figures show output from

²By focusing on this question, we are actually following a suggestion that was already made in Hegselmann and Krause (2002). Toward the end of that paper, the authors consider the possibility of running the BC dynamics on underlying static network structures. They mention in particular checkerboard-based networks, conjecturing that these would greatly affect the outcomes of the BC dynamics and specifically that the BC dynamics on such networks would result in less polarisation for comparatively small neighborhoods, such as von Neumann or Moore neighborhoods (Hegselmann & Krause, 2002, p. 29). For an overview of other relevant work combining BC dynamics with static networks, see Parasnis et al. (2018, 2021). See also Fortunato (2004), Riegler and Douven (2009), Ahat et al. (2011), and Schawe et al. (2021) for extensions of the BC model that place the dynamic impact network on certain forms of independent static access networks. For other work using static networks in the context of different approaches to opinion dynamics, see DeMarzo, Vayanos, and Zwiebel (2003), Golub and Jackson (2010), Grim et al. (2015), and Friedkin et al. (2016).




Figure 6: A complete network (left), a wheel (middle), and a cycle (right).

simulations involving communities of 50 agents updating via (BCC), where the worldly evidence always consists of data from two arms, one with a bias for success of 0.5 and the other having a bias for success of 0.501, which—not coincidentally—are the same success rates assumed in Zollman’s simulations. In all simulations represented in Figure 7, the agents have the same α value of 0.2 and the same ϵ value of 0.15; in the simulations represented in Figure 8, these values are 0.1 and 0.05, respectively. The panels show, at each time step, average Euclidean distances from the truth—which is $\langle 0.5, 0.501 \rangle$, as said—where the average is over all agents in the community; error bars represent one standard deviation from the mean.

Comparisons are along two dimensions: different access networks are imposed on the communities, and the members of the communities use different rules to choose which arm to play. To compare the impact of network structure, we must compare rows in both figures: the top rows give the results for agents placed on a completely connected network (the usual situation in all versions of the BC model), the middle rows give the results for agents placed on a wheel, and the bottom rows do the same for agents placed on a cycle. The columns compare the different choice rules, the left columns showing the results for agents which always play the arm they believe to be best, the middle columns showing the results for agents which are δ -greedy (with $\delta = 0.1$) in the sense of Kummerfeld and Zollmann (2016), and the right columns showing the results for the proportional choice rule defined above.

Just eye-balling these results, it is obvious that the rule agents use to choose which arm to play has the biggest effect on average distance from the truth. At least in these simulations, there is no need to run any statistics to appreciate that agents using the proportional rule had greatly diminished error rates in comparison to virtually all other agents. One might want to decide to continue only with the proportional rule, but of course results could be very different for different parameter settings.

The type of access network imposed on a community would, from these results, appear to have a more subtle effect, if any. Of course, these results concern single runs, and the outcomes might depend on the random estimates with which the agents start. To gain a better understanding of the relation between network structure and accuracy, we ran 50 simulations for each combination of choice rule and network structure, for both parameter settings that were used in the single runs. In each simulation, we again calculated the mean distance from the truth at each time step. Further averaging the scores over the 50 simulations yields the results shown in Figure 9 for the setting with $\alpha = 0.3$ and $\epsilon = 0.25$ and the results shown in Figure 10 for the setting with $\alpha = 0.1$ and $\epsilon = 0.05$.

What can already be said on the basis of the plots is that there is no indication that the kind of network effect found by Zollman, in which sparser networks increase the overall accuracy of a community of agents, replicates for the new extension of the BC model: for any of the choice rules, a fully connected access network leads to scores that are never higher and typically lower

