DOI: 10.1017/psa.2025.11
This is a manuscript accepted for publication in *Philosophy of Science*.
This version may be subject to change during the production process.

How the structure of scientific communities could impact the public uptake of uncertain science

Sacha Ferrari[1], Wouter Lammers[2] and Sylvia Wenmackers[1*]

[1]Center for Logic and Philosophy of Science, KU Leuven, Belgium, [2]Public Governance Institute, KU Leuven, Belgium
[*]Corresponding author: Sylvia Wenmackers; Email: sylvia.wenmackers@kuleuven.be

Abstract
We present an agent-based model to study how the network structure of a scientific community could impact the public uptake of science, and how this impact is influenced by scientific uncertainty and affinity bias. For unbiased agents, a highly connected scientific network decreases the probability that the public favors the correct theory. For biased agents, however, a moderately connected scientific network causes the public to favor the correct theory more often. This results from the competition between the scarcity of information (for poorly connected agents) and the spread of misleading information (for highly connected agents). Adding more scientists strengthens both effects.

1. Introduction

In contemporary society, science plays an important role in many aspects of life, such as healthcare, energy, and education. However, it can be challenging for individuals to determine the most credible scientific theory when making personal or policy decisions. Factors such as literacy level, ideological orientation, and the manner of science communication can influence their judgments (Miller, 1998; Rekker, 2021; Knight, 2006; Harker, 2015).

This article focuses on topics lacking scientific consensus, a common stage in the scientific process (Shwed and Bearman, 2010). Even perfectly rational scientists may endorse differing theories due to inherent variability in research findings. Hence, consensus is more likely when all research results are shared, but the speed of sharing and processing such information by peers has limits. Internal communication channels are thus vital for scientific progress. Moreover, scientists—like all humans—are susceptible to affinity bias, where information uptake is influenced by the source's affinity. This bias can affect consensus formation and sometimes even increase polarization.

Previous agent-based studies have shown that the structure of scientific networks affects scientists' beliefs, influencing the formation of consensus or polarization(Zollman, 2007; O'Connor and Weatherall, 2018). Empirical research also indicates that citizens react differently to scientific results when they perceive a lack of consensus among scientists. In particular, a lack of perceived consensus among scientists has been shown to have a slightly negative effect on citizens' belief in findings reported in science communication (Chinn et al., 2018; Gustafson and Rice, 2019; van Stekelenburg et al., 2022). So, there is a complex interplay of individual and network-level factors in the formation of scientific consensus and the effects on citizens' beliefs. So far, we know of little research simulating the effects of this interplay on citizens' uptake of scientific findings.

The core innovation of this article is that we investigate how network features of the scientific community affect citizens' uptake of scientific findings. We do so with computational simulations, extending the model of Zollman (2007). Our extended model includes two groups of actors: scientists and other citizens. We study the effect on citizens' belief of the different types of networks that scientists may form. We also include two additional variables: the uncertainty of the evidence and the affinity bias of scientists and citizens. In the next subsection, we discuss what we know about the four main variables.

1.1. Four main variables

First, the main dependent variable is the citizens' uptake of scientific theories: it is through this success rate among citizens that we can assess if the public gets a good understanding of science. In this article, we quantify the public uptake of science by with a single number: the success rate of a correct theory in the citizen community. This success rate is given by the proportion of the number of citizens favoring the correct theory over the total number of citizens (see below). For the sake of simplicity, we adopt what one calls in public communication of science and technology (PCST) the *deficit model* of science communication, which focuses on unilateral knowledge transfer from scientists to other citizens (Wynne, 1991; Burns et al., 2003). In this model, citizens are relatively passive receivers of evidence. We are aware of the limitations of this model (Trench, 2008; Seethaler et al., 2019), but we consider this minimal model here as a first step toward a more comprehensive understanding of the impact of scientific uncertainty on the public uptake of science (Schmid-Petri and Bürger, 2020). One limitation of our model is that while the scientists search for evidence is influenced by their prior beliefs (as explained below), the citizens are modeled as receiving the same evidence, to which they may respond differently depending on their prior beliefs.

Second, the main independent variable of our model is the structure of scientific networks. Here, we understand the term 'structure' as the shape of the network of epistemic relations that exist between scientists. In particular, two scientists share an epistemic connection in the network when they exchange their empirical results. Bibliometric analysis has shown that many scientists are just a few links away from each other (Newman, 2001*a*). Authors' positions in networks affect the uptake of their results (Uddin et al., 2013; Kumar, 2015). Thus, network structures directly affect the dissemination of newly produced scientific knowledge among scientists and potentially among citizens as well. Next, we consider two moderating variables.

Third, the acceptance of a scientific theory by citizens can depend on how uncertain this theory is. Uncertainty is inherent to scientific inquiry (Kampourakis and McCain, 2019; Pellizzoni, 2003) and can be due to the limited accuracy of the experimental setup (e.g., a PCR test with aleatory false positive results or a telescope with a low-resolution lens), the nature of the studied object itself (e.g., a complex social phenomenon or a

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

stochastic quantum effect), or both. The communication of scientific uncertainty to a public audience has received ample attention (Giles, 2002; Fischhoff and Davis, 2014; Broomell and Kane, 2017; Van Der Bles et al., 2019). Indeed, making explicit the scientific uncertainties can impact the acceptance of a scientific hypothesis or theory by citizens (Gustafson and Rice, 2019). To contribute to the existing literature, we aim to assess this impact in a more systematic and quantitative way.

Fourth, scientists and citizens alike are susceptible to psychological biases. One such bias is affinity bias, where individuals give more weight to evidence coming from people with whom they share similar beliefs, regardless of whether the new evidence confirms their own beliefs. So, affinity bias is a form of 'homophily', understood here as a preference for interacting with like-minded people (see Dandekar et al., 2013); it pertains to the source rather than the content. As such, it differs from biased assimilation or confirmation bias (whereby people selectively accept evidence that confirms their prior beliefs while rejecting disconfirming findings; see, e.g., Lord et al., 1979).[1] Affinity bias seems especially relevant for modeling scientists who revise their beliefs in response to evidence and who make decisions on whether or not further experiments are required. Moreover, the bias of individual scientists may impact the whole scientific community through peer interaction, as well as the rest of society through public communication. The impact of biases has been studied in scientific communities, both in psychology and in the philosophy of science (Peters, 2021; Mahoney, 1977; Wilholt, 2009; Schumm, 2021; Peters, 2022; Kelly, 2008; Dorst, 2023).

Biases have also been implemented in numerical models. For instance, Baumgaertner and Justwan (2022) modeled how people's beliefs are influenced by homophily. As mentioned, this bias is similar to what we call 'affinity bias' in the current article. However, Baumgaertner and Justwan (2022) only considered a single group of agents (modeled after online groups) with full beliefs, whereas we investigate two groups of agents with graded degrees of belief. An earlier example of a computational study focused on homophily is Dandekar et al. (2013), who started from DeGroot's (1974) model. Individuals update their subjective probability assignments by taking a weighted average over the opinions of others. This can be understood as an agent-based model on a total graph with weighted edges that can be chosen to represent homophily. Dandekar et al. (2013) pointed out that homophily alone does not lead to polarization in such a model (while biased assimilation does).

Our work aims to contribute to this debate by evaluating the role of affinity bias in shaping beliefs of scientists and citizens, especially under scientific uncertainty. Key questions include: How does affinity bias influence scientists' beliefs when results are uncertain? Is affinity bias overcome with more certain evidence? Additionally, our model tests whether individually problematic dispositions (e.g., affinity bias) are equally problematic at the group level. Some cognitive biases can be problematic at the individual level but turn out to be beneficial on group level (Peters, 2021); this is known as 'Mandevillian Intelligence' (Smart, 2018).

