



Free information disrupts even Bayesian crowds

Jonas Stein^{a,1} , Shannon Cruz^{b,1} , Davide Grossi^{c,d,e,1} , and Martina Testori^{f,1,2}

Affiliations are included on p. 6.

Edited by Simon Asher Levin, Princeton University, Princeton, NJ; received July 11, 2025; accepted February 23, 2026

A core tenet underpinning the conception of contemporary information networks, such as social media platforms, is that users should not be constrained in the amount of information they can freely and willingly exchange with one another about a given topic. By means of a computational agent-based model, we show how even in groups of truth-seeking and cooperative agents with perfect information-processing abilities, unconstrained information exchange may lead to detrimental effects on the correctness of the group's beliefs. If unconstrained information exchange can be detrimental even among such idealized agents, it is prudent to assume it can also be so in practice. We therefore argue that constraints on information flow should be carefully considered in the design of communication networks with substantial societal impact, such as social media platforms.

wisdom of crowds | homophily | epistemic inequality | social learning | Bayesian agent based model

A cornerstone view in the contemporary conception of digital communication is that no piece of information between two individuals is ever one piece too many, when individuals are willing to communicate with one another. Crucially, this idea still pervades the narrative projected by social media platforms around their societal role, rooted in the digital activism mantra that “information wants to be free” (1). Even arguments in favor of, or opposed to, specific technical design features of social media platforms, such as full-text search of posts, quote-posts, or even fact-checking, are often framed with respect to the conviction that information generation by users should be easy and information provision to users should be seamless. This idea extends beyond social media, permeating broader discussions around openness in the digital age. Exchanges between practitioners in education (2), patients and clinicians in healthcare (3), or researchers collaborating across institutions (4) are frequently described in similar terms: The more information is shared among individuals, the more epistemically effective the interaction becomes. In these domains, the free flow of information is seen as essential for fostering mutual learning, improving patient treatment plans, and enhancing the quality and impact of research, thereby reinforcing its inherent value.

In this paper, we scrutinize the above tenet from a wisdom-of-crowds perspective (5). In their simplest form, these settings involve a group of agents, each endowed with private information about the value of a binary variable of interest. The group then needs to aggregate the members' private information and make a collective decision about the correct value of the (binary) variable. A substantial body of work using both experimental and computational techniques (see ref. 6 for a recent overview) has already shown that information sharing networks strongly influence the epistemic performance of groups in such settings and that, crucially, different network topologies induce different epistemic effects. Within such a context, our paper analyzes epistemic performance of groups as a function of two parameters of peer interaction that abstract away from the explicit consideration of network topology: homophily and information-exchange capacity. Specifically, we minimally enrich the classical wisdom-of-crowds setting by allowing our agents to reveal their private information to one another in bilateral encounters. In doing so, we follow recent work (7–10) but enrich it in one specific dimension by manipulating how much private information can be revealed in such encounters.

This simple feature allows us to investigate the effects of various levels of information-sharing capacity and homophily on the epistemic performance of crowds. Importantly, we investigate such information-sharing effects under the most ideal information processing and sharing conditions during encounters: a crowd consisting of perfect, i.e., Bayesian, information-processing agents, who communicate truthfully and fully cooperatively. At the same time, we place these agents in interaction conditions that capture a pervasive feature of many real-world information networks: homophily. The choice of these specific model features is methodologically motivated (we do, however, relax them to ascertain

Significance

A foundational assumption in the design of digital communication platforms is that information sharing is always beneficial. Using an agent-based model of truth-seeking and cooperative agents with perfect information-processing abilities, we show that unconstrained information exchange can reduce the accuracy of group beliefs, particularly in socially homophilous settings. Because these negative effects emerge even in crowds of agents that process information perfectly and cooperatively, a plausible risk exists of them being even more severe in less ideal, real-world contexts. Our findings challenge the common design principle that information should always “be free,” in the sense of being unlimitedly shareable, and suggest that constraining the amount of information shared—rather than maximizing it—may actually improve collective epistemic outcomes.

Author contributions: J.S., S.C., D.G., and M.T. designed research; J.S., S.C., D.G., and M.T. performed research; J.S. implemented the model; S.C. analyzed data; and J.S., S.C., D.G., and M.T. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2026 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹J.S., S.C., D.G., and M.T. contributed equally to this work.

