

Exploring Opinion Diversity and Epistemic Success with an Argumentative Model

Louise Dupuis de Tarlé¹, Gabriella Pigozzi¹, and Juliette Rouchier¹

LAMSADE, Université Paris-Dauphine PSL
louise.dupuis@dauphine.eu
gabriella.pigozzi@lamsade.dauphine.fr
juliette.rouchier@lamsade.dauphine.fr

Abstract. Abstract Argumentation Frameworks are a useful tool for modelling argumentative dynamics in agent-based models, and the recent introduction of gradual argumentation semantics expands their potential. Our research focuses on using these tools to examine an epistemic community of agents who use the result of their experimentation to produce arguments and exchange them with their peers. The central objectives are to analyze the stability and convergence of opinions among agents and evaluate individual and collective epistemic success.

Keywords: Opinion Diffusion · Social Epistemology · Argumentation.

1 Introduction

Social Epistemology is a branch of philosophy which explores the impact of social life and organisation on the pursuit of truth [9]. Agent-based simulation is a compelling tool for examining this relationship. Philosophers have recently used argumentation to understand the group dynamics of truth-seeking agents [13]. By modelling argumentative processes, we aim to analyse and understand the effect of some features of group discussion on the epistemic outcome.

We present a model that depicts an epistemic community, a group of agents seeking to uncover the truth regarding a specific issue. The agents engage in a dynamic argumentative discussion to discover this truth. It is assumed that experience grants them access to certain facets of reality and that the discussion entails a repeated exchange of conflicting and supportive arguments. The assumption that reasoning is done through the exchange of arguments is in line with the argumentative theory of reasoning defended by Mercier [14]. The arguments, organized in abstract argumentation frameworks and equipped with a novel gradual semantic, constitute the information base of the agents and can be shared between them in a mechanism of social influence.

Our model shows many similarities with opinion diffusion models: like in the seminal bounded confidence-type models [16, 6], agents' opinions are numbers between 0 and 1, and agents influence each other through pairwise interaction. One major novelty of our model is that such behavior is modelled via the explicit

representation of information as an organized set of objects, the arguments, which can be generated and transmitted between agents.

Our model merges distinct areas of multi-agent simulation.

The foundations of social epistemology, which studies the knowledge and beliefs of groups of individuals, were introduced by Goldman [9]. Subsequently, many works have used multi-agent systems to study scientific communities as examples of epistemic communities. [23] and [15] use models of Bayesian agent networks that interact with an environment and exchange the results of their experiments.

With the development of abstract argumentation as an effective way to represent debates as graphs, many agent-based models have been created to simulate argumentative discussions [12, 2, 20, 3]. Our work is partially based on [22], which presents a study of the dynamics of a gradual semantics, a type of semantics applied to abstract argumentation frameworks recently proposed in the literature [1]. The representation of knowledge among agents is similar to that of [22], although we introduce the notions of graph generation and interaction between agents and the environment. We use a different type of argumentation framework: bipolar graphs, and propose a novel semantic for these graphs. [22] show that when agents learn each other’s arguments, their opinions converge. Our model challenges these results, as we show that exchanging arguments can also lead to a greater diversity of opinions.

In our model, agents discuss the validity of a central question, the *issue*, and their opinions vary between 0 and 1. A major assumption of our model is that we assume that the *issue* and all subsequently generated arguments are characterized by a *truth value* in the interval $[0, 1]$. This way of modelling the truth has precedents: [10] uses a similar truth value to extend the bounded confidence model to a community of truth-seekers. Fuzzy logic uses continuous truth values interpreted as partial truths.

Our modelling approach is outlined in three sections. The first provides a brief overview of the formal tools utilized: abstract argumentation and the gradual semantic we propose. The second section outlines our model and protocol. The third section delves into two crucial aspects of our model: the evaluation of epistemic success (i.e. the ability to approach the truth), and examination of opinion convergence. We show that the relation between the number of connections between agents and the diversity of their opinions is not monotonous, and, departing from the seminal results of Bayesian models, that opinion diversity is harmful for their epistemic success.