$$\alpha = 0.3; \epsilon = 0.25; \tau = \langle 0.5, 0.501 \rangle$$

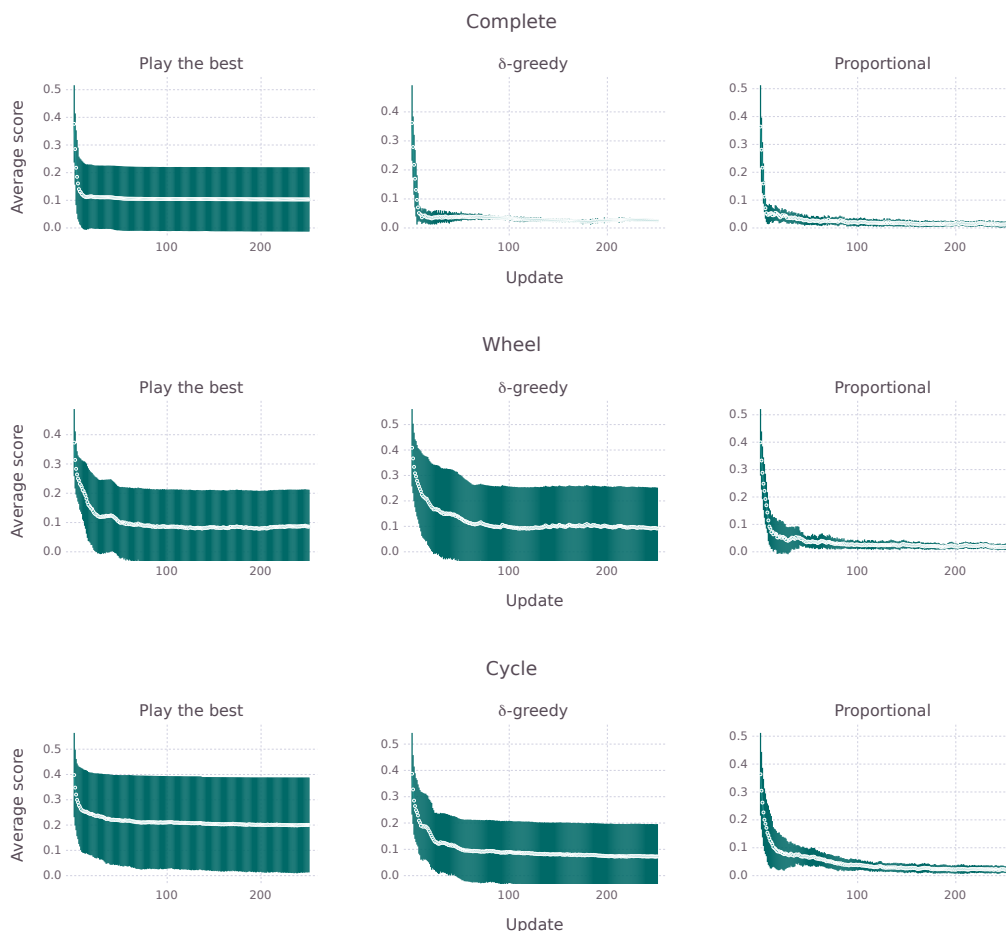


Figure 7: Average Euclidean distances from the truth, $\tau = \langle 0.5, 0.501 \rangle$, for 250 updates, with error bars showing one standard deviation from the mean. For all agents, $\alpha = 0.3$ and $\epsilon = 0.25$. Results for agents placed on a completely connected access network (top row), on a wheel access network (middle row), and on a cycle access network (bottom row), and for communities with agents always playing the arm they believe to be best (left column), playing a δ -greedy strategy (middle column), and playing a proportional strategy (right column).

than those the other network structures give rise to. That means that, typically, agents in a fully connected community have more accurate beliefs than agents in communities with fewer connections. Not only that: agents forming a wheel network were never less accurate than agents forming a cycle network, assuming the same choice rule.

For a more exact assessment, we ran a series of ANOVAs, with (where appropriate) follow-up t -tests for pairwise comparisons. Specifically, for the two sets of simulations—the one in which all agents had an α of 0.3 and an ϵ of 0.25 and the one in which they had an α of 0.1 and an ϵ of 0.05—we ran, for each choice rule separately, and per time step, an ANOVA with average scores as dependent variable and network type (complete/wheel/cycle) as independent variable and, whenever the results were significant, followed this up by a series of pairwise t -tests.

For the first parameter setting, this showed that the complete network always yielded sig-

$$\alpha = 0.1; \epsilon = 0.05; \tau = \langle 0.5, 0.501 \rangle$$

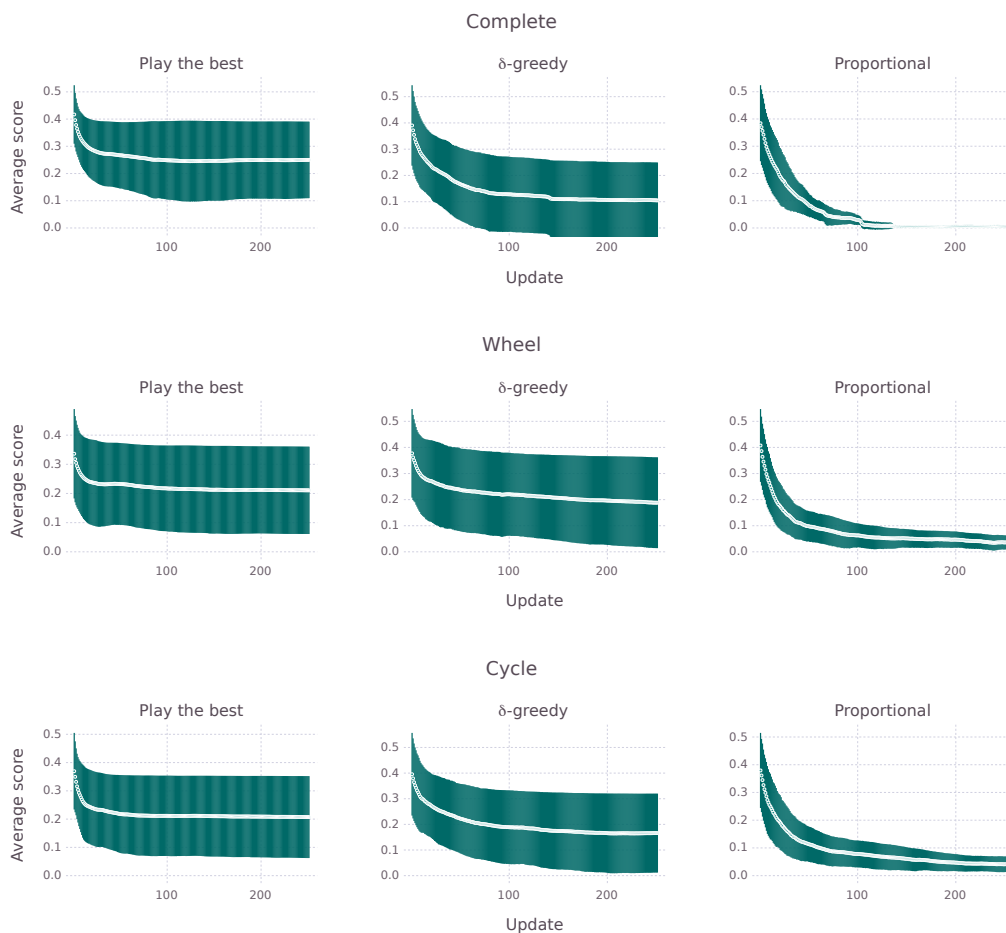


Figure 8: Same as Figure 7 but with $\alpha = 0.1$ and $\epsilon = 0.05$ for all agents.

nificantly more accurate results, regardless of the choice rule, all ANOVAs having an associated p -value smaller than 0.0001. For the play-the-best rule, follow-up tests also came out highly significant for all pairwise comparisons; for the δ -greedy rule this was true except for the comparison between the wheel and the cycle network, which was not significant; and for the proportional rule the t -tests were mostly significant.