²To whom correspondence may be addressed. Email: testori.martina@gmail.com.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2518472123/-/DCSupplemental>.

Published April 1, 2026.

robustness of our findings, see *SI Appendix, section 3*): If unlimited information exchange has negative impact under even the most ideal assumptions about individuals, it should be questioned altogether as a foundational design principle for information-sharing networks.

Through agent-based simulations, we observe that the ability to exchange more information negatively affects the decision-making performance of the group when information is shared in the context of homophilous encounters, even under such ideal individual information processing and sharing conditions. This largely aligns with previous findings (6), while focusing on the combined effects of homophily and information-sharing capacity. Our contribution should be interpreted as a computationally aided thought experiment, showing how the “information wants to be free” principle in the design of digital communication networks should be subject to qualification: being constrained in how much information can be shared may actually improve collective decision-making when homophily is high.

Model. We model crowds using an agent-based model (11) and investigate the interplay of two key parameters of the model: *communication capacity* (i.e., the number of pieces of evidence that agents are able to exchange in each interaction) and the level of *homophily* that guides agents’ interactions (Fig. 1). Homophily, the tendency to interact with similar others, is a strong force in humans (12) and one that deeply shapes interaction in social and digital information networks, especially when algorithmically curated. Algorithmic curation (i.e., the computational process that filters the information to which users are exposed) is a pervasive design feature of contemporary social media and one that has been repeatedly linked to homophilous information exchange (13–15).

Our simulation experiments assume a crowd of Bayesian agents seeking to learn the true state of the world. Each agent is endowed with one piece of evidence that either points toward the true state *A* or the false state *B* of the world (see *Materials and Methods* and *SI Appendix* for more details). Agents in the model form beliefs by applying the Bayes’ rule on the pieces of evidence they have access to. They influence each other by exchanging pieces of evidence. Each piece of evidence consists of two elements: a signal (pointing to either state *A* or state *B*) and the quality of that signal (i.e., the probability that the signal is correct). This setup can be illustrated by thinking of the agents as carrying out experiments about the state of the world using imperfect experimental equipment of known accuracy. A piece of evidence is thus the outcome of an observation together with information about the accuracy of the equipment that generated the observation. It is important to stress that equipment with higher accuracy can still occasionally return wrong observations, albeit less frequently than equipment with lower accuracy, and, vice versa, equipment with low accuracy can still provide correct observations. As a consequence, accurate evidence may still point to the incorrect state of the world, and inaccurate evidence to the correct one. This richer model of what constitutes information sets our model apart from most existing literature on information dynamics in groups (see *Discussion* section).

At model initialization, agents observe their own pieces of evidence. Even though we assume that these observations are more likely to be correct than not, a minority of agents will typically hold incorrect beliefs because their imperfect equipment returns the wrong signal by chance. After initialization, agents seek information about others’ evidence. The level of homophily present in the group affects how likely these exchanges are to