2 Abstract Argumentation

2.1 Bipolar Argumentation Frameworks

A bipolar abstract argumentation framework (BipoAF), formally $\mathbf{B} = \langle \mathcal{A}, \mathcal{R}, \mathcal{S} \rangle$, consists of a finite set of arguments \mathcal{A} , equipped with two binary relations on this set: the attack relation \mathcal{R} and the support relation \mathcal{S} [5]. BipoAFs can be

represented by graphs whose nodes are the arguments and edges represent the relation between them. In the rest of this article, we will use the terms BipoAF and bipolar argumentation graphs independently, and often refer simply to “argumentation graphs”. Figure 1 presents a small debate made up of five arguments, and their representation as an argumentation graph with the relations that binds them: we used full lines to represent an attack, and dashed lines to represent support. It can be noted that the *abstract* argumentation graphs say nothing in the general case of the content of the arguments but concern themselves with describing the general structure that the debate forms.

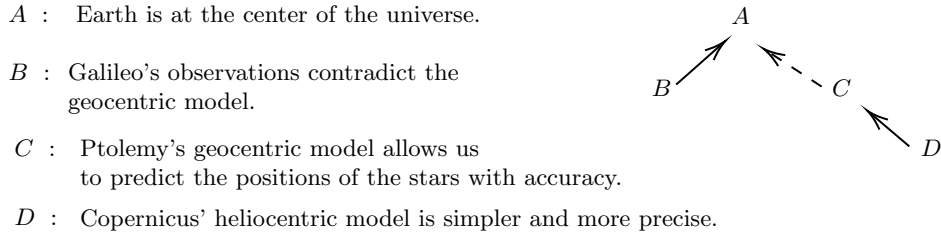


Fig. 1: Example of a debate and its representation as a bipolar abstract argumentation graph.

Abstract argumentation is especially concerned with the question of the *acceptability* of arguments. From a given argument graph, the **semantics** are functions that determine the acceptability of the arguments: here, acceptable can be understood as ‘rationally defensible’. Recently, new semantics have been proposed to allow a finer evaluation of the concept of acceptability: *gradual* semantics (see [1]) assign a **score** to each argument.

2.2 Logistic Sum Semantic for Bipolar Graphs

[21] presents a novel gradual semantic for bipolar argumentation graphs, the **Logistic Sum** gradual semantic (LSS). This function, which they justify as being well adapted to represent the opinion of agents, verifies a number of desirable properties and assigns to every argument a **score** between 0 and 1. Their semantic is defined for *weighted* bipolar argumentation graphs, where arguments are equipped with *weights*; numerical values which represent an intrinsic quality of the argument. We use the version of their semantic which is adapted for non-weighted graphs, by assuming a weight of 0.5 for all of our arguments.

Definition 1 (Non-weighted Logistic Sum Semantic). Let $\mathbf{B} = \langle \mathcal{A}, \mathcal{R}, \mathcal{S} \rangle$ be an acyclic bipolar graph and $a \in \mathcal{A}$ one of its arguments. LSS is the score function $\mathfrak{Z}(a)$ recursively defined as :

$$\sigma(B, a) = 1 - \frac{1}{1 + e^{E(B, a)}}, \quad E(B, a) = \sum_{x \in \text{Supp}(B, a)} \sigma(x) - \sum_{x \in \text{Att}(B, a)} \sigma(x) \quad (1)$$

with $Supp(a)$ and $Att(a)$ the sets of arguments which are respectively supporting and attacking a .

We see that arguments which are neither attacked nor supported, or whose attacks are compensated by supports, have a score of 0.5. The more attacked an argument is, the closer its score is to 0, and conversely, arguments that are well supported have a score that tends towards 1. This semantic thus allows us to derive, for each agent, an opinion about an issue which is a number between 0 and 1 based on what they know of the attack and support towards this issue. The following section details this mechanism.

3 The Model

3.1 Setting

In our model, agents generate arguments which form bipolar graphs. These BipoAF, like those introduced by [22], are *issue-oriented* argumentation graphs. A special argument, the *issue*, is the root of the graph, and all other arguments are part of a path to this *issue*. Intuitively, the issue represents the central proposition of the debate, and all agents' moves are aimed at (indirectly) attacking or supporting that issue.