Methodologically, we chose the public uptake of science as our dependent variable, since this is the effect on which PCST focuses generally. The structure of the scientific community, the scientific uncertainty, and the affinity bias could in principle all be

[1] The simultaneous effect of social informational sharing and confirmation bias on polarization versus consensus has been studied by Del Vicario et al. (2017).

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

considered as independent variables. We selected the structure of the scientific community as our main independent variable, though, because our goal is to understand, for a specific structure of the scientific interactions, how changes in individual behaviors (i.e., affinity bias) and the accuracy of experiments (i.e., scientific uncertainty) affect the dependent variable.

1.2. Interaction of the four variables

Previous models have studied these variables in isolation or have focused on the interaction of some of them. In practice, however, these factors operate simultaneously and likely interact in complex ways, so their net effect seems impossible to determine a priori. We are not aware of any model or theory that has incorporated all these variables together. Therefore, we opted for a comprehensive simulation model to study dynamic interactions between these variables. This approach helps us to develop a more nuanced understanding of the effects of these factors on public science communication.

Complex interplays of parameters on individual and network levels can be simulated with agent-based modeling (Hedström and Ylikoski, 2010; Bruch and Atwell, 2015). Agent-based modeling is used in the social sciences to understand the dynamics of social phenomena (Šešelja, 2023). Such models consider social entities (individuals, institutions, etc.) like agents forming a network. Each of these agents can share information or influence others in other ways through the agent community. In particular, the network epistemology framework of Bala and Goyal (1998) has been adapted in the context of science by Zollman (2007), and has been further developed in several publications to describe the dynamics of scientific communities (Weatherall et al., 2020; O'Connor and Weatherall, 2018). Our article aims to adapt this model in a new direction, to simulate how a scientific community exchanges knowledge with a non-scientific audience.

Our model represents an undecided scientific community hesitating between two theories, *A* and *B*. We assume that one theory is, in fact, correct, but the scientists only have fallible means for determining this empirically. Some scientists perform experiments; they make their outcomes public to inform other scientists as well as citizens. In response, the members of both groups progressively change their degrees of belief concerning theory *A* and *B*. As far as we know, our extension of Zollman's model is the first one to consider two distinct epistemic communities: scientists and other citizens.

As we will see in the next section, our four variables can be implemented numerically, so their influence can be quantified. The outcomes of our model can be read both descriptively and normatively. On the one hand, we *describe* how agents react to various combinations of the aforementioned variables and parameters. On the other hand, we can use this knowledge to assert how a scientific network *should* be organized in order to maximize public uptake of the correct scientific theory. Our article can also be considered as a first step toward a comprehensive computational study of the deficit model in science communication. Our methodology can be extended to more complex science communication paradigms, such as the dialogue approach in PCST, but doing so falls outside the scope of our present work.

Our paper is structured as follows. In section 2, we introduce Zollman's model and present our modifications to it. In section 3 we run the model with varying input values for the main parameters. In section 4, we summarize our key findings and suggest directions for follow-up studies.

2. The model

In this section, we introduce the model of Zollman (2007) and our extension of it. We explain how we implemented the adapted model numerically to address our research question.

2.1. Zollman's model and our application of it

There are several agent-based models that aim to describe how individuals create, share, and update their knowledge (also known as opinion dynamics; for a review, see e.g., Fischbach et al., 2021). In recent years, Zollman's model and its improvements raised specific attention. In his influential article, Zollman (2007) applied the economic model of Bala and Goyal (1998) to epistemic communities in order to understand their communication structures. Such a model only considers one type of agents, which represent scientists. Each scientist is a node of a communication network: a scientist can interact with other scientists (if some communication channel links them directly) or can stay isolated from other scientists (if no direct channel exists between them).

Zollman (2007) described quantitatively how this network of interactions influences each scientist in their beliefs in given theories. Agents can have degrees of beliefs about which of two options, A and B, is best. From round 1 onwards, agents have to decide between two statements: 'Treatment A is better than treatment B' or 'Treatment B is better than treatment A'. In our article, we apply the model to two scientific theories (rather than treatments, although this interpretation remains admissible, too). We consider a scientific community in which two competing theories, A and B, have been proposed to explain a given phenomenon. The first theory, A, is a well-known theory that has been confirmed by a large number of experiments. The second theory, B, is either a theory that has so far been ignored—for instance, because its predicted effects were too small compared to available measurement resolution—or an improved version of theory A. We consider the phase in which a new empirical method has just become available that might give more strength to theory B (relative to theory A).[2] Note that, in reality, two scientific theories are rarely each other's negation.[3] Our model merely compares the relative merits of two theories, in a context where those are the main or only contenders.

To model such situations, in which dissensus about two rival theories has started to emerge within the scientific community, we assume that each scientist has a personal degree of belief in which theory is better. We assume that these degrees of belief are rational in the sense that they adhere to the axioms of probability. Hence,

[2] As a historical example, theory A may represent geocentrism and theory B may be heliocentrism. The latter theory had been suggested in Antiquity, but had been ignored because there were no measurable effects. Early telescopic observations in the seventeenth century showed evidence of moons that revolve around planets other than the Earth, which provided direct empirical confirmation of heliocentrism relative to geocentrism.

[3] In the previous example, a hybrid theory was indeed proposed: geo-heliocentrism (see, e.g., Blair, 1990).

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

we also call the degrees 'credences'. For a given agent at a given time, we denote these degrees of belief respectively by $P(A) \equiv P($'Theory A is better than theory B'$)$ and $P(B) \equiv P($'Theory B is better than theory A'$)$. They are both real numbers between 0 and 1, with $P(A) = 0$ denoting the agent's subjective certainty that the theory B is better than theory A and $P(A) = 1$ denoting their certainty that theory A is better than B; mutatis mutandis for $P(B)$. Rational coherence requires that these credences obey the normality requirement: $P(A) + P(B) = 1$. So, for instance, if a scientist has a credence of 80% that theory B is better than A ($P(B) = 0.8$), their credence that theory A is better must be 20% ($P(A) = 1 - 0.8 = 0.2$). If an agent has a credence above 50% that either theory is better than the other (at a given time), we say the agent *favors* that theory.

As mentioned earlier, in our model, theory A is initially far more established than theory B. Scientists are unlikely to challenge theory A without significant belief in theory B. However, some *dissident* scientists may doubt the established theory A and conduct new experiments to confirm their belief. Meanwhile, their *conservative* colleagues strongly favor theory A and will not perform additional experiments. Stated differently, we assume that only dissident scientists who have a prior degree of belief in the superiority of theory B greater than 50% ($P(B) > 0.5$, or equivalently $P(A) < 0.5$) will deem it relevant to run further experiments in order to further confirm theory B and to convince their colleagues that theory B warrants more support than theory A. Conservative scientists who favor the established theory A (i.e., having $P(A) > 0.5$ and thus $P(B) < 0.5$) lack incentive to run extra experiments due to A's established empirical adequacy and their low belief in B. However, a conservative scientist may update their beliefs based on dissidents' results, and if they come to favor B ($P(A) < 0.5$), they may run experiments to confirm their new belief, becoming dissidents themselves.