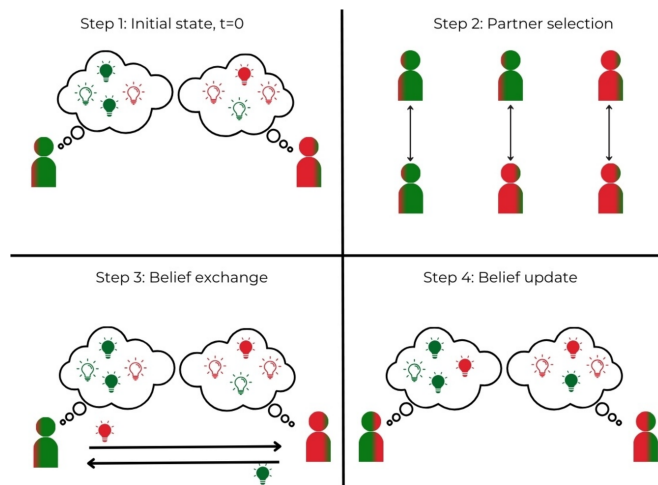


Fig. 1. Each agent is represented by a colored figure, with the color indicating their belief: Green-leaning agents are more confident that the true state is *A*, while red-leaning agents are more confident that the true state is *B*. Step 1: Initial state ($t = 0$). Each agent is endowed with a private piece of evidence about the state of the world. Each piece of evidence (green lightbulb for *A*, red lightbulb for *B*) comes with a known quality—the probability that the signal is accurate (full/empty lightbulb for high/low quality). Agents use Bayes’ rule to form an initial belief based on this single observation. Some agents start with a correct belief (favoring *A*), others with an incorrect one (favoring *B*) due to randomness in their initial evidence. Step 2: Partner selection. At each time step, agents are stochastically matched with another agent to exchange information. The likelihood of being paired increases with belief similarity, a dynamic controlled by homophily: When homophily is high, agents are more likely to interact with others who hold beliefs similar to their own. Step 3: Belief exchange. Agents exchange up to k pieces of evidence, determined by their communication capacity. Each agent selects their best available evidence—i.e., the most accurate signals—and shares them in proportion to their current belief. For example, an agent strongly favoring *A* will mostly share high-quality signals supporting *A*, but may also share some evidence supporting *B*. Note: Exchanges occur in one direction, meaning that one agent might not share information with the same agent they received information from. Step 4: Belief update. After the exchange, each agent updates their belief by applying Bayes’ rule to the new evidence received. Over time, these updates can lead agents to converge toward the true state *A*, though this depends on both who they interact with (influenced by homophily) and how much information they can exchange (communication capacity).

occur between agents with similar beliefs about the state of the world. Furthermore, the communication capacity determines the number of pieces of evidence that can be exchanged in each interaction. Importantly, evidence exchange is cooperative: Agents share only their highest-quality evidence and select evidence in favor of *A* or *B* with probability proportional to their current belief. In other words, agents try as much as possible, within the given communication constraints, to convey their best available evidence. Agents interact until a stable state is reached at which no further belief change occurs. Despite the above idealizations on the information processing ability and cooperation of agents, homophily drives agents with incorrect beliefs to interact with agents that hold similarly incorrect beliefs. And if no constraints are in place on how much information can be exchanged during such encounters, agents with incorrect beliefs will necessarily exchange more evidence that points to the incorrect state of the world. This is, in essence, the phenomenon that our experiments set out to scope, quantify, and analyze.

Focus and Metrics: Effects of Information-Exchange Capacity.

Our main focus is on the group of agents whose initial observations point toward the wrong state, which we refer to as Group *B* (as opposed to Group *A*). We are interested in assessing the extent to which, under the above ideal conditions, this group of “unlucky” observers is able to improve the quality

of their beliefs through interaction, that is, what we refer to as *epistemic gain*. Intuitively, the larger the capacity individuals have to communicate evidence with each other, the more likely it is that agents with incorrect initial beliefs will shift toward the correct belief. We will see, however, that this is not necessarily the case. Limiting communication capacity, and therefore the amount of evidence that can be exchanged in each interaction, serves the epistemic gain of the disadvantaged group better under conditions of high homophily.

Results

Our simulation experiments assume a crowd of $N = 100$ agents interacting according to the above model (see *Materials and Methods* and *SI Appendix* for more details). Even sample runs of the model illustrate that effects of communication capacity are contingent upon homophily (Fig. 2). Crucially, beliefs only converge around the true state of the world (A) when communication capacity is high but homophily is low (Fig. 2, panel *iv*). In this situation, all agents exchange all relevant information, thereby carrying agents toward the correct belief, regardless of their initial belief. Epistemic performance of the population (i.e., how close the average belief is to the correct belief that A is the true state) is lowest in the high-homophily, high-capacity regime (Fig. 2, panel *ii*), exhibiting a population split into polarized groups at opposite ends of the belief spectrum. Low levels of communication capacity result in moderate epistemic performance, independent of the level of homophily (Fig. 2, panels *i* and *iii*).