Each agent $k \in N$ is equipped with her own argumentation graph $\mathcal{O}_k = \langle \mathcal{A}_k, \mathcal{R}_k, \mathcal{S}_k \rangle$, which will contain the argument that she generates and those that she receives from her neighbours. This graph is called the **opinion graph** of the agent. The agents apply the non-weighted Euler based semantic to their opinion graph and derive a *score* $\mathfrak{J}_k(a)$ for each argument $a \in \mathcal{A}_k$. The score of the issue is the **opinion** of the agent, o_k .

Each argument a is also equipped with an intrinsic **truth value** T_a . The agents cannot access these truth values. The truth value of the issue, T_i , is of particular interest to our model and is sometimes referred to as the truth value of the debate. We measure the epistemic success of the agents by how distant their opinion (assessment of the issue) is from T_i .

Network Communication Agents are placed in a **social network** \mathcal{N} , and share all the arguments that they produce with their direct neighbours. We use Erdos-Renyi non-directed random graphs [7] to model this network: a network \mathcal{N} is characterized by the number of agents N and p_{ER} , the probability that two agents are linked.

3.2 Protocol

At the beginning, all agent's opinion graphs are composed only of the issue. Therefore, the initial opinion of all agents is the same: $\forall k \in N, o_k = 0.5$.

Each step t of our protocol is composed of two main stages.

Stage 1 : Argument Generation

- Every agent k chooses randomly an argument a from her opinion graph.
- The agent investigates the chosen argument and, according to the process detailed below, generates a new argument¹ a' and either an attack or a support relation from a' towards a .

Stage 2 : Argument Sharing

- Every agent adds to their own opinion graph the argument and relation they generated.
- Every agent adds to their own opinion graph the argument and relations generated by their neighbours in \mathcal{N}^2 .

This two-stage step allows us to simulate the simultaneous activation of the agents. Note that the opinion graph (and thus the opinion) of the agents changes due to their own generation of argument as well as that of their neighbors. This means that even in situations where agents do not communicate with each other, their opinions evolve.

Argument Generation The generation of attacks and supports is inspired by the one-armed bandit mechanisms used in Bayesian networks of agents [23, 17].

In our argument generation method, agents sample a result from a distribution centered on the investigated argument's truth value. A natural interpretation is that the sampled result obtained corresponds to an "experiment" with the world. The result of this experiment is compared to the agent's previous perception of the investigated argument. If it is higher, the result has exceeded the agent's expectation, and the agent will produce a support for the investigated argument. Conversely, if it is lower, the agent will produce an attack of the investigated argument. The agents only have direct access to their sampled results and their own opinion.

Formally, let us consider agent k , whose opinion graph is $\mathcal{O}_k = \langle \mathcal{A}_k, \mathcal{R}_k, \mathcal{S}_k \rangle$. Agent k investigates argument $a \in \mathcal{A}_k$, let T_a be the truth value of a , and $\mathfrak{Z}_k(a)$ the score of a in the agent's opinion graph. We introduce σ , the parameter which controls the standard deviation of Gaussian distributions.

1. The agent samples r_k from a Gaussian distribution $G_1 = \Gamma(T_a, \sigma)$. Argument a' is created. If $r_k > \mathfrak{Z}_k(a)$, the agent generates a support relation $(a', a) \in \mathcal{S}_k$ otherwise she generates an attack relation $(a', a) \in \mathcal{R}_k$.
2. T'_a is sampled from $\begin{cases} G_2 = \Gamma(T_a, \sigma) & \text{if } (a', a) \text{ is a support} \\ G_2 = \Gamma(1 - T_a, \sigma) & \text{if } (a', a) \text{ is an attack} \end{cases}$ restricted to the $[0, 1]$ interval.

¹ Each new generated argument is unique, and identified by a unique id; two agents may both generate an argument attacking a , which will translate as two new arguments attacking a . Note that there is no limit as to the number of attacks and supports an argument can have.