Zollman's model assumes that the second treatment is better than the first one. Analogously, we assume that theory B has better predictive success than theory A. However, the experimental device is not perfect and leaves room for uncertainty. That is, the device does not lead to a positive result in favor of B 100% of the time. Although not perfect, we expect it to have an accuracy of more than 50%. A lower value would imply, given our assumption that B is indeed better than A, that the device is not a suitable one. A 50% accuracy would be equivalent to assessing the truth of theory B by flipping a coin. We define the accuracy of the experimental device, p, as the sensitivity of the device: the probability of producing a true positive experimental result (given that B is the correct theory).[4] For example, in a counterfactual case where the geocentrism versus heliocentrism debate took place with nineteenth-century telescope technology, p could represent the probability of measuring stellar parallax (which is a true positive, given that heliocentrism is correct). Formally, we use the notation

$$p = 0.5 + \varepsilon, \tag{1}$$

with ε a real number between 0 and 0.5. If $\varepsilon = 0.5$, then the device is 100% accurate and its results leave no room for uncertainty.

Since p is a probability, its (Bayesian) interpretation can be extended from the experimental accuracy to encompass other forms of uncertainty. Indeed, dispersion in the experimental outcomes is not necessarily caused by the measurement device alone, but

[4]Like Zollman (2007), we assume that the device never produces false negative results (100% specificity). Hence, we use the terms accuracy and sensitivity interchangeably.

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

can also result from the intrinsic stochasticity of the system under study itself. Having that in mind, one can apply our model to many other fields dealing with inherent uncertainty, such as the social sciences, medicine, statistical physics, and quantum mechanics. For instance, one can cite sampling error in a population survey (in that case, one of the theories could be 'The majority of the population is smaller than 1m70.'), chaos in weather simulation ('It will rain tomorrow at 2:34 pm.'), the detection of an electron outside an electron trap ('The electron stays at least 30 min in the trap.'), etc.

In Zollman's model, in order to reduce statistical error, each dissident scientist chooses to run the experiment n times (always with a probability of success of $0.5 + \varepsilon$ for each run). E denotes the event of k positive results out of n runs. The probability of this event given that theory B is true is given by the binomial distribution:

$$P(E|B) = P(k, n, p) = \frac{n!}{(n-k)!k!}\, p^k (1-p)^{n-k}, \tag{2}$$

with $p = 0.5 + \varepsilon$.

When faced with the evidence E of such a run of n experiments, each agent updates their prior credences according to Bayes' rule:

$$P_{\text{new}}(B) = P(B|E) = \frac{P(E|B)P(B)}{P(E)}, \tag{3}$$

with $P(B)$ the agent's prior credence in the superiority of theory B, $P(E|B)$ the probability of the evidence E given that the theory B is true, and $P(E)$ the absolute probability of E. By the law of total probability, the latter probability can be rewritten as

$$P(E) = P(E|B)P(B) + P(E|A)P(A), \tag{4}$$

which says that the probability that E occurs (without knowing whether proposition A is true or B) is proportional to the probability that it occurs on theory A or on theory B, weighted by the probability that the given theory is correct. Assuming E corresponds with k successes out of n experiments, we obtain by combining eqs. 2–4:

$$\begin{aligned} P_{\text{new}}(B) &= \frac{p^k(1-p)^{n-k}P(B)}{p^k(1-p)^{n-k}P(B) + (1-p)^k p^{n-k}(1-P(B))} \\ &= \frac{1}{1 + \frac{1-P(B)}{P(B)}\left(\frac{0.5-\varepsilon}{0.5+\varepsilon}\right)^{2k-n}}. \end{aligned} \tag{5}$$

Each dissident scientist will perform an experimental run and update their prior belief according to eq. 5. If ε is very small (meaning that an unreliable experimental device is used), the outcome of the run has a non-negligible probability of disconfirming theory B. Then, after updating their belief, the dissident scientist can end up with a degree of belief in B below 0.5. The agent will then disfavor their former favorite theory B and become a conservative scientist, who favors theory A. This scientist will not perform any new experiments because, as mentioned, scientists are reluctant to perform a costly experiment in favor of a new theory in which they have low credence, while there are already a lot of old experiments in favor of A.

But all the scientists, both conservatives and dissidents, aim to improve their knowledge and are open to listening to neighbor scientists located in their direct network. Thus, even if conservative scientists will not perform an experiment themselves, they

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

will consider the experimental results of dissident colleagues who are direct neighbors and update their prior belief according to eq. 5. In Zollman's model, all pieces of evidence have the same weight, regardless of whether it comes from the scientist themselves or from dissident colleagues.

2.2. Extending the model

Some extensions of Zollman's original (2007) model have been proposed in the literature. O'Connor and Weatherall (2018) added a social bias (similar to affinity bias), related to the source of the evidence: in their model, scientists treat the evidence of peers as more uncertain when their credences are further apart. The authors found that this promotes polarization, but their model only concerns the scientific community and does not include citizens. Wu (2023) set up a variant of the model including two groups of agents, in which members of one group ignored the testimony of members of the other group. Zollman's later (2010) model represents scientists who have a choice again between two methods, but now, instead of one having a known success rate and the other unknown compared to it, both methods have unknown success rates (modeled by binomial distributions). Gabriel and O'Connor (2024) added confirmation bias to this model and found that it may improve group learning. After each experimental round, agents have some probability to accept or reject these outcomes. This probability is driven by a beta-binomial distribution which depends on the history of success and failure of each theory and the new outcomes. In another version of the model, Weatherall et al. (2020) considered an epistemic community made of scientists, policy makers, and a propagandist. The propagandist aims to shift public opinion in one direction by cherry-picking among the experimental results confirming their prior beliefs and massively sharing them. Even though both communities (i.e., scientists and citizens) are considered, Weatherall et al. (2020) did not give a systematic study of the impact of the scientific network. They only considered two types of networks among the thousands possible: the cycle graph and the complete graph. In the next section, we will discuss the interpretation of these graphs in more detail.

In our model, affinity bias influences how agents (both scientists and citizens) adapt their degrees of belief in response to the testimony of (other) scientists. Specifically, if the agent is prone to affinity bias, their trust in the scientist's testimony will be high if their prior credence on a particular topic (in this case, whether they favour theory B) is very similar. The closer the agent's prior degree of belief is to that of the scientist, the more the agent will trust the reported evidence.

To represent this type of belief revision, we must deviate from Bayes' rule (eq. 5) that was part of Zollman's base model, because it assumes that all evidence is learned with certainty. Instead, we start from Jeffrey's (1990) generalization of Bayes' rule, as did O'Connor and Weatherall (2018). In addition, we modify the way agents respond to testimony under the influence of affinity bias. To achieve this, we essentially use the same equation as O'Connor and Weatherall (2018), but with one component fewer.

Formally, we consider a scientist j who report their evidence E to another agent i, who does not fully believe this testimony. By 'testimony', we consider an observation report (i.e., a scientist's testimony on their experimental evidence), rather than an expert's posterior degree of belief, which has been studied, e.g., by Steele (2012) and Roussos (2021). According to Jeffrey's (1990) conditioning, the posterior of agent i is as follows:

$$P_i'(B) = P_i(B|E)P_i'(E) + P_i(B|-E)P_i'(-E), \tag{6}$$

where $P_i'(B)$ is agent i's posterior credence that theory B is better than theory A, $P_i(B|E)$ and $P_i(B|-E)$ are the conditional probabilities of theory B being better than theory A given that E or $-E$ occurred, respectively (see eq. 5), and $P_i'(E)$ and $P_i'(-E)$ are agent i's posterior credence that E or $-E$ occurred, respectively, after accounting for the testimony of scientist j. These final two factors are influenced by the affinity bias, as defined in eq. 7 below.