Across model runs, agents' epistemic performance depends not only on homophily and communication capacity, but also on which beliefs agents draw from their private observations at initialization. To examine these effects, we conducted a

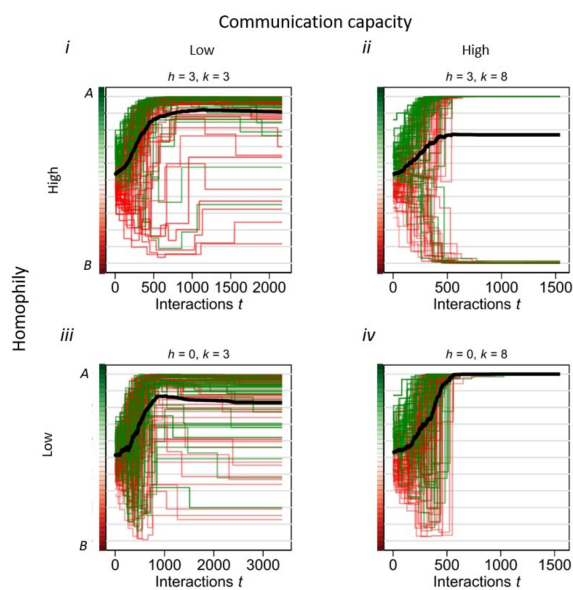


Fig. 2. Belief evolution over interactions by homophily h and communication capacity k . Panel (*i*): high homophily, low communication capacity; Panel (*ii*): high homophily, high communication capacity; Panel (*iii*): low homophily, low communication capacity; Panel (*iv*): low homophily, high communication capacity. Panel (*i* and *iii*): Low communication capacity results in moderate belief diversity and epistemic performance, regardless of the level of homophily. Panel (*ii*): High communication capacity and high homophily result in belief polarization and low epistemic performance. Panel (*iv*): Epistemic performance is maximal at high capacity and no homophily. Green (red) lines represent agents with an initial belief in favor of A (B).

multilevel analysis with agents (Level 1, $N = 100$ per run) nested within runs (Level 2, $N = 70,000$). Individual-level epistemic gain (ρ_i) was then regressed on initial beliefs (B_i , a level-1 variable), communication capacity (k , a standardized level-2 variable), homophily (h , a standardized level-2 variable), and the interactions among these predictors. The results (see *SI Appendix* for mathematical details and full tables) suggest that communication capacity by itself has a very limited effect on realized epistemic gain (unstandardized $\beta = -0.01$). The interaction between communication capacity and homophily ($\beta = -0.12$), however, suggests that communication capacity tends to exacerbate homophily's already negative effect on epistemic gain ($\beta = -0.48$). These effects also differ for agents with different initial beliefs. Agents with initial beliefs that are closer to the correct state of the world (Fig. 3, *Top Left*) tend to have greater epistemic gain regardless of communication capacity ($\beta = 0.63$). More importantly, the interaction terms indicate that these agents tend to benefit more from increased communication capacity ($\beta = 0.72$) and suffer less from the negative effects of homophily ($\beta = 0.17$) relative to agents with incorrect initial beliefs (Fig. 3, *Bottom Left*).

We then probed the group- and population-level outcomes of these dynamics by conducting a series of linear regressions in which we regressed group- and population-level mean epistemic gain and inequality on communication capacity, homophily, and their interaction. At the group level, the results indicate that the mean epistemic gain in both the community of initially correct agents (Group A ; W_A) and the community of initially incorrect agents (Group B ; W_B) is positively associated with communication capacity, but negatively associated with homophily (Fig. 3, *Central column*). However, whereas the positive effect of communication capacity outweighs the negative effect of homophily in group A (standardized $\beta = 0.26$ vs. -0.16), the opposite is true for group B ($\beta = 0.17$ vs. -0.47). The interaction effects also indicate that, similar to what was observed at the individual level, a higher communication capacity exacerbates the negative effects of homophily more for group B ($\beta = -0.17$) than for group A ($\beta = -0.12$). As a result of these dynamics, the epistemic inequality between the two groups ($W_A - W_B$) tends to get larger as capacity increases ($\beta = 0.11$) and as homophily increases ($\beta = 0.65$), particularly when both are high simultaneously ($\beta = 0.13$).

Because the effects of homophily are particularly harmful to mean epistemic gain for group B , the negative effect of homophily ($\beta = -0.34$) also tends to outweigh the positive effect of communication capacity ($\beta = 0.22$) on population-level mean epistemic gain (W ; see Fig. 3, *Top Right*). Indeed, the interaction between these variables indicates that increased communication capacity exacerbates the negative effects of homophily at the population level ($\beta = -0.16$). The effect of homophily on epistemic inequality, measured as the Gini coefficient of individual epistemic gains (G ; $\beta = 0.38$; see Fig. 3, *Bottom Right*), also outweighs the equalizing effect of communication capacity ($\beta = -0.28$). When homophily is high, higher capacity also tends to make inequality worse ($\beta = 0.17$).