For the second parameter setting, none of the ANOVAs for the play-the-best rule turned up significant, while for the δ -greedy rule the complete network gave significantly more accurate results from the eighteenth update onward; for the ANOVAs, all $ps < 0.05$, and after the 39-th update, all $ps < 0.0001$, while the follow-up tests were highly significant (most $ps < 0.001$) except for the wheel versus cycle comparison, for which none of the t -tests was significant. For the proportional rule, finally, results for the complete network were significantly more accurate from the eighth update onward; for the ANOVAs, all $ps < 0.05$, after the seventeenth update, all $ps < 0.0001$, while the follow-up tests came out highly significant as well (most $ps < 0.001$).

There is a clear trend in these results, which appears to indicate that, in our extension of the BC model, there is no analogue of the effect of connectivity on accuracy that was reported



Figure 9: Average scores for 50 runs of communities of 50 agents, with $\alpha = 0.3$ and $\epsilon = 0.25$, for each possible combination of choice rule and access network structure. The averages are first over the 50 agents per community and then over the 50 simulations that were run for the combination of choice rule used by the agents and access network imposed on their community.

in Zollman's (2010) and Kummerfeld and Zollman's (2016). Quite the opposite indeed: in our model accuracy tends to increase with more connections among agents. We have two comments on this finding.

First, we already mentioned Rosenstock and colleagues' observation that the network effect first discovered by Zollman only holds when the success rates of the two arms are very close to each other. We also mentioned the more general concerns about whether the results from agent-based modeling are sufficiently robust to support the kind of conclusions that researchers have wanted to draw from them. In the above, we assumed the exact same success rates as Zollman did and so our results might suffer from the same defect as his, and might also raise the same general concern. Indeed, our model has more parameters than Zollman's, so the question of how well our results generalize is even more pressing for us. In fact, we are not just interested in generalizations concerning the success rates of the arms and the values of the α and ϵ parameters; we also want to generalize the network structures beyond the three types used above. In Sections 5 and 6, we undertake a more systematic investigation of the new model in combination with the access network approach that is meant to directly address the aforementioned concerns about robustness.

Second, in Hegselmann and Krause (2005), which was the first paper in which these authors integrated truth-seeking (as a black box mechanism) into the BC model, the success of the agents' joint truth-seeking is measured via the square root of the mean square distance of the opinions to the true value τ .³ In this paper, we adapted Hegselmann and Krause's measure in




Figure 10: Same as Figure 9 but now for agents with $\alpha = 0.1$ and $\epsilon = 0.05$.

³This measure of success is also used in Zollman (2012), in which Zollman explicitly relates the approach from his earlier work (Zollman, 2007, 2010) to the BC model (see Zollman, 2012, p. 31 ff), although he gives few details about the actual simulations that he carried out in that model, which makes a direct comparison with our results

the seemingly most natural way to a two-dimensional opinion space, where one then still has the choice between taking the mean Euclidean distance or the mean Manhattan distance (both of which were defined above) of the opinions from the true value τ . Either way, we are using a measure different from the one used in Zollman (2010, 2012). In these papers, the measure of success is rather the proportion of the simulation runs which end with all members of the network playing the better arm.⁴ In Zollman’s setup, each agent always plays the arm that it currently considers to be the better one. The proportion of runs in which everyone ends up playing the better arm is then referred to by Zollman as the “probability of successful learning” (2007, p. 580) and used as a measure of success. To emphasize how different Zollman’s measure of success is from ours, we note that it is possible for an epistemic community to play the worse arm overall despite having a low mean Euclidean (or Manhattan) distance to truth. Similarly, it is possible for the community to have a relatively large mean Euclidean (or Manhattan) distance from the truth and yet uniformly play the better arm. The Supplementary Materials explains this in greater detail. There, we also look at convergence in the manner of Zollman, finding that, in our model, there is no noteworthy effect of network connectivity on convergence either.

5 Robustness

Hegselmann and Krause (2002) investigate the robustness of the most basic version of the BC model, with ϵ as its only parameter, by going in small steps through the range of that parameter. With the number of parameters we want to consider, such an approach is computationally well-nigh unfeasible. Therefore, we take our cue from Douven and Hegselmann (2021), who looked at a related extension of the BC model, which also had several parameters. To investigate the effect the individual parameters had on their outcome variable of interest (which also was accuracy, understood in terms of Euclidean distance from the truth), they ran thousands of simulations each of which used a random setting sampled from parameter space. Douven and Hegselmann collected the results from each simulation and then ran a regression analysis on the outcomes from all the simulations taken together. Thereby, they obtained a *general* understanding of how changes in the central parameters impacted the outcome variable.

For our present model, this same approach was implemented by running simulations that first picked at random each of the following:

1. one of the three choice rules: the play-the-best rule, the δ -greedy rule (always with $\delta = 0.1$), and the proportional rule;
2. a success rate $r_1 \sim \mathcal{U}(0, 1)$ for the first arm;
3. a success rate $r_2 \sim \mathcal{U}(0, 1)$ for the second arm;
4. an α value, with $\alpha \sim \mathcal{U}(0, 1)$;
5. an ϵ value, with $\epsilon \sim \mathcal{U}(0, 1)$;
6. a λ value, with $\lambda \sim \mathcal{U}(0, 20)$;
7. a random access network for a community of 50 agents, generated by creating a directed Erdős–Rényi graph with 50 nodes and with, for any given pair of nodes, a probability s for the first member of the pair to be connected to the second member, where $s \sim \mathcal{U}(0, 1)$.

We chose to make the access graphs *directed*—meaning that one node being connected to

difficult.