In our model, when scientist j claims that they received evidence E, the posterior credence of agent i depends on the affinity bias, as follows:

$$P_i'(E) = 1 - \min(1, \max(0, \alpha\,|P_i(B) - P_j(B)|)), \tag{7}$$

where $P_i'(E)$ is agent i's posterior credence in E, $|P_i(B) - P_j(B)|$ is the distance between the prior credences of agent i and scientist j in theory B being better than A, and α is a positive real parameter that represents *the degree of affinity bias* of agent i. $P_i'(-E)$ is obtained as $1 - P_i'(E)$.

If α is 0, the agent is not prone to affinity bias and $P_i'(E)$ will be 1, so the agent will trust another scientist regardless of how different their beliefs are. As α increases, the agent is more prone to affinity bias, so the agent will distrust experimental results coming from other scientists, except those for which $|P_i(B) - P_j(B)|$ is smaller than $1/\alpha$. We notice as well that when the credences of scientists i and j in the theories get closer, the subjective probability $P_i'(E)$ approaches 1: scientist i approaches full belief in the occurrence of E as reported by scientist j. So, there are two ways in which an agent i may fully trust the testimony of scientist j: if α is 0 or if agent i happens to have the same prior credence in B as scientist j. In both cases, $P_i'(E) = 1$ and Jeffrey's formula reduces to Bayes' rule.

Our eq. 7 is structurally similar to the expression introduced by O'Connor and Weatherall (2018), except that we left out the additional factor of $(1–P_i(E))$—a useful simplifying assumption.[5] So, our approach assumes that the uptake of the testimony only depends on the difference of degree of belief between i and j and the intensity of the affinity bias, regardless of how probable this piece of evidence is.

We have described how each scientist updates their degree of believe according their own experiment's outcomes and those of their epistemic neighbors. These pieces of evidence are communicated to the citizens via a communication channel, called a *mediator*. In our article, we only consider a rapporteur in the role of a mediator, who publishes all the scientific outcomes.[6] Unlike the dialogue model in PCST (Trench, 2008), there is no scientist–citizen knowledge co-production. The citizens merely receive information from the mediator, a one-way communication channel from scientists to citizens. Once new evidence has been produced by any scientist, it reaches *every* citizen. Like the scientists, each citizen will use eq. 6 to update their degree of belief. We assume that, realistically, citizens, like scientists, are prone to affinity bias.

[5] We also compared our results with those obtained by using O'Connor and Weatherall (2018)'s expression for affinity bias. The results differ only slightly quantitatively and the qualitative conclusions remain the same.

[6] In general, there are other options, such as a journalist who publishes the most interesting research results, a science educator, or an opinion maker (Burns et al., 2003).

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

2.3. From the theoretical model to its numerical implementation

As stated in the introduction, our aim is to understand how the structure of scientific communities, scientific uncertainty, and affinity bias impact the public uptake of science. The agent-based model we just reviewed gives us a useful tool to approach this question. We implemented our model in a Python algorithm (publicly available: Ferrari, 2025).

For each simulation, several parameters are fixed: the structure of the scientific community, the sensitivity of the experimental device, the affinity bias, the number N_{sc} of scientists in the scientific community, the number N_{cit} of citizens in the public, and the number of experiments n done by each dissident scientist in each run.

The N_{cit} citizens, like the N_{sc} scientists, all start with a prior degree of belief $P(B)$ at time $t = 0$. These degrees of belief (between 0 and 1) are randomly generated by the computer. Each dissident scientist (with $P(B) > 0.5$) will run n experiments with a probability of success of $0.5 + \varepsilon$ for each trial. At the next time increment ($t = 1$), each of these dissident scientists will update their personal degrees of belief according to the outcomes of their own experiment by using Bayes' rule. In addition, they will share their results with scientists located in their neighborhood. Remember that the network of scientists is a graph, where each scientist is represented by a vertex and each connection by an edge. Each of the scientists (conservative or dissident) of the neighborhood will update on each of the upcoming pieces of evidence coming from their neighborhood according to Jeffrey's rule (eq. 6). Then, each of the citizens will update their degree of believe with all the piece of evidence produced by the scientific community according to Jeffrey's rule as well.

We reiterate this process for $t = 2$, $t = 3$, etc. until all agents (both scientists and citizens) stabilize their beliefs: not changing them for subsequent time t. The time after which beliefs stabilize is called the *stabilization time* τ. Once all beliefs are stabilized, the simulation stops. This is the halt condition of our algorithm. We can now count how many scientists and citizens favor theories A and B. From these numbers, we can conclude what is the public uptake of science for this specific community.

So far, we described a single simulation for a specific values of the independent variables and a random degree of belief assignation. In order to have a general picture of the impact of a given choice of parameters (our independent variables), we would like to make this result independent of the prior beliefs of the agents (i.e., the $P(B)$ at $t = 0$). To do so, we randomized the initial distribution of beliefs of agents and simulate the same epistemic network with the same parameters a large number of times. More specifically, we start with a random distribution of scientists with degrees of belief between 0 and 1 and a distribution of citizens with degrees of belief between 0 and 0.5 (so, no citizens favor B at $t = 0$ because we assume that the conservative scientists had enough time in the past to convince all the citizens to favor theory A). Then, we average the proportion of agents favoring the correct theory (i.e., theory B) at the end of the interaction process (i.e., once all scientists' beliefs have stabilized). This average ratio is called the *success rate*. We use this success rate among the citizens to assess if the public gets a good understanding of science (assuming the deficit approach of PCST). This is why we quantify the dependent variable of this article (i.e., the public uptake of science) with the the success rate of theory B (i.e., the correct theory) in the citizen community.

We summarised the independent and dependent variables in Table 1. The main four variables of this article are written in bold.

Table 1. Independent and dependent variables of the model; main variables of interest indicated in bold.

Independent variable name	Symbol	Range of value
Number of scientists	N_{sc}	Natural number
Number of citizens*	N_{cit}	Natural number
Number of experiments at each run	n	Natural number
Network structure	None	All possible graphs representing connections between N_{sc} agents
Sensitivity of the experimental device	$0.5 + \varepsilon$	$\varepsilon \in [0, 0.5]$
Affinity bias of agents*	α	Positive real number
Dependent variable name	Symbol	Range of value
Success rate of scientists	None	$[0, 1]$
Success rate of citizens	None	$[0, 1]$
Stabilization time	τ	Natural number

⋆ = added to Zollman's model

Our model is now complete. In the next section, we will investigate how the choice of the scientific network affects both the scientists and the public in their beliefs: we study the success rate in these two communities.

3. Network structure

The structure of a scientific community can be represented by a graph in which each vertex represents an agent and each edge represents the connection existing between two agents. For instance, the graph of all members of the same university department is usually a complete graph: each agent stands in a direct epistemic connection with any other one. In other cases, the graphs might not be connected sets, such as when there is an accidental or forced segregation of two (or more) scientific communities (e.g., due to language barriers). In this case, the graph of the whole scientific community consists of at least two disconnected subgraphs. Such a setup does not imply that the subcommunities cannot reach the right conclusion independently. A more extreme case consists of a society of isolated agents with no communication between any of them. One can think of independent scholars non-affiliated to any university and lacking some covering for their research, or scholars during Antiquity when manuscripts were often unaffordable and communication means were very slow or nonexistent. Some authors also consider two other kinds of networks: the cycle and the wheel (Zollman, 2007; O'Connor and Weatherall, 2018). In a cycle network, each agent is connected to two other agents. The resulting connected graph is a loop. Such a network is one of the most economical one to link all agents together. However, the path between one agent to another can be long and has to transit through a lot of peers who could modify the message. The wheel is an improved version of the cycle with an agent at the cycle's center and connected to all other agents. This agent is like a postman providing shortcuts for communication between any pair of agents. An illustration of these four networks is presented in Figure 1.