² Cases of seemingly conflicting arguments (eg if an agent receives both a support and an attack towards the same argument from her neighbors) do not pose any problem, since the gradual semantic is built to aggregate multiple supports and attacks.

This mechanism has a *corrective* effect: if the agent’s opinion of the argument is higher than its true strength, the agent will likely produce an attack and therefore lower her own opinion of the argument, and vice versa. Therefore, this mechanism improves each agent’s perception of the argument’s strength, and its precision is controlled by σ . For this reason, we consider σ to be the agent’s **experimental accuracy**.

The truth value of a supporter is similar to that of the supported argument; for attackers, the relation is reversed. Although we do not claim that this property would at any rate be verified by real-life argumentative discussions, it is a feature of our model because we wish to represent epistemic investigation and consider that agents’ experiments give them access to some aspects of the world. How much the truth value of an argument influences that of its attackers and supporters is also controlled by the experimental accuracy σ .

Table 1: Model parameters and variables. The quantities indexed by a * change at each time step t .

Symbol	Description
N	Number of agents.
p_{ER}	Probability that two agents are linked in the social network.
\mathcal{N}	Social network of the agents.
T_i	Truth value of the issue of the debate.
σ	Accuracy of the agent’s experiments.
\mathcal{O}_k^*	Opinion graph of agent k .
$\mathfrak{J}_k(a)^*$	Score of argument a in agent k ’s opinion graph.
o_k^*	Opinion of agent k ; score of the the issue in agent k ’s opinion graph.
T_a	Truth value of argument a .
r_k^*	Result sampled by agent k .

3.3 Examples

Tables 1 presents an overview of every parameter and variable in our model, along with their description. Here, we present an example step of our protocol.

Example 1. Let us consider $N = 3$ agents, members of the social network $\mathcal{N} = (1, 2), (1, 3)$ shown below. In the beginning, all agents’ opinion graphs are composed only of the issue i . We fix $T_i = 0.8$, $\sigma = 0.5$.



During the first step, all agents $k \in \{1, 2, 3\}$ investigate the issue i , and they all have the same opinion $o_k^0 = 0.5$. They each sample a result r_k^0 from the same Gaussian distribution $\Gamma(0.8, 0.5)$. Table 2 details an example of the first step of such a run, where we indicate the sampled result, the name of the generated argument and its relation with the issue computed as stated above,

its truth value, the argument that the agents learn from their neighbors and their updated opinion graph and opinion. Note that all relevant quantities are indexed for the agents k and time step $t = 0$.

Table 2: Example first step of a debate with three agents.

k	Opinion o_k^0	Sampled Result r_k^0	Generated Argument	Truth value	Arguments Learned	Updated Opinion Graph	Updated Opinion o_k^1
1	0.5	0.78	a , support	1.00	b, c		0.62
2	0.5	0.62	b , support	0.65	a		0.73
3	0.5	0.45	c , attack	0.3	a		0.5

At the end of this step, the three agents each have a distinct opinion graph and opinion of the issue. Agents 1 and 2's opinions are closer to the truth value of the issue $T_i = 0.8$, while agent 3's opinion has not changed, although her graph has. Listening to more agents is not always beneficial for convergence to the truth, as we see in the case of agent 1.

Two questions arise naturally from the observation of our model :

- Can the agents approach the truth, and which conditions are favorable to their epistemic success?
- Can agents influence each other in a way that makes their opinions converge?

These questions are explored in the following section.

4 Results

In theory, the agents' opinions can keep evolving forever. In practice, they vary less and less as steps go by and can reach states of ‘stability’ where no significant variation happens for hundreds or thousands of steps. Based on our experimental results, we decided to stop all simulations after 500 steps, as in the majority of runs we observed, the opinion of agents have reached a state of stability by then ³.

³ Due to the lack of place, we do not develop this notion of stability here. In practice, we measure the rolling standard deviation of the opinion of one agent as a measure of how much it varies across the steps.