To be clear, these results focus on how realized epistemic gains compare to the maximum possible epistemic gains. Even in the conditions where gains are lowest, they are rarely zero; information sharing even under the least favorable conditions in our model still results, on average, in some benefit for the population as a whole. The important finding is that under high homophily and high information capacity, epistemic gains are consistently lower *relative* to other conditions. In other

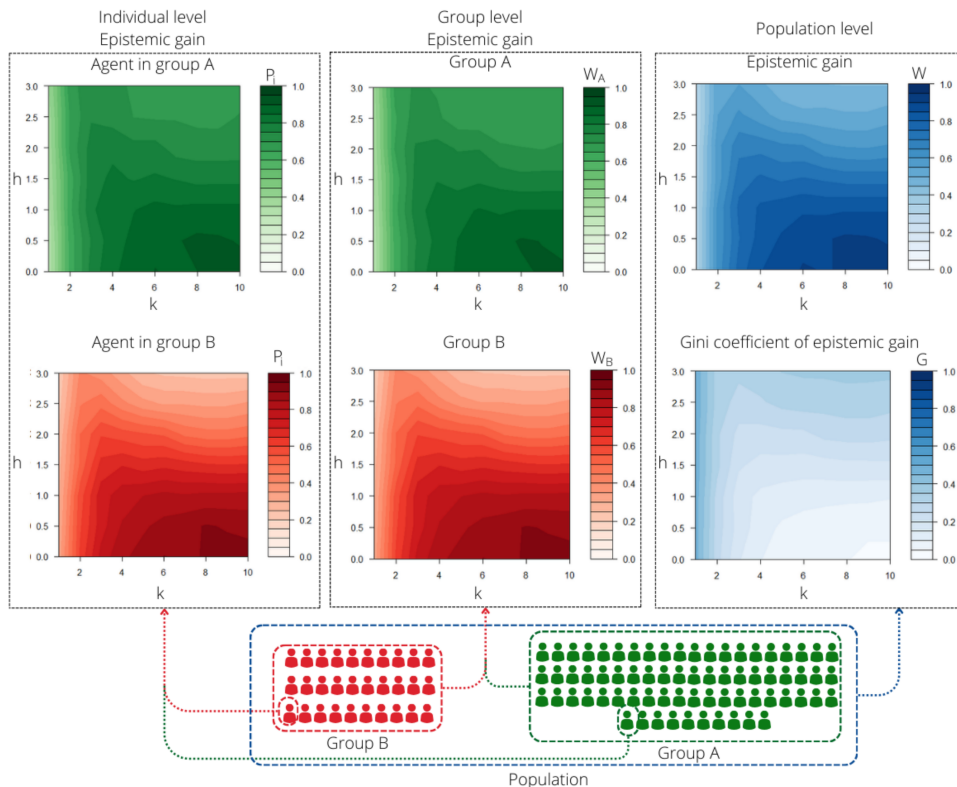


Fig. 3. Overview of (1) agents' A and B epistemic gain (ρ_i ; first column); (2) groups' A and B epistemic gain (W_A, W_B ; Central column); (3) population epistemic gain (W ; Top row, last column); (4) inequality across the population in the epistemic gain (G ; Bottom row, last column), as a function of homophily h and communication capacity k .

words, although it is encouraging that epistemic gains occur even under the worst-case scenario in our model, this finding is less positive when viewed against the substantially greater gains that could have been achieved under different information-sharing conditions.

We also conducted model variations and robustness checks (*SI Appendix, section 3*) to assess whether the results depend on specific modeling choices. The findings are robust to changes in population size, to shifts in the distribution and number of initial signals assigned to agents, and to modifications of structural features. In particular, results remain mostly consistent when allowing information to be shared simultaneously with multiple agents (one-to-many broadcasting), when restricting agents to share only arguments that support their own posterior beliefs (noncooperative sharing), and when salient group identities make it more likely that agents reject information from outgroup members and accept information from ingroup peers (group salience).