In this section, our object of investigation is the effect of the network structure of the scientific community on the success rates within the communities of both scientists and citizens. As mentioned before, we assume here that the communication channel is a rapporteur, such that all experimental outcomes produced by the scientists are made

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

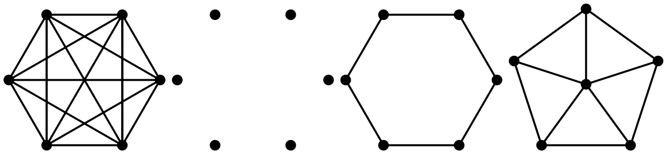

Figure 1. The complete, isolated, cycle, and wheel networks.

public to the citizens (which may be viewed as the ideal of open science) and that the latter take this information into account (an even less realistic modeling assumption). Stated differently, at each round, each citizen will update (in a Bayesian way) their degree of belief based on all the experiment outcomes produced during this round.

3.1. Complete, isolated, cycle, and wheel graphs

We begin by examining the four common network graphs—complete, isolated, cycle, and wheel—to understand their impact on the public uptake of science. We model a society of 20 scientists and 20 citizens, considering both societies of agents without affinity bias ($\alpha = 0$) and with agents prone to affinity bias ($\alpha > 0$). Initially, scientists' prior degrees of belief are randomly distributed between 0 and 1, and citizens' between 0 and 0.5.

3.1.1. Unbiased case: $\alpha = 0$

The result of the unbiased case is depicted in Figure 3.1.1. We notice that the wheel network is the most successful graph for making scientists favor the correct theory, followed by the complete network and the cycle network. For the isolated network even if half of the scientists start by favoring the correct theory on average, less than half of them end up with the right conclusion. We can explain this by noticing that a small value of ε implies a high probability of failure ($k < n/2$). In the case of an update with false negative results, a scientist starting with a prior degree of belief above 0.5 can have a posterior degree of belief below 0.5 at the next time increment. As a consequence, this scientist will then stop running experiments. Because the agent is isolated, they will not have any new experimental outcomes for updating their erroneous belief. Such an agent will stay stuck below 0.5 forever.

Concerning the impact of these four structures on citizens, we notice quite similar behaviors in each case. No citizen will be convinced to favor the correct theory if the accuracy of the experimental device is 0.5 ($\varepsilon = 0$). But the number of citizens that favors the correct theory grows rapidly and reaches the maximal value even for a poor accuracy of the device. It is surprising to see that the isolated network now performs as good as the other networks.

In order to study the robustness of our results, we varied the number of scientists and the number of citizens. We discovered that varying the number of citizens does not affect their success rate. However, as depicted in Figure 3.1.1, a larger scientific community leads to a better success rate both for itself and for the citizens. Although we exemplified it for a complete graph, this statement is valid for all the four graphs considered here.

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

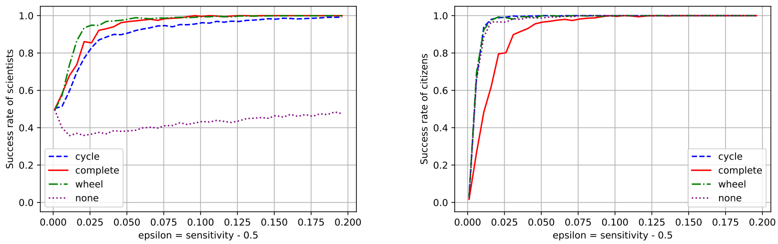

Figure 2. Fraction of scientists and citizens that reached the correct conclusion in a society without affinity bias in function of the increasing experimental accuracy (or sensitivity) $0.5+\varepsilon$ and the graph geometry. In these simulations: $N_{sc}=N_{cit}=20$, $\alpha=0$, $n=10$ and number of runs $=500$.

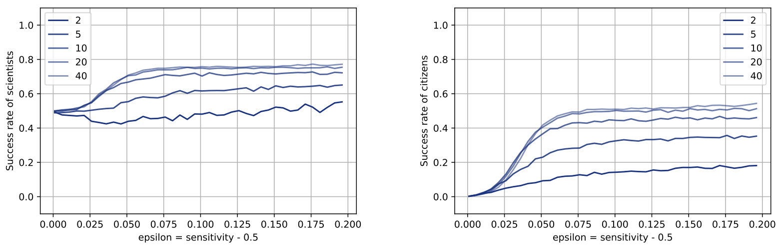

Figure 3. Fraction of scientists and citizens that reached the correct conclusion in a society without affinity bias in function of the increasing experimental accuracy (or sensitivity) $0.5+\varepsilon$ and the number of scientists in the case of a complete graph. In these simulations: $N_{cit}=20$, $\alpha=0$, $n=10$ and number of runs $=1000$.

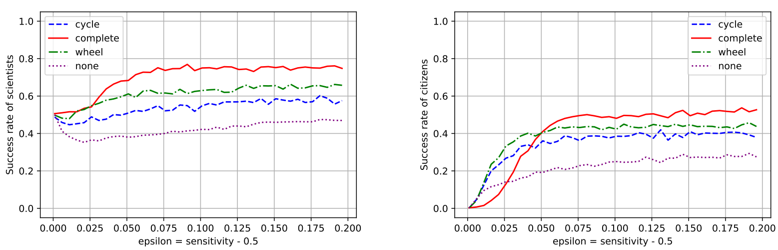

Figure 4. Fraction of scientists and citizens that reached the correct conclusion in a society with affinity bias in function of the experimental accuracy $0.5+\varepsilon$ and the graph geometry. In these simulations: $N_{sc}=N_{cit}=20$, $\alpha=2$, $n=10$ and number of runs $=200$.

3.1.2. Biased case: $\alpha>0$

We run the same algorithm, now considering a society with affinity bias ($\alpha=2$). In Figure 3.1.2, we notice that the success rate of the four networks is lowered and none of them achieve convincing either the scientists or the citizens to favor the correct theory. This effect is even more prominent in the case of citizens; only the complete graph reaches slightly more than 50%. If we add more scientists to the network, the success rate for the scientists rises but never surpasses 75%, and the citizens' success rate never rises above 50%. For readability, we omitted these plots.

The geometry with the lowest success rates is once again the isolated network. We notice here that the more connected a graph is, the higher its success rates are. The next subsection investigates whether that statement can be true in general.

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

3.2. General graphs

Even though the complete, isolated, cycle, and wheel graphs are often implemented in Zollman's model (Zollman, 2007; Weatherall et al., 2020; O'Connor and Weatherall, 2018), there are few systematic studies covering all the possible graphs. Zollman (2007) investigated all the possible graphs for $N_{sc} = 3, 4, 5, 6$. (We will discuss his conclusions later on.) Zollman's approach aims to be analytical: for a fixed number of scientists (i.e., vertices), he computed all the ways of linking them. He ended up with 2, 6, 21, and 112 possible graphs, respectively. This number grows exponentially with the number of scientists (Sloane, 2024): 853 for 7 scientists, 11,117 for 8, 261,080 for 9, 11,716,571 for 10, etc. Hence, an intrinsic limitation for this type of research is the exponential increase of computation time needed for exploring larger graphs. However, it is worth investigating beyond graphs of 6 vertices, since a scientific society is rarely limited to 6 individuals and important differences are to be expected for larger networks. Like in the previous example, we would like to simulate all possible graphs for a community of 20 scientists. For this case, there are roughly 10^{37} possible graphs (Sloane, 2024). Since our script takes 0.1 seconds for simulating 10 graphs in one CPU, it will take around 10^{35} seconds or 10^{27} years (i.e., more than one billion times the current age of the Universe) to go through all possible graphs. Clearly, this is far beyond the capacity of current computers. Instead, we modestly simulated 10,000 random graphs. We will show later that this tiny sample seems to suffice for studying the trend of the results.