⁴In Zollman (2007), the simulation runs stop when all agents play the same arm or are certain with a probability greater than 0.9999 that one arm is the better one (see p. 579). In Zollman (2010), the simulation stops after 10,000 updates, after which it is checked whether all agents play the better arm or not (p. 28 ff).




Figure 11: Three directed Erdős–Rényi graphs, corresponding to different connection probabilities.

another is independent of whether the latter is connected to the former—to reflect the fact that, in reality, it is quite common for one person to have access to the opinions of another without the latter having access to the opinions of the former. Figure 11 shows three networks (for ten agents, to avoid clutter) generated in this manner, with different associated probabilities for an edge from one node to another to be present.⁵

Each simulation then had all agents update for one hundred time steps, in the way described earlier, but given the parameter settings and within the access network that had been randomly determined for that simulation. For each agent, we determined the score at each time (the Euclidean distance from what had been randomly determined to be the truth in that simulation), and in the end we summed all the scores, so across agents and across time steps. The simulation terminated by recording the parameter settings (so the values for r_1 , r_2 , α , ϵ , and λ , as randomly chosen at the beginning of the simulation), the connection probability that had been randomly chosen for generating the access network, and the total score.

We ran 5000 such simulations and then fitted a linear model with the 5000 resulting scores forming the response variable and with choice rule as a categorical predictor (play-the-best/ δ -greedy/proportional) and α value, ϵ value, λ value, absolute difference in success rates between the two arms, $|r_1 - r_2|$, and connection probability s , as continuous predictors. Table 1 shows the outcomes from the analysis.

A first thing to notice about these results is that all predictors are significant. Furthermore, they are in line with the single-run simulations reported above. In particular, we see that if the δ -greedy rule is used instead of the play-the-best rule—which served as the reference in the regression analysis—then this leads, on average, and all else being equal, to a reduction in total score (and so an increase in overall accuracy) of about 2. Using the proportional rule instead of the play-the-best rule leads to an even bigger—much bigger—average reduction of over 8.

Why does the absolute difference between the success rates of the two arms have a significant impact on scores? By looking at the plot of $|r_1 - r_2|$ in Figure 12, one sees immediately that if $|r_1 - r_2|$ is bigger, then if an agent’s belief state moves toward the truth, the frequency of playing one arm relative to the frequency of playing the other arm will go up, which will mean that it will receive increasingly more worldly evidence relevant to the success rate of the former arm, but increasingly less worldly evidence relevant to the success rate of the latter arm. By contrast, if $|r_1 - r_2|$ is small, then if an agent’s belief state moves toward the truth, it

⁵One could consider other types of graphs to define access networks. For instance, Schawe, Fontaine, and Hernández (2021) use in their simulations, next to Erdős–Rényi graphs, also Watts–Strogatz graphs as well as Barabási–Albert graphs. The Julia code we made available allows interested readers to explore the effects of adding access networks represented by Watts–Strogatz and Barabási–Albert graphs, as well as by some other graph types. In our experiments, using other graphs did not lead to interestingly different results.




Figure 12: Difference in success rates between arms.

will receive worldly evidence about the success rates of the two arms in about equal measures, which will allow it to become highly accurate over time (and so lower its score). In the latter case, the agent will be able to improve its estimate of both arms over time, while in the former case, the agent will rapidly improve its estimate of the success rate of one of the arms, but that will go at the expense of improving its estimate of the success rate of the other arm, which means that even in the end, the agent may, overall, be quite inaccurate.

That, in the model we are considering, updating is not just based on worldly evidence, but also has a social part to it, amplifies the said effect. Figures 13 and 14 show, for selected time steps, how the update process unfolds for a community of fifty agents using (BCC). All agents use the play-the-best rule to determine which arm to play, and all have an α value of 0.2, an ϵ value of 0.15, and a λ value of 2. The agents' belief states had the same random starting positions in both runs, and both runs also assumed the same access network (which was created by generating an Erdős–Rényi directed graph with connection probability 0.5). The difference between the two figures is that, in the first, $\tau = \langle 0.2, 0.8 \rangle$, while, in the second, $\tau = \langle 0.2, 0.21 \rangle$. Comparing the figures, we see that where there is a large gap between the

Table 1: Results from the regression analysis of simulation outcomes.

Predictor	b	SE	t	$\Pr(> t)$	Lower 95 %	Upper 95 %
(Intercept)	22.23	0.50	44.06	<0.0001	21.24	23.22
Rule: 2	-2.08	0.29	-7.05	<0.0001	-2.65	-1.50
Rule: 3	-8.15	0.29	-27.78	<0.0001	-8.73	-7.58
$ r_1 - r_2 $	17.20	0.51	33.63	<0.0001	16.19	18.20
α	-9.08	0.42	-21.56	<0.0001	-9.90	-8.25
ϵ	-6.93	0.42	-16.38	<0.0001	-7.76	-6.10
λ	0.51	0.02	24.40	<0.0001	0.47	0.56
s	-1.02	0.41	-2.45	<0.05	-1.83	-0.20

Note. “Rule: 2” compares the δ -greedy rule to the play-the-best rule, which served as reference, and “Rule: 3” does the same for the proportional rule; $|r_1 - r_2|$ is the absolute difference between the success rates of the two arms; s is the probability with which a given node in an access network is connected to another given node.