4.1 Opinion Convergence

[22] show that agents' opinions converge when they exchange arguments, which hints that argumentation paired with gradual semantics could be a form of opinion-merging mechanism. Like [22], we study the convergence of agents' opinions by reporting their standard deviation, which we denote the **opinion diversity** metric.

Definition 2 (Opinion Diversity).


Let us consider a set of N agents and let o_k^t denote the opinion of agent k at time step t . The opinion diversity is the standard deviation of the opinion of agents:

$$D^t = \sqrt{\frac{1}{N} \sum_{k=1}^N (o_k^t - \bar{o}^t)^2}$$

where \bar{o}^t is the mean opinion of all agents at time step t :

$$\bar{o}^t = \frac{1}{N} \sum_{k=1}^N o_k^t$$

Formal Discussion Mathematically, exchanging arguments does not automatically lead to a convergence of the opinions of the agents. Consider the following example:



Let us consider two agents 1 and 2 whose opinion graphs, represented above, are respectively composed of two arguments and none. Their opinions are the same: $o_1 = o_2 = 0.5$. Let us consider what happens if agent 2 adds one of the arguments of agent 1 to her opinion graph: this is the situation denoted as agent 2' above. Her opinion $o_{2'} = 0.37$ is more distant to that of agent 1, even though their opinion graphs are now more similar⁴.

A similar property was shown by [22] for a different semantic and class of graphs. The question of whether the agents' opinions will converge experimentally is not a trivial one.

Simulation Results Figure 2a reports the opinion diversity of the community when we vary p_{ER} , the probability that two agents are connected in the social network. We observe a surprising threshold phenomenon: the opinion diversity *increases* with p_{ER} as long as $p_{ER} \leq 0.5$, then *decreases* when $p_{ER} \geq 0.5$ to reach 0 in the case of a complete network.

Observation 1 : Communities become more diverse when p_{ER} increases below 0.5 and less diverse when p_{ER} increases above 0.5.

⁴ The notion of similarity between graphs is not developed here, but consider for example a graph-edit distance [18].

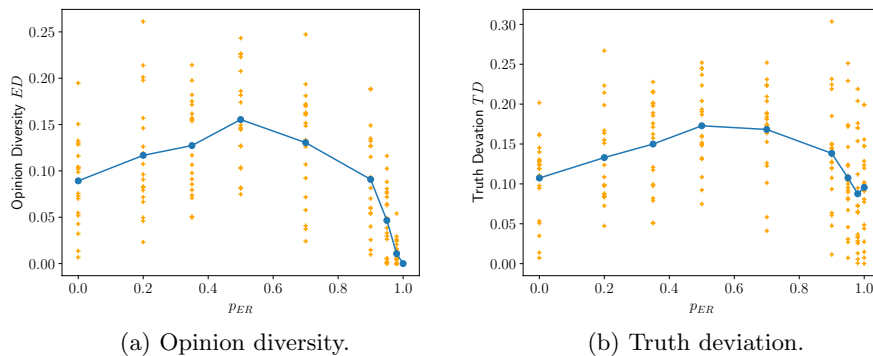


Fig. 2: $N = 10$, $\sigma = 0.15$, T_i is random, p_{ER} varies. The orange points are individual simulation results, blue points are averaged for each p_{ER} .

These results are surprising. When p_{ER} is low, agents interact with a limited number of others, but there is no obvious reason that this should lead them to have more diverse opinions as a whole: indeed, in the classical bounded confidence models for example, interaction between agents always leads to an averaging of their opinion and thus, to a decrease of global diversity. Our model's dynamics are more complex, as we have formally shown that communication may lead to divergence, and as the opinion of our agents evolve even when they do not interact with each other. Further analysis is warranted to understand these effects.

4.2 Epistemic Success

The epistemic success of agents refers to the degree to which they are able to arrive at accurate opinions. In our case, this accuracy is measured as the distance between an agent's opinion and the truth value T_i . We use the **truth deviation** metric introduced by [10] to measure the epistemic success of our community.

Definition 3 (Truth Deviation).