Discussion

Much has already been written about the effects of homophily on opinion formation in groups, since at least social judgment theory (16). Directly relevant to our work is modeling research on the effects concerning polarization (17–20), especially within wisdom-of-crowds settings (7, 8, 21, 22). In the bulk of this literature, but with the notable exception of refs. 21 and 22, which we come back to below, information exchange is typically modeled as a form of belief averaging, thereby abstracting away from an explicit model of the evidence that is exchanged and determines belief formation. Our modeling of information exchange aims instead to represent explicitly each individual

piece of evidence that is exchanged via social interaction. This modeling choice enables observations that target the combined effect of homophily and the amount of information that agents are able to exchange in interaction. We model pieces of evidence as observations from Bernoulli trials with a specific quality (*SI Appendix*). In doing this, we model information in a slightly more stylized manner than previous work has proposed (21, 22), which is, to the best of our knowledge, the only other work attempting to model evidence explicitly. There, information is modeled using richer structures known as Bayesian networks, of which our model is a special case. The more stylized model we use allows us to represent the amount of information exchanged in a parsimonious manner, enabling us to develop an angle on the interplay between amount of information exchanged and homophily.

Our findings are driven by an agent-based model methodology. They do, however, align with observations from empirical studies on the wisdom of crowds and information exchange. For example, negative epistemic effects of social influence among peers with similar ideological priors were empirically observed (23), although information capacity was not treated as an explicit experimental parameter. Related empirical research has focused also on the role of partisan identities, showing that group salience can decrease influence from out-group members (24). While not a part of our main model, additional findings in *SI Appendix* show that such measures of identity salience produce effects similar to homophily, and can complicate the epistemic progress of populations with free information flow. In showing how moderate amounts of communication can produce large epistemic gain despite high levels of homophily, our results further align with empirical findings suggesting that the wisdom of crowds is robust to homophilous partitioning (25).

A feature of opinion dynamics that is incorporated in some empirical as well as agent-based modeling research (17) is negative influence, according to which contrasting ideological views can sometimes amplify polarization tendencies (26). Our analysis shows that even in the absence of strong polarization drivers such as negative influence, epistemic group performance declines under high homophily and information-capacity regimes.

Finally, we note that our model is circumscribed by a setting in which agents are confronted with a binary decision [cf. (27)]. While this may be applicable to many circumstances, such as forming beliefs on whether a law should be passed or not, whether a drug should be approved or disapproved, or whether a defendant is guilty or innocent, it is important to point out that other decision environments may be more complex. A natural direction for future work is to investigate if the results obtained here also apply to ordinal (28, 29) or multidimensional (30) decision problems.

In general, our findings identify a core trade-off in the design of information networks between homophily and communication capacity. When communication is homophilous, the ability to exchange large amounts of information penalizes precisely the agents who would benefit most from accessing accurate information. If such negative effects can be established in groups of ideal agents like the ones we study, caution suggests that they should remain a concern when agents' information processing abilities and cooperative tendencies are less than ideal, like in real-world information networks.

The above insight aligns with previous findings, analytical (31, 32), simulation-based (7, 8), and empirical (10), that highlight how communication may both improve crowd-wisdom, by disseminating high-quality information within the group, as well as damage it, by introducing too much consistency across agents' beliefs and thereby disrupting effective information aggregation.

To conclude, our results align with previous work in suggesting that groups' epistemic performance crucially depends on structural features of individuals' interactions, such as homophily. They suggest, in addition, that tradeoffs exist between epistemic performance, interaction structure, and the amount of information that individuals are able to exchange in interaction. Any implementation of the "information wants to be free" tenet in the design of information networks should, therefore, not ignore the structural features of interaction that constrain or facilitate information flow.

Materials and Methods

Experimental Design. We implemented our model in NetLogo, version 6.2.0 (33). The complete replication package can be found on OSF <https://osf.io/vg3y2/overview>. The model is instantiated for a crowd of $N = 100$ agents and each experiment stopped after 1,000 interactions during which no agent crossed the 0.5 posterior threshold. We conducted 1,000 simulation runs per parameter combination of h and k . Sensitivity analyses in *SI Appendix* confirm that results stay similar in populations of larger sizes.

Environment. There are two states of the world: A and B . We assume that A is the true state and that each agent has a prior $\Pr(A) = \Pr(B) = 0.5$ as typically done in maximum-likelihood estimation settings (34). We work with a discrete time stochastic model with a finite population.