$$\alpha = 0.2; \epsilon = 0.15; \lambda = 2; \tau = \langle 0.2, 0.8 \rangle$$

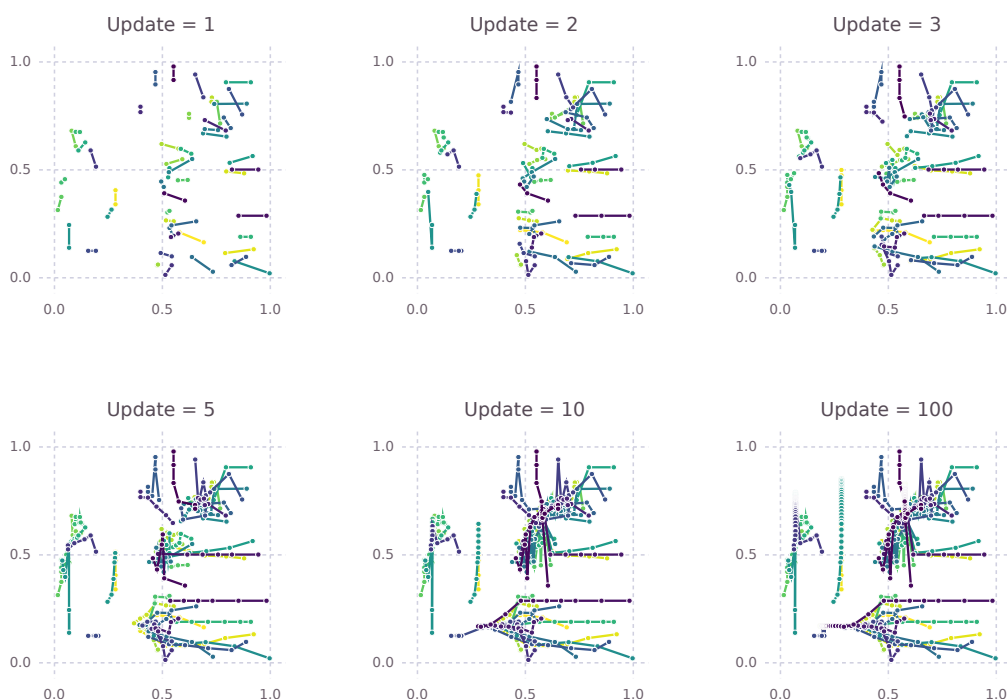


Figure 13: Trajectories of belief states of 50 agents updating via (BCC), and using the play-the-best rule, on data from two-armed bandits, after selected time steps, with the success rates of the arms being far apart from each other.

success rates, groups of agents can more easily become isolated, and thereby become unable to benefit from others whose evolving belief states might otherwise help them to move toward τ .

Furthermore, the coefficients for α and ϵ indicate that giving more weight to one's neighbors' opinions, and being more inclusive in counting others as one's neighbor, has a positive effect on accuracy. This finding is consistent with all previous findings about the relation between an agent's α and ϵ values and its accuracy in the context of BC updating (Douven & Wenmackers, 2017; Douven, 2019, 2022a, 2022b; Douven & Hegselmann, 2021). On the other hand, λ , while significant, has only a minor impact on scores.

Finally, there appears to be a relatively small yet significant network effect, which is opposite to the one found by Zollman,⁶ but in line with what we saw above: greater connectivity tends to lower scores, or equivalently, tends to increase accuracy.

6 Individual-level parameters

So far, α , ϵ , and λ parameters were set at the community level. However, given that we may be differently inclined to regard others as our epistemic neighbors—people whose opinions we value—and also differently inclined to balance worldly evidence and our epistemic neighbors' opinions in how they impact our belief changes, a more realistic formal model of BC updating

⁶Recall, however, that Zollman is assuming a measure of success different from ours.

$$\alpha = 0.2; \epsilon = 0.15; \lambda = 2; \tau = (0.2, 0.21)$$




Figure 14: Same as Figure 13 but now with the success rates of both arms being near to each other.

allows agents to have their own α , ϵ , and λ values.

Equally as important, in going beyond the community level we can also take into consideration properties of the access network at a finer grain of detail. In the previous simulations, we took into account a global property of access networks, viz., their connectivity, that is, the probability to find a link from one node to another, for any pair of nodes. But within the same network, different nodes can have different properties, depending on how they are connected to other nodes. In particular, there is a whole range of measures all aimed at capturing the idea that a node can be “central” to a network to varying degrees. In special cases, the idea of centrality is easy to understand. For instance, in another look at Figure 6, the node located at the center of the wheel is in an intuitively clear sense more central to the network than any of the other nodes: it is in direct “touch” with every node in the network (recall that the networks are reflexive but that this is not indicated in the plots) and vice versa (recall that the absence of arrow tips indicates that the graphs are undirected); the same is not true for any of the other nodes in the wheel. By contrast, both in the complete network and in the cycle, each node is as central to the network as any other node: in the former, every node is directly connected to every other node, while in the latter, every node is directly connected to just two other nodes besides itself, is not more than two connections removed from four other nodes (five nodes, itself included), not more than three connections from six other nodes, and so on. Given a more realistic network, intuitions are not as clear, and different measures of centrality might give different verdicts on how central a node is to a network.

The possibly simplest measure is that of *degree centrality*, which for directed graphs splits into two: the *indegree centrality* of a given node is the number of nodes that link to that node,

Table 2: Results from the regression analysis of simulation outcomes.

Predictor	b	SE	t	$\Pr(> t)$	Lower 95 %	Upper 95 %
(Intercept)	17.95	0.07	241.21	<0.0001	17.80	18.10
α	-1.41	0.04	-31.56	<0.0001	-1.50	-1.32
ϵ	-2.33	0.04	-52.18	<0.0001	-2.42	-2.24
λ	0.19	0.00	82.79	<0.0001	0.18	0.19
closeness	-5.07	0.18	-28.59	<0.0001	-5.42	-4.73
betweenness	-7.10	0.93	-7.67	<0.0001	-8.91	-5.28
indegree	0.19	0.16	1.21	=0.23	-0.12	0.51
outdegree	2.20	0.20	11.29	<0.0001	1.82	2.59

while the *outdegree centrality* of a node is the number of nodes the node links to. Another commonly used measure is *closeness centrality*, which is a measure of the average shortest path length from a node to the other nodes in the network; to be exact, the closeness centrality of a given node is the total number of nodes in the network diminished by one, all over the sum of the lengths of the shortest paths from the given node to the other nodes. A third popular measure, and the final one we look at, is that of *betweenness centrality*. This measure indicates for a given node on how many shortest paths between pairs of other nodes in the network it occurs.