Like Zollman (2007) and earlier authors (see, e.g., Newman, 2001*a*,b), we synthesize the graph identity with one unique number: the *clustering coefficient* (also called transitivity).[7] This coefficient aims to describe how vertices tend to be clustered. For each agent, the local clustering coefficient is proportional to the number of connections the agent's neighbors form. The more the neighbors are connected, the higher the local clustering coefficient. These local coefficients are computed for each vertex (i.e., for each agent) and are averaged. This final number (between 0 and 1) is called the global clustering coefficient or simply the *clustering coefficient* of the graph. For example, a completely isolated community has a graph with a null coefficient, and a fully connected community has a coefficient of 1. The cycle structure has a coefficient of 0.5. The higher the coefficient, the denser and more connected the community.

In Figure 3.2.2, we simulated 10,000 random graphs and computed their clustering coefficient, their stabilization time, and their ratio of scientists that reached the correct conclusion. The upper plots pertain to an unbiased society ($\alpha = 0$) and the lower ones to a more biased one ($\alpha = 2$). The complete, isolated, cycle, and wheel graphs are represented by specific symbols as well.

Our results for scientists agree with earlier work in this area. In addition, we consider the effect of the network in one community (the scientists) on the credences of another group (the citizens), for which no such studies exist. Moreover, we study the interaction with affinity bias, as discussed below.

[7] One could also describe the network with average path length: the average number of steps to connect two nodes by the shortest path. Real-world communities tend to have a small-world network: a high clustering coefficient and a low average path length.

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

3.2.1. Unbiased society

In the case of an unbiased society, the more clustered the graph is (i.e., the higher the clustering coefficient), the more likely the scientists will reach the correct conclusion. Concerning the citizens, however, it is the exact opposite: the more disconnected a graph is, the more likely the citizens will favor theory B. We notice, out of the four common graphs, that the complete one can not lead all the citizens to favor theory B. However, it is the quickest one: the community stabilizes after only a few iterations. The isolated graph lies in the bottom left and scores a success rate below 0.5 for the scientists although scoring at 1 for the citizens. We can explain this by noticing that once a dissident scientist runs an experiment whose outcomes lower their degree of belief below 0.5, they will never do an experiment again, nor update their belief based on another scientist's experiment. At the same time, each isolated scientist will share their knowledge with the audience and the latter will reach the correct conclusion. In general, we notice that increasing the clustering of a graph improves its stabilization time and the ratio of dissident scientists but lowers the chance of getting all the citizens unequivocally favoring theory B (i.e., being in one of the horizontal strips of the second graph). We can see it as a trade-off between a successful scientific community and a successful citizen community. The link between connectivity and stabilization time is consistent with the results of Zollman (2007).

3.2.2. Biased society

We ran the simulation again with a non-null level of affinity bias ($\alpha = 2$). We first notice that the three dotted clouds in the three lower charts are, on average, convex. This time, no graph achieves a success rate of 1 for the scientists, and, in a few graphs only, the success rate for the citizens is above 0.5 The most successful graphs are located around a clustering coefficient of 0.6. The success rate of the isolated graph is one of the worst ones even though its stabilization time is very low. The three other classical graphs have low success rates, especially the cycle and the wheel which lie below the majority of points. This specific convex shape of the curve can be understood as the result of two competing phenomena: epistemic isolation of the agents due to high affinity bias and the fast dissemination of false pieces of evidence in highly connected graphs. The first phenomenon that takes place is poorly connected graphs. Agents are isolated due to the lack of connection with other agents and have fewer opportunities to receive information from other agents. This effect is even more stringent with the affinity bias: even though an agent receives an experiment outcome from one of their rare peers, they are more easily prone to discard it. That explains the low success rate for poorly connected networks. This rate increases when the connectedness increases. However, a second phenomenon will counteract this increase. In a highly connected graph, information spreads very fast and very easily to all the agents. This may sound as beneficial to the success rate. However, even though true positive results (i.e., in favor of theory B) spread fast, false positive results (i.e., in favor of theory A) do as well. Such false positive results are difficult to correct once they have been communicated to a large number of agents. This effect has already been pointed out by Zollman (2007) and is known as the *Zollman effect* (Šešelja, 2023). This effect diminishes for poorly connected networks. The convex shape is thus understood as the result of these two competing phenomena.

We notice here a trade-off between accuracy and speed. On average, adding or removing some vertices to change the clustering coefficient of the graph in order to

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

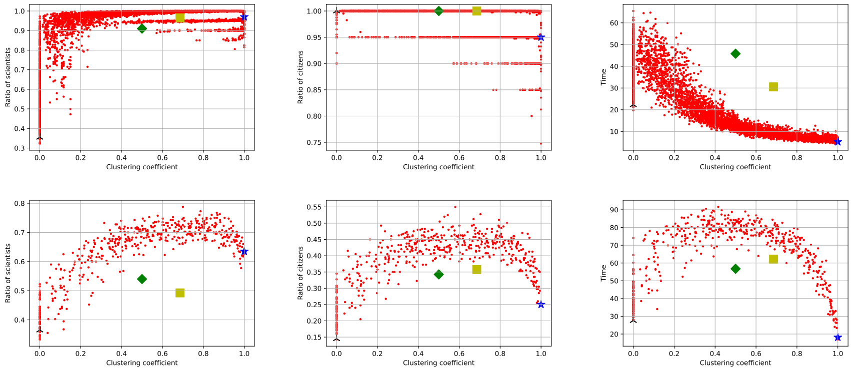

Figure 5. Fraction of scientists and citizens that reached the correct conclusion and the stabilization time in function of the clustering coefficient. $\alpha = 0$ for the two upper graphs and $\alpha = 2$ for the three lower ones. We fixed $\varepsilon = 0.05$, $n = 5$, number of generations $= 20$, and number of graphs $= 1000$. The blue stars denote the complete graph, the green diamonds the cycle, the yellow squares the wheel, and the black tripods the isolated graph.

reach the value of 0.6 (i.e., to maximize the success rate of scientists and citizens) will increase the stabilization time. Stated differently, slower graphs will perform better.

To assess the model's sensitivity to affinity bias, we also run the script for $\alpha = 4$. In this case, the curve of the two first graphs is shifted downwards: fewer scientists and citizens reach the correct conclusion. One could have expected this result: due to their strong affinity bias, all the agents will rarely update their degrees of belief and will stay stuck not far from their prior beliefs. The top of the curve lies around 0.5 on average for scientists and around 0.25 on average for citizens. The first value can be understood as follows. Because scientists who initially believe in theory B will never change their minds, their proportion stays the same throughout the interaction process (i.e., 50%). So half of the scientists of the initial and final community favors theory A whereas the other half favors theory B.

In this section, we studied the impact of the scientific network on both the scientists' and citizens' beliefs. We stressed that a society prone to affinity bias (i.e., a biased society) performs poorly and never achieves to make more than half the citizen population favor theory B. Even if these limitations are unavoidable, a poor result can be improved either by hiring more scientists (raising N_{sc}), or by reorganizing the scientific network in such a way that its clustering coefficient is near to 0.6 (i.e., moderately connected). In the case of unbiased societies, we saw that there is a trade-off between making either scientists or citizens favoring the correct theory. These results are especially interesting since they illustrate how the network of one community (i.e., the scientists) impacts the uptake by another (i.e., the citizens). This suggest that citizens' uptake is not only driven by the content of scientific information (i.e., the experimental outcomes) but also by the temporal variations of the flow of information. These variations are caused by the conversion of conservative scientists into dissident scientists and vice versa during all the simulation. In addition, the network's structure directly impacts the likelihood of such conversions.

https://doi.org/10.1017/psa.2025.11 Published online by Cambridge University Press

4. Conclusion and outlook

In this article, we investigated how the structure of the scientific community impacts citizens' uptake of science. We proposed an adapted version of the Zollman agent-based model including not only the structure of the scientific community and citizen uptake of scientific findings but also scientific uncertainty and the agents' propensity to affinity bias. The latter, as defined in eq. 7, is one of the major contributions of this article.