Let us consider a set of N agents and let o_k^t denote the opinion of agent k at time step t . The truth deviation is:

$$TD^t = \sqrt{\frac{1}{N} \sum_{k=1}^N (o_k^t - T_i)^2}$$

with T_i the truth value of the issue of the debate.

Experimental Accuracy We find that σ is correlated with the truth deviation: the smaller it is, the better agents are at approximating the truth (see Figure 3). This result is not unexpected given the role of σ as the experiment accuracy of the agents. It validates our model as one where the success of truth-seeking agents can be controlled by an accuracy parameter.




Fig. 3: Truth deviation. $N = 1$, σ varies, T_i is random. The orange points are individual simulation results, blue points are averaged for each p_{ER} .





Fig. 4: Truth deviation against opinion diversity. $N = 10$, $\sigma = 0.15$, T_i is random, p_{ER} varies.

Probability of connection When we vary p_{ER} , the probability that two agents are connected in a community of $N = 10$ agents, we see in Figure 2b that the truth deviation *increases* for $p_{ER} \leq 0.5$ then *decreases* for $p_{ER} \geq 0.5$. We observe the same threshold of $p_{ER} = 0.5$ as in the opinion diversity metric. This hints of a relation between the diversity of the opinions and the truth deviation. To explore this link, we plot in Figure 4 the opinion diversity of each of the debates with varying p_{ER} and the truth deviation of the community, and we see clearly that higher diversity is associated with a worse epistemic success.

Observation 2: Diverse communities are less epistemically successful.

Remarkably, this result presents similarities with the unintuitive Zollman paradox which states that more communication is sometimes detrimental to the epistemic success of communities [23]. In Zollman’s model, however, less communication is correlated with more diversity, which gives communities a boost in epistemic success. In our model, the relationship between connectivity and diversity is more complex, and diversity worsens the outcome of agents. Zollman’s model is composed of Bayesian agents, who exchange directly the result of their investigations, whereas our agents transform their investigations into arguments. Given how we modelled argument generation, an argument improves the opinion of the agent who generated it, because it corrects her current opinion. However, if it is shared with another agent who holds a different opinion, it can mislead her: this could explain the harmful effect of diversity.



Consider the following situation: the truth value of the issue is $T_i = 0.8$, two agents are equipped with opinion graphs (without the purple argument), thus have opinions $O_1 = 0.88$ and $O_2 = 0.67$. Now, let's consider what happens if agent 1 investigates the issue, and collects a sample result $r_1 = 0.8$. Because $O_1 > r_1$, this sample result is transformed into an attacking argument by agent 1. After generating and adding this new argument (in purple), the opinion of agent 1 is now 0.81, closer to the truth value of the issue. However, if agent 1 shares this new argument with agent 2, then the opinion of agent 2 decreases as well, to 0.55: it is now further away from the truth value of the issue. The difference between their opinions made the exchange of arguments counter-productive.

5 Discussion and Conclusion

We present a novel modelling of an epistemic community using abstract argumentation as the representation of information. The analysis of our model presents us with two stylized facts.

1. Communities become more diverse when p_{ER} increases below 0.5 and less diverse when p_{ER} increases above 0.5.
2. Diverse communities are less epistemically successful.

Observation 2 is of particular interest. One limit of classical opinion diffusion models is that they cannot account for phenomenons of polarisation as an increase in the distance between people's opinions which have been reported by real-life experiments [4]. Observation 2 suggests that our model may account both for an increase in the diversity of opinions when agents interact more, as well as the formation of a consensus in tight-knit communities.

Observation 3 enriches the discussion about the counter-intuitive effects of communication in epistemic communities, opened by Zollman's paradox [23]. Our results hint at a more complex link between the rate of communication between agents and the diversity of their opinions, and our observation is in direct contrast with agent-based models from the literature showing that opinion diversity benefits truth-seeking communities [11, 8]. One direction for our future work will be to explore how our model can inform the role of the mode and quality of communication on the effect of diversity in epistemic communities, as explored in [19].