We ran again 5000 simulations with communities of 50 agents updating via (BCC) on data from a two-armed bandit, trying to determine the success rates of the two arms. This time, each simulation started by selecting

1. a success rate $r_1 \sim \mathcal{U}(0, 1)$ for the first arm;
2. a success rate $r_2 \sim \mathcal{U}(0, 1)$ for the second arm;
3. for each agent i , an α_i value, with $\alpha_i \sim \mathcal{U}(0, 1)$;
4. for each agent i , an ϵ_i value, with $\epsilon_i \sim \mathcal{U}(0, 1)$;
5. for each agent i , a λ_i value, with $\lambda_i \sim \mathcal{U}(0, 20)$;
6. a random access network for 50 agents, again formalized as a directed Erdős–Rényi graph, with connection probability $s \sim \mathcal{U}(0, 1)$.

Each simulation consisted again of 100 time steps, at each of which an agent chose to play one of the two arms, the choice always determined on the basis of the proportional choice rule, which came out on top in our previous simulations.

In each simulation, we registered, for all agents i , their α_i , ϵ_i , and λ_i values, their in- and outdegree centrality, their closeness and betweenness centrality, and finally, their total score over the 100 time steps, that is, the sum of the Euclidean distances from $\tau = \langle r_1, r_2 \rangle$ measured at each time step.

These scores gave us 25,000 data points (50 agents \times 5000 simulations), which served as the response variable for a regression analysis with agents' α_i , ϵ_i , and λ_i values as well as their various centrality indices as predictor variables. This yielded the outcomes shown in Table 2.

We see that per-agent α and ϵ values still have a highly significant negative impact on scores, and hence a highly significant positive impact on accuracy, albeit that their effect sizes are smaller than the corresponding population-level parameters taken into account in the previous analyses. Furthermore, the per-agent λ values have a highly significant positive effect on scores, just like the population-level λ -values were seen to have. More interestingly, while at the population level the network effects were somewhat unexciting, here we see that, first,

apart from indegree centrality, all centrality indices had a highly significant impact on scores, and second, closeness and betweenness centrality actually had by far the *biggest* impact of all the predictors. As for indegree centrality, it is perhaps not too surprising that how many community members have access to your belief state has no reliable effect on your accuracy. By contrast, the number of community members whose belief states you have access to *would* seem to matter to your accuracy, and indeed outdegree centrality does have a significant impact. Nonetheless, the direction of the effect is at first puzzling: outdegree centrality appears to *positively* affect score, and hence to *negatively* affect accuracy. At the population level, more densely connected access network tended to improve accuracy. Now we seem to find the opposite at the individual level.

The mystery is resolved by looking at smaller regression models. When we remove from the above model as predictors all centrality indices except outdegree centrality, the latter is again highly significant but now does impact score negatively, in line with what one would expect ($b = -0.91, t = -20.79, p < 0.0001$). If, however, we add to this smaller model closeness centrality and in particular also the interaction between outdegree centrality and closeness centrality, we find that all predictors as well as the interaction term ($b = 0.74, t = 2.72, p < 0.01$) are significant but now the coefficient of outdegree centrality is again positive ($b = 1.79, t = 5.35, p < 0.0001$). This indicates that being well-connected does help to improve one's accuracy, but *given that one is well-connected*, it is better to have access to fewer fellow agents.

7 Conclusion

In this paper, we looked at a version of the BC model from the perspective of network epistemology. While dynamic impact networks had (implicitly) been part of the BC model from the very first start, the same is not true for the kind of static access networks that are central to much recent work in network epistemology. That could be regarded as a shortcoming, given that it is unrealistic to suppose that there are no constraints on the access agents have to the belief states of other agents. This motivated us to add static access networks as a new layer to the BC model.

We focused on a version of the BC model that features agents which update on worldly evidence via a Carnapian λ rule. This model allowed us to study the bandit problems that occupy center stage in Zollman's influential work on network epistemology. We were especially interested in whether the surprising effect that Zollman reported for Bayesian agents organized in an access network—the finding that better connected networks may sometimes lead to worse epistemic outcomes—has an analogue in BC models featuring social Carnapian agents.

We did find network effects, both at a community level and at an individual level. As to the former, we found that the accuracy of communities as a whole tends to be affected by the density of their access network. As to the latter, we found that the accuracy of agents can be affected by their place in an access network, notably their centrality (how well-connected they are). But the effects were in the direction that is more in line with expectations: greater connectivity, both at the community level and at the individual level, tended to be beneficial to accuracy.

Why did we fail to replicate Zollman's finding in our model? First, as we said, others discovered that the found effect was not robust to begin with; even in Zollman's own model, it occurs only under very special circumstances, circumstances which, for most practical purposes, are of limited interest at best. But second, if we look at Zollman's explanation of his finding—agents whose belief states are “off” have, in sparser networks, a lower chance of

negatively affecting agents whose belief states are closer to the truth—then we realize that the dynamic impact networks that were all along implicit in the BC model provide a similar protection. Just look back at Figure 1. Especially the left panel very clearly shows the clustering process that finds place as the agents update their estimates of τ . For quite a while, agents starting with random guesses at a distance from τ will not be able to impact the estimates of agents starting with random guesses close to τ ; that happens only after sufficient corrective exposure of the former agents to worldly evidence. But if the dynamic impact networks in our model offer the kind of protection that Zollman’s sparser access networks offer in his, then it becomes an empirical question what the effect of placing access networks on top of communities relying on BC updating will be. We used computer simulations to answer that question, the answer being that there is no additional benefit, and that more densely connected access network tend to be better than sparser ones, in case they are combined with BC updating.