By doing an extensive study of the influence of the structure of the scientific network, we found that in unbiased societies, on average, most of the scientists and citizens arrive at believing the correct theory. We also noticed a trade-off between successfully making either scientists or citizens favoring theory B over theory A. Highly connected scientific communities will lead more scientists than citizens to believe in theory B. Less connected scientific communities will lead more citizens than scientists to believe in theory B. In contrast, we found that a society prone to affinity bias (i.e., biased society) performs poorly and never ends up with more than half of the citizen population favoring the true theory (i.e., theory B). Two interventions are possible if one wants to improve this ratio: (1) hiring more scientists or (2) reorganizing the scientific network in such a way that it is just moderately connected (clustering coefficient around 0.6). Our findings suggest that maximal connectivity is not always the best way to produce better science, which is in line with the findings of Zollman (2007).

The previous results give us more insight into how the choice of parameters influences the public uptake of science in the deficit model. By carefully adjusting these parameters, one can improve not only the success rate of the scientific community but also the public uptake of science. Some changes in the model are suggestive of interventions that can be tested experimentally and that can be influenced through policies for the organization of science and for science communication. For instance, one can change the number of connections per scientist in the model as well as in reality (e.g., by incentives that either promote or discourage team science). The effect of these choices will depend on other parameters as well (modifiable or not) such as the degree of affinity bias in society, the number of agents, or the experimental accuracy.

This article is a first contribution toward a comprehensive understanding of the interaction between scientists and the public in science communication. This model can also serve as a starting point for studying the limitations of the deficit model. For instance, we only considered a one-way interaction from the scientists to the other citizens and no interaction between the citizens. A possible improvement would be to move from a deficit model to a dialogue model, which allows two-way communication between the two types of agents as well as communication between the citizens. Other possible improvements would be to consider other psychological biases or to test our model with a more realistic network structure, for instance, based on citation patterns reported in empirical bibliometric studies. Different communication channels between the scientists and the citizen can also be implemented, as was done by Weatherall et al. (2020). Lastly, we assumed that once a dissident becomes conservative after running experiments favoring theory A, they will not perform any new experiment (by our definition of a conservative scientist). In the real world, however, one may expect that scientists do not give up so easily and keep experimenting for several iterations.[8]

We have used an expression for the posterior degree of belief in evidence reported by a scientist that depends on the agent's affinity bias and the difference between their prior credences on which theory to favour, but—unlike O'Connor and Weatherall (2018)—not on the prior probability of the evidence (eq. 7). Qualitatively, this simplification did not seem to affect our results, but we flag a systematic robustness study of different implementations of this bias as well as empirical validation as potential avenues for future research.

Our results suggest that the structure and size of the scientific community affect the uptake of correct theories by citizens but also that the direction of this effect depends on the degree of affinity bias. Without this bias, the probability that the public ends up favoring the correct theory decreases as the connectivity of the scientific network increases. When affinity bias is present, however, the probability that the public favors the correct theory is highest for a moderately connected scientific network. Both effects are more pronounced when the number of scientists increases.

CRediT author statement. Sacha Ferrari: Conceptualization, Investigation, Methodology, Software, Visualization, Writing – original draft. Wouter Lammers: Writing – review & editing. Sylvia Wenmackers: Funding acquisition, Project administration, Supervision, Writing – review & editing.

Funding information. This work was supported by the KU Leuven Internal Funds (grant C14/20/029).

Acknowledgements. We are grateful for helpful comments and suggestions from two anonymous reviewers. The authors thank Valérie Pattyn and Steven Van de Walle for their valuable feedback, and C. O'Connor and J. O. Weatherall for insightful discussions.

References

Bala, V. and Goyal, S. (1998) Learning from neighbours. *The Review of Economic Studies*. *65*(3), 595–621. https://doi.org/10.1111/1467-937X.00059.

Baumgaertner, B. and Justwan, F. (2022) The preference for belief, issue polarization, and echo chambers. *Synthese*. *200*(5), 412. https://doi.org/10.1007/s11229-022-03880-y.

Blair, A. (1990) Tycho Brahe's critique of Copernicus and the Copernican system. *Journal of the History of Ideas*. *51*(3), 355–377. https://doi.org/10.2307/2709620.

Broomell, S. B. and Kane, P. B. (2017) Public perception and communication of scientific uncertainty. *Journal of Experimental Psychology: General*. *146*(2), 286–304. https://doi.org/10.1037/xge0000260.

Bruch, E. and Atwell, J. (2015) Agent-based models in empirical social research. *Sociological Methods & Research*. *44*(2), 186–221. https://doi.org/10.1177/0049124113506405.

Burns, T. W., O'Connor, D. J. and Stocklmayer, S. M. (2003) Science communication: A contemporary definition. *Public Understanding of Science*. *12*(2), 183–202. https://doi.org/10.1177/09636625030122004.

Chinn, S., Lane, D. S. and Hart, P. S. (2018) In consensus we trust? Persuasive effects of scientific consensus communication. *Public Understanding of Science*. *27*(7), 807–823. https://doi.org/10.1177/0963662518791094.

Dandekar, P., Goel, A. and Lee, D. T. (2013) Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences*. *110*(15), 5791–5796. https://doi.org/10.1073/pnas.1217220110.

DeGroot, M. H. (1974) Reaching a consensus. *Journal of the American Statistical association*. *69*(345), 118–121. https://doi.org/10.1080/01621459.1974.10480137.

[8]Moreover, as indicated in the introduction, our current model does not aim to represent confirmation bias, which does impact how real-world agents deal with uncertain evidence and may be crucial to understand belief polarization (see, e.g., Kelly, 2008; Dorst, 2023).

Del Vicario, M., Scala, A., Caldarelli, G., Stanley, H. E. and Quattrociocchi, W. (2017) Modeling confirmation bias and polarization. *Scientific Reports*. *7*(1), 40391. https://doi.org/10.1038/srep40391.

Dorst, K. (2023) Rational polarization. *Philosophical Review*. *132*(3), 355–458. https://doi.org/10.1215/00318108-10469499.

Ferrari, S. (2025) The impact of scientific networks, affinity bias and scientific uncertainty on the public uptake of science (version 1.0.0). `https://www.comses.net/codebases/a0db6176-e1a4-4450-ac7e-a6cd78b9b235/releases/1.0.0/`. Accessed: 2025-01-28.

Fischbach, K., Marx, J. and Weitzel, T. (2021) Agent-based modeling in social sciences. *Journal of Business Economics*. *91*, 1263–1270. https://doi.org/10.1007/s11573-021-01070-9.

Fischhoff, B. and Davis, A. L. (2014) Communicating scientific uncertainty. *Proceedings of the National Academy of Sciences*. *111*(supplement_4), 13664–13671. https://doi.org/10.1073/pnas.1317504111.

Gabriel, N. and O'Connor, C. (2024) Can confirmation bias improve group learning? *Philosophy of Science*. *91*(2), 329–350. https://doi.org/10.1017/psa.2023.176.