Initialization of Agents' Beliefs. At initialization (time $t = 0$), each agent i is endowed with a set $S_i (|S_i| < k)$ of private pieces of evidence about the state of the world. These pieces of evidence are pairs (s, q) , where

- s is a random variable which can take either value A (the information supports state A) or B (the information supports state B);

- q is a parameter defining the Bernoulli distribution

$$q = \Pr(s|A) = 1 - \Pr(s|B) \in (0.5, 1). \quad [1]$$

from which s is drawn. q describes the quality of the specific piece of evidence, that is, the conditional probability of receiving signal s given the state of the world. We assume these quality parameters to be drawn from a beta distribution and transformed according to

$$q \sim \frac{\text{Beta}(2, 9)}{2} + 0.5. \quad [2]$$

Upon observation of these pieces of evidence (s, q) , agents obtain their initial belief about the state of the world via Bayes' rule:

$$\mathcal{B}_{i,0} = \Pr(A | S_{i,0}), \quad [3]$$

that is, the belief of agent i for state A at time 0. Clearly, the belief of agent i for state B at time 0 is $1 - \mathcal{B}_{i,0}$.

In choosing a right-skewed distribution of $\text{Beta}(2, 9)$ with a relatively low average signal quality $\bar{q} = 0.59$, we ensure that a sizable minority of agents starts off with initial evidence pointing toward B as the true state. Sensitivity analyses in the *SI* reveal that although general patterns remain unchanged, results become weaker in less right-skewed Beta distributions. Here, higher average signal quality makes it easier for agents to infer the true state of the world A and harder for subpopulations to develop stable incorrect beliefs. Results are robust to a larger endowment of initial signals per agent $1 < |S| < 10$.

Dynamics of Agents' Beliefs. At each subsequent time step $t > 0$, each agent is paired with another agent with whom to interact and exchange information. This happens stochastically and agents are more likely to interact with agents similar to themselves in terms of their view of the true state of the world at the previous time $t - 1$: $\mathcal{B}_{*,t-1}$. The similarity score $\sigma_{ij,t}$ between two agents i and j at time t is calculated as

$$\sigma_{ij,t} = 1 - |\mathcal{B}_{j,t} - \mathcal{B}_{i,t}|. \quad [4]$$

The two agents are then matched with probability equal to

$$\frac{\sigma_{ij,t}^h}{\sum_{j \in J} \sigma_{ij,t}^h}, \quad [5]$$

where $h \in [0, 10]$ is the homophily parameter and $J = N \setminus \{i\}$ is the set of agents minus agent i . Once the partner has been selected, agents exchange k signals (i.e., pieces of evidence) so that

$$S_{i,t+1} = S_{i,t} \cup C_j^t, \quad [6]$$

where:

- $C_j^t = \text{TOP}_\alpha^t(A) \cup \text{TOP}_\beta^t(B)$, with TOP_α^t standing for the best (by quality) α signals supporting state A in j 's observation set at time t (similarly for B);
- $\alpha = \lceil \mathcal{B}_{j,t}(A) \cdot k \rceil$ and $\beta = \lfloor \mathcal{B}_{j,t}(B) \cdot k \rfloor$.

Intuitively, j passes to i the best (highest quality) signals she has for A and B proportionally to her posterior at time t . This information update models a form of informational "good faith" by agents, who are willing to exchange pieces of evidence that go against their current belief.

After each interaction, agents iterate over the set of newly received signals $S_{i,t} \setminus S_{i,t-1}$ to update their belief. For each new signal supporting A , they use Bayes' rule to obtain:

$$\mathcal{B}_{i,t+1} = \frac{q \cdot \mathcal{B}_{i,t}}{q \cdot \mathcal{B}_{i,t} + ((1 - q) \cdot (1 - \mathcal{B}_{i,t}))}, \quad [7]$$

where q is the quality of the signal received, giving thus also the belief for state B : $1 - \mathcal{B}_{i,t+1}$. As a result of this operationalization, as populations polarize, similarity scores σ increase for like-minded agents and decrease for agents with opposing beliefs. Consequently, receivers become more likely to select interaction partners who share their views. At the same time, senders gain

confidence in their beliefs and selectively share arguments more in line with their posterior. These dynamics generate a feedback loop: Rising polarization increases the likelihood of interactions with similar others, who in turn provide information that further reinforces existing views.