Should the general conclusion be that, to the extent that the BC model in the version explored in this paper is a good approximation of how real people update, in response to a continuing inflow of worldly evidence as well as interactions with others, we should make an effort to make the access network of our community as close to complete as we can? Not necessarily. For consider that you may have to invest some time and perhaps even money in order to get, and remain, in touch with others (Zollman, 2013). Thus, we should ask whether the possible rewards we reap from being better connected are worth the costs. We have not addressed this question in the above and flag it here as an important avenue for future research.⁷

References

- Ahat, M., Bennani, O., Bui, M. & Amor, S. B., (2011). Consensus and relation networks. In G. Eichler, A. Küpper, V. Schau, H. Fouchal, & H. Unger (eds.), *11th International conference on innovative internet community systems* (pp. 26–37). Bonn: Gesellschaft für Informatik.
- Bala, V. & Goyal, S. (1998). Learning from neighbours. *Review of Economic Studies*, 65, 595–621.
- Berry, D. A. & Fristedt, B. (1985). *Bandit problems: Sequential allocation of experiments*. London: Chapman and Hall.
- Bertsekas, D. P. & Tsitsiklis, J. N. (2008). *Introduction to probability* (2nd ed.). Belmont MA: Athena Scientific.
- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, 59, 65–98.
- BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge MA: Harvard University Press.
- Borg, A., Frey, D., Straßer, C., & Šešelja, D. (2019). Theory-choice, transient diversity and the efficiency of scientific inquiry. *European Journal for Philosophy of Science*, 9, 26, <https://doi.org/10.1007/s13194-019-0249-5>.
- Carnap, R. (1952). *The continuum of inductive methods*. Chicago: University of Chicago Press.
- Chazelle, B. (2011). The total s -energy of a multiagent system. *SIAM Journal on Control and Optimization*, 49, 1680–1706.
- Chen, G. & Lou, Y. (2019). *Naming game: Models, simulations and analysis*. Cham: Springer.
- Crosscombe, M. & Lawry, J. (2016). A model of multi-agent consensus for vague and uncertain beliefs. *Adaptive Behavior*, 24, 249–260.
- Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interact-

⁷We are greatly indebted to two anonymous referees for valuable comments on a previous version.

- ing agents. *Advances in Complex Systems*, 3, 87–98.
- De Langhe, R. (2013). Peer disagreement under multiple epistemic constraints. *Synthese*, 190, 2547–2556.
- DeMarzo, P. M., Vayanos, D., & Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118, 909–968.
- Douven, I. (2010). Simulating peer disagreements. *Studies in History and Philosophy of Science*, 41, 148–157.
- Douven, I. (2019). Optimizing group learning: An evolutionary computing approach. *Artificial Intelligence*, 275, 235–251.
- Douven, I. (2022a). *The art of abduction*. Cambridge MA: MIT Press.
- Douven, I. (2022b). Explaining the success of induction. *British Journal for the Philosophy of Science*, in press.
- Douven, I. & Hegselmann, R. (2021). Mis- and disinformation in the bounded confidence model. *Artificial Intelligence*, 291, 103415.
- Douven, I. & Kelp, C. (2011). Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis*, 75, 271–283.
- Douven, I. & Riegler, A. (2010). Extending the Hegselmann–Krause model I. *Logic Journal of the IGPL*, 18, 323–335.
- Douven, I. & Wenmackers, S. (2017). Inference to the best explanation versus Bayes’ rule in a social setting. *British Journal for the Philosophy of Science*, 68, 535–570.
- Foley, R. (1993). *Working without a net*. Oxford: Oxford University Press.
- Fortunato, S. (2004). The Krause–Hegselmann consensus model with discrete opinions. *International Journal of Modern Physics C*, 15, 1021–1029.
- Frey, D. & Šešelja, D. (2018). What is the epistemic function of highly idealized agent-based models of scientific inquiry? *Philosophy of the Social Sciences*, 48, 407–433.
- Frey, D. & Šešelja, D. (2020). Robustness and idealizations in agent-based models of scientific interaction. *British Journal for the Philosophy of Science*, 71, 1411–1437.
- Fricker, E. (2019). Knowing full well from testimony? *Episteme*, 16, 369–384.
- Friedkin, N. E., Proskurnikov, A. V., Tempo, R., & Parsegov, S. E. (2016). Network science on belief system dynamics under logic constraints. *Science*, 354, 321–326.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Golub, B. & Jackson, M. O. (2010). Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2, 112–149.
- Grim, P., Singer, D. J., Reade, C., & Fisher, S. (2015). Germs, genes, and memes: Function and fitness dynamics on information networks. *Philosophy of Science*, 82, 219–243.
- Hahn, U., Hansen, J. U., & Olsson, E. J. (2020). Truth tracking performance of social networks: How connectivity and clustering can make groups less competent. *Synthese*, 197, 1511–1541.
- Hegarty, P. & Wedin, E. (2016). The Hegselmann–Krause dynamics for equally spaced agents. *Journal of Difference Equations and Applications*, 22, 1621–1645.
- Hegselmann, R. (2004). Opinion dynamics: Insights by radically simplifying models. In D. Gillies (ed.) *Laws and models in science* (pp. 19–44). London: King’s College Publications.
- Hegselmann, R. (2020). Polarization and radicalization in the bounded confidence model: A computer-aided speculation. In V. Buskens, R. Corten, & C. Snijders (eds.) *Advances in the sociology of trust and cooperation: Theory, experiment, and field studies* (pp. 197–226). Berlin: De Gruyter.
- Hegselmann, R., König, S., Kurz, S., Niemann, C., & Rambau, J. (2015). Optimal opinion