Giles, J. (2002) Scientific uncertainty: When doubt is a sure thing. *Nature*. *418*(6897), 476–479. https://doi.org/10.1038/418476a.

Gustafson, A. and Rice, R. E. (2019) The effects of uncertainty frames in three science communication topics. *Science Communication*. *41*(6), 679–706. https://doi.org/10.1177/1075547019870811.

Harker, D. (2015) *Creating Scientific Controversies: Uncertainty and Bias in Science and Society*. Cambridge University Press.

Hedström, P. and Ylikoski, P. (2010) Causal mechanisms in the social sciences. *Annual Review of Sociology*. *36*, 49–67. https://doi.org/10.1146/annurev.soc.012809.102632.

Jeffrey, R. C. (1990) *The Logic of Decision*. University of Chicago Press.

Kampourakis, K. and McCain, K. (2019) *Uncertainty: How It Makes Science Advance*. Oxford University Press.

Kelly, T. (2008) Disagreement, dogmatism, and belief polarization. *The Journal of Philosophy*. *105*(10), 611–633. https://doi.org/10.5840/jphil20081051024.

Knight, D. (2006) *Public Understanding of Science: A History of Communicating Scientific Ideas*. vol. 26. Routledge.

Kumar, S. (2015) Co-authorship networks: A review of the literature. *Aslib Journal of Information Management*. *67*(1), 55–73. https://doi.org/10.1108/AJIM-09-2014-0116.

Lord, C. G., Ross, L. and Lepper, M. R. (1979) Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*. *37*(11), 2098–2109. https://doi.org/10.1037/0022-3514.37.11.2098.

Mahoney, M. J. (1977) Publication prejudices: An experimental study of confirmatory bias in the peer review system. *Cognitive Therapy and Research*. *1*, 161–175. https://doi.org/10.1007/bf01173636.

Miller, J. D. (1998) The measurement of civic scientific literacy. *Public Understanding of Science*. *7*(3), 203–223. https://doi.org/10.1088/0963-6625/7/3/001.

Newman, M. E. J. (2001a) Scientific collaboration networks. I. Network construction and fundamental results. *Physical Review E*. *64*(1), 016131. https://doi.org/10.1103/PhysRevE.64.016131.

Newman, M. E. J. (2001b) The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*. *98*(2), 404–409. https://doi.org/10.1073/pnas.98.2.404.

O'Connor, C. and Weatherall, J. O. (2018) Scientific polarization. *European Journal for Philosophy of Science*. *8*(3), 855–875. https://doi.org/10.1007/s13194-018-0213-9.

Pellizzoni, L. (2003) Knowledge, uncertainty and the transformation of the public sphere. *European Journal of Social Theory*. *6*(3), 327–355. https://doi.org/10.1177/13684310030063004.

Peters, U. (2021) Illegitimate values, confirmation bias, and Mandevillian cognition in science. *The British Journal for the Philosophy of Science*. *72*, 1061–1081. https://doi.org/10.1093/bjps/axy079.

Peters, U. (2022) What is the function of confirmation bias? *Erkenntnis*. *87*(3), 1351–1376. https://doi.org/10.1007/s10670-020-00252-1.

Rekker, R. (2021) The nature and origins of political polarization over science. *Public Understanding of Science*. *30*(4), 352–368. https://doi.org/10.1177/0963662521989193.

Roussos, J. (2021) Expert deference as a belief revision schema. *Synthese*. *199*(1-2), 3457–3484. https://doi.org/10.1007/s11229-020-02942-3.

Schmid-Petri, H. and Bürger, M. (2020) Modeling science communication: From linear to more complex models In *Science Communication*, Leßmöllmann, A., Dascal, M., and Gloning, T. (eds). vol. 17

of *Handbooks of Communication Science*. De Gruyter Mouton. Berlin, Germany. pp. 105–122. https://doi.org/10.1515/9783110255522-005.

Schumm, W. R. (2021) Confirmation bias and methodology in social science: An editorial. *Marriage & Family Review*. *57*(4), 285–293. https://doi.org/10.1080/01494929.2021.1872859.

Seethaler, S., Evans, J. H., Gere, C. and Rajagopalan, R. M. (2019) Science, values, and science communication: Competencies for pushing beyond the deficit model. *Science Communication*. *41*(3), 378–388. https://doi.org/10.1177/1075547019847484.

Šešelja, D. (2023) Agent-based modeling in the philosophy of science In *The Stanford Encyclopedia of Philosophy*, Zalta, E. N. and Nodelman, U. (eds). Metaphysics Research Lab, Stanford University. Winter 2023 edition.

Shwed, U. and Bearman, P. S. (2010) The temporal structure of scientific consensus formation. *American Sociological Review*. *75*(6), 817–840. https://doi.org/10.1177/0003122410388488.

Sloane, N. J. A. (2024) A001349: Number of connected graphs with n nodes. `https://oeis.org/A001349`. Accessed: 2024-06-18.

Smart, P. R. (2018) Mandevillian intelligence. *Synthese*. *195*, 4169–4200. https://doi.org/10.1007/s11229-017-1414-z.

Steele, K. (2012) Testimony as evidence: More problems for linear pooling. *Journal of Philosophical Logic*. *41*, 983–999. https://doi.org/10.1007/s10992-012-9227-5.

Trench, B. (2008) Towards an analytical framework of science communication models In *Communicating Science in Social Contexts*, Cheng, D., Claessens, M., Gascoigne, T., Metcalfe, J., Schiele, B., and Shi, S. (eds). Springer. Dordrecht, The Netherlands. pp. 119–135. https://doi.org/10.1007/978-1-4020-8598-7_7.

Uddin, S., Hossain, L. and Rasmussen, K. (2013) Network effects on scientific collaborations. *PloS One*. *8*(2), e57546. https://doi.org/10.1371/journal.pone.0057546.

Van Der Bles, A. M., Van Der Linden, S., Freeman, A. L. J., Mitchell, J., Galvao, A. B., Zaval, L. and Spiegelhalter, D. J. (2019) Communicating uncertainty about facts, numbers and science. *Royal Society Open Science*. *6*(5), 181870. https://doi.org/10.1098/rsos.181870.

van Stekelenburg, A., Schaap, G., Veling, H., van't Riet, J. and Buijzen, M. (2022) Scientific-consensus communication about contested science: A preregistered meta-analysis. *Psychological Science*. *33*(12), 1989–2008. https://doi.org/10.1177/09567976221083219.

Weatherall, J. O., O'Connor, C. and Bruner, J. P. (2020) How to beat science and influence people: Policymakers and propaganda in epistemic networks. *The British Journal for the Philosophy of Science*. *71*, 1157–1186. https://doi.org/10.1093/bjps/axy062.

Wilholt, T. (2009) Bias and values in scientific research. *Studies in History and Philosophy of Science Part A*. *40*(1), 92–101. https://doi.org/10.1016/j.shpsa.2008.12.005.

Wu, J. (2023) Epistemic advantage on the margin: A network standpoint epistemology. *Philosophy and Phenomenological Research*. *106*(3), 755–777. https://doi.org/10.1111/phpr.12895.

Wynne, B. (1991) Knowledges in context. *Science, Technology, & Human Values*. *16*(1), 111–121. https://doi.org/10.1177/016224399101600108.

Zollman, K. J. S. (2007) The communication structure of epistemic communities. *Philosophy of Science*. *74*(5), 574–587. https://doi.org/10.1086/525605.

Zollman, K. J. S. (2010) The epistemic benefit of transient diversity. *Erkenntnis*. *72*(1), 17–35. https://doi.org/10.1007/s10670-009-9194-6.