References

1. Amgoud, L., Ben-Naim, J., Doder, D., Vesic, S.: Acceptability semantics for weighted argumentation frameworks. In: Proceedings of the Twenty-Sixth Int. Joint Conference on Artificial Intelligence, IJCAI-17. pp. 56–62 (2017). <https://doi.org/10.24963/ijcai.2017/9>
2. Banisch, S., Olbrich, E., et al.: An argument communication model of polarization and ideological alignment. *Journal of Artificial Societies and Social Simulation* **24**(1), 1–1 (2021)

3. Borg, A., Frey, D., Šešelja, D., Straßer, C.: An argumentative agent-based model of scientific inquiry. In: IEA/AIE 2017, Arras, France, June 27-30, 2017, Proceedings, Part I 30. pp. 507–510. Springer (2017)
4. Bramson, A., Grim, P., Singer, D.J., Berger, W.J., Sack, G., Fisher, S., Flocken, C., Holman, B.: Understanding polarization: Meanings, measures, and model evaluation. *Philosophy of science* **84**(1), 115–159 (2017)
5. Cayrol, C., Lagasque-Schiex, M.C.: On the acceptability of arguments in bipolar argumentation frameworks. In: ECSQARU 2005, Barcelona, Spain, July 6-8, 2005. Proceedings 8. pp. 378–389. Springer (2005)
6. Deffuant, G., Amblard, F., Weisbuch, G., Faure, T.: How can extremism prevail? a study based on the relative agreement interaction model. *Journal of artificial societies and social simulation* **5**(4) (2002)
7. Erdős, P., Rényi, A., et al.: On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci* **5**(1), 17–60 (1960)
8. Gabriel, N., O’Connor, C.: Can confirmation bias improve group learning? (2022)
9. Goldman, A.I.: *Knowledge in a Social World*. Oxford University Press (1999)
10. Hegselmann, R., Krause, U., et al.: Truth and cognitive division of labor: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation* **9**(3), 10 (2006)
11. Hong, L., Page, S.E.: Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* **101**(46), 16385–16389 (2004)
12. Mäs, M., Flache, A.: Differentiation without distancing. explaining bi-polarization of opinions without negative influence. *PloS one* **8**(11), e74516 (2013)
13. Mercier, H., Heintz, C.: Scientists’ argumentative reasoning. *Topoi* **33**(2), 513–524 (2014)
14. Mercier, H., Sperber, D.: Why do humans reason? arguments for an argumentative theory. *Behavioral and brain sciences* **34**(2), 57–74 (2011)
15. O’Connor, C., Weatherall, J.O.: Scientific polarization. *European Journal for Philosophy of Science* **8**(3), 855–875 (2018)
16. Rainer, H., Krause, U.: Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation* **5**(3) (2002)
17. Rosenstock, S., Bruner, J., O’Connor, C.: In epistemic networks, is less really more? *Philosophy of Science* **84**(2), 234–252 (2017)
18. Sanfeliu, A., Fu, K.S.: A distance measure between attributed relational graphs for pattern recognition. *IEEE transactions on systems, man, and cybernetics* (3), 353–362 (1983)
19. Sulik, J., Bahrami, B., Deroy, O.: The diversity gap: when diversity matters for knowledge. *Perspectives on Psychological Science* **17**(3), 752–767 (2022)
20. Taillandier, P., Salliou, N., Thomopoulos, R.: Introducing the argumentation framework within agent-based models to better simulate agents’ cognition in opinion dynamics: application to vegetarian diet diffusion. *Journal of Artificial Societies and Social Simulation* **24**(2), 6 (2021)
21. Dupuis de Tarlé, L.: A gradual semantic to model opinion from bipolar argumentation graphs. In: Proceedings of the 16èmes Journées d’Intelligence Artificielle Fondamentale (JIAF 2023) (2023)
22. Dupuis de Tarlé, L., Bonzon, E., Maudet, N.: Multiagent dynamics of gradual argumentation semantics. Proceedings of the 21th Int. Conference on Autonomous Agents and Multi Agent Systems (AAMAS) (2022)
23. Zollman, K.J.: The communication structure of epistemic communities. *Philosophy of Science* **74**(5), 574–587 (2007)