- control: The campaign problem. *Journal of Artificial Societies and Social Simulation*, 18, <http://jasss.soc.surrey.ac.uk/18/3/18.html>.
- Hegselmann, R. & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis, and simulations. *Journal of Artificial Societies and Social Simulation*, 5, <http://jasss.soc.surrey.ac.uk/5/3/2.html>.
- Hegselmann, R. & Krause, U. (2005). Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25, 381–405.
- Hegselmann, R. & Krause, U. (2006). Truth and cognitive division of labor: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation*, 9, <http://jasss.soc.surrey.ac.uk/9/3/10.html>.
- Hegselmann, R. & Krause, U. (2009). Deliberative exchange, truth, and cognitive division of labour: A low-resolution modeling approach. *Episteme*, 6, 130–144.
- Hegselmann, R. & Krause, U. (2015). Opinion dynamics under the influence of radical groups, charismatic leaders, and other constant signals: A simple unifying model. *Networks and Heterogeneous Media*, 10, 477–509.
- Hegselmann, R. & Krause, U. (2019). Consensus and fragmentation of opinions with a focus on bounded confidence. *American Mathematical Monthly*, 126, 700–716.
- Jacobmeier, D. (2004). Multidimensional consensus model on a Barabási–Albert network. *International Journal of Modern Physics C*, 16, 633–646.
- Krause, U. (2000). A discrete nonlinear and non-autonomous model of consensus formation. In S. Elaydi, G. Ladas, J. Popenda, & J. Rakowski (eds.) *Communications in difference equations* (pp. 227–236). Amsterdam: Gordon and Breach.
- Krause, U. (2015). *Positive dynamical systems in discrete time: Theory, models, and applications*. Berlin: De Gruyter.
- Kummerfeld, E. & Zollman, K. J. S. (2016). Conservatism and the scientific state of nature. *British Journal for the Philosophy of Science*, 82, 956–968.
- Kurz, S. (2015). Optimal control of the freezing time in the Hegselmann–Krause dynamics. *Journal of Difference Equations and Applications*, 21, 633–648.
- Kurz, S. & Rambau, J. (2011). On the Hegselmann–Krause conjecture in opinion dynamics. *Journal of Difference Equations and Applications*, 17, 859–876.
- Lackey, J. (2015). Reliability and knowledge in the epistemology of testimony. *Episteme*, 12, 203–208.
- Lehrer, K. (1974). *Knowledge*. Oxford: Oxford University Press.
- Lorenz, J. (2003). *Mehrdimensionale Meinungsdynamik bei wechselndem Vertrauen*. Diploma thesis, University of Bremen, available at <http://nbn-resolving.de/urn:nbn:de:gbv:46-dipl000000564>.
- Lorenz, J. (2008). Fostering consensus in multidimensional continuous opinion dynamics under bounded confidence. In D. Helbing (ed.) *Managing complexity: Insights, concepts, applications* (pp. 321–334). Berlin/Heidelberg: Springer.
- O'Connor, C. & Weatherall, J. O. (2017). Scientific polarization. *European Journal for Philosophy of Science*, 8, 855–875.
- O'Connor, C. & Weatherall, J. O. (2019). *The misinformation age: How false beliefs spread*. New Haven CT: Yale University Press.
- Olsson, E. J. (2008). Knowledge, truth, and bullshit: Reflections on Frankfurt. *Midwest Studies in Philosophy*, 32, 94–110.
- Parasnis, R., Franceschetti, M., & Touri, B. (2018). Hegselmann–Krause dynamics with limited connectivity. *2018 IEEE Conference on Decision and Control*, 5364–5369.

- Parasnis, R., Franceschetti, M., & Touri, B. (2021). On the convergence properties of social Hegselmann–Krause dynamics. *IEEE Transactions on Automatic Control*, <https://doi.org/10.1109/TAC.2021.3052748>.
- Pluchino, A., Latora, V., & Rapisarda, A. (2006). Compromise and synchronization in opinion dynamics. *European Physical Journal B*, 50, 169–176.
- Riegler, A. & Douven, I. (2009). Extending the Hegselmann–Krause model III: From single beliefs to complex belief states. *Episteme*, 6, 145–163.
- Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In epistemic networks, is less really more? *Philosophy of Science*, 84, 234–252.
- Schawe, H., Fontaine, S., & Hernández, L. (2021). Network bridges foster consensus: Bounded confidence models in networked societies. *Physical Review Research*, 3, 023208.
- Šešelja, D. (2022). Some lessons from simulations of scientific disagreements. *Synthese*, in press.
- Stone, M. (1961). The opinion pool. *Annals of Mathematical Statistics*, 32, 1339–1342.
- Wenmackers, S., Vanpoucke, D., & Douven, I. (2012). Probability of inconsistencies in theory revision: A multi-agent model for updating logically interconnected beliefs under bounded confidence. *European Physical Journal B*, 85, <https://doi.org/10.1140/epjb/e2011-20617-8>.
- Wenmackers, S., Vanpoucke, D., & Douven, I. (2014). Rationality: A social-epistemology perspective. *Frontiers in Psychology*, 5, 581, <https://doi.org/10.3389/fpsyg.2014.00581>.
- Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74, 574–587.
- Zollman, K. J. S. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72, 17–35.
- Zollman, K. J. S. (2011). Social network structure and the achievement of consensus. *Politics, Philosophy & Economics*, 11, 26–44.
- Zollman, K. J. S. (2013). Network epistemology: Communication in epistemic communities. *Philosophy Compass*, 8, 15–27.