We computed the *realized information gain* as

$$\rho_i = \frac{\alpha_i}{\phi_i}, \quad [8]$$

where $\phi_i = 1 - B_i$ is the difference between the full-information posterior 1 and i 's initial belief. Analysis at the agent-level comprised multilevel models with agents ($N = 100$) nested within runs ($N = 70,000$), with individual-level epistemic gain (ρ_i) as the dependent variable. Analysis at the group-level and the population-level were performed through linear regressions.

Data, Materials, and Software Availability. Replication package and original data have been deposited in Open Science Framework (<https://osf.io/vg3y2/overview>) (35).

ACKNOWLEDGMENTS. We would like to thank the Lorentz Centre, the organizers—Dr. Francesca Giardini, Dr. Jens Madsen, Prof. Jos Hornikx—and the participants of the “Computational modelling of social complexity in

argumentation” for giving us the opportunity to meet each other and develop this project. Davide Grossi acknowledges support from the [Hybrid Intelligence Center](#), a 10-y program funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research (NWO) and by the European Union under the Horizon Europe project [Percycles](#) (Participatory Democracy that Scales). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.



**Funded by
the European Union**

Author affiliations: ^aDepartment of Sociology, Faculty of Behavioral and Social Sciences, University of Groningen, Grote Rozenstraat 31, 9712 TG, Groningen, The Netherlands; ^bDepartment of Communication Arts and Sciences, The Pennsylvania State University, University Park, PA 16802; ^cBernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, Faculty of Science and Engineering, University of Groningen, Groningen 9700 AK, The Netherlands; ^dAmsterdam Center for Law & Economics, University of Amsterdam, Amsterdam 1001 ND, The Netherlands; ^eInstitute for Logic, Language and Computation, University of Amsterdam, Amsterdam 1098 XH, The Netherlands; and ^fBusiness, Operations and Strategy, Greenwich Business School, University of Greenwich, London SE10 9LS, United Kingdom

1. D. E. Denning, Concerning hackers who break into computer systems. *High noon on the electronic frontier: Conceptual issues in cyberspace* **137164** (1996).
2. United Nations Educational, Scientific and Cultural Organization (UNESCO), Open educational resources (oer) (2025). <https://www.unesco.org/en/open-educational-resources>. Accessed 20 February 2026.
3. J. L. Wolff *et al.*, Inviting patients and care partners to read doctors' notes: Opennotes and shared access to electronic medical records. *J. Am. Med. Inform. Assoc.* **24**, e166–e172 (2017).
4. M. Arentoft *et al.*, Agreement on reforming research assessment (2022).
5. B. Grofman, G. Owen, S. L. Feld, Thirteen theorems in search of the truth. *Theor. Decis.* **15**, 261–278 (1983).
6. D. Centola, The network science of collective intelligence. *Trends Cogn. Sci.* **26**, 923–941 (2022).
7. U. Hahn, M. von Sydow, C. Merdes, How communication can make voters choose less well. *Top. Cogn. Sci.* **11**, 194–206 (2019).
8. U. Hahn, J. U. Hansen, E. J. Olsson, Truth tracking performance of social networks: How connectivity and clustering can make groups less competent. *Synthese* **197**, 1511–1541 (2020).
9. J. Becker, D. Brackbill, D. Centola, Network dynamics of social influence in the wisdom of crowds. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E5070–E5076 (2017).
10. J. Lorenz, H. Rauhut, F. Schweitzer, D. Helbing, How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 9020–9025 (2011).
11. A. Flache, C. A. de Matos Fernandes, “Agent-based computational models” in *Research handbook on analytical sociology* (Edward Elgar Publishing, 2021), pp. 453–473.
12. M. McPherson, L. Smith-Lovin, J. M. Cook, Birds of a feather: Homophily in social networks. *Ann. Rev. Sociol.* **27**, 415–444 (2001).
13. E. Bakshy, S. Messing, L. A. Adamic, Exposure to ideologically diverse news and opinion on Facebook. *Science* **348**, 1130–1132 (2015).
14. A. Boutyline, R. Willer, The social structure of political echo chambers: Variation in ideological homophily in online networks. *Polit. Psychol.* **38**, 551–569 (2017).
15. F. P. Santos, Y. Lelkes, S. A. Levin, Link recommendation algorithms and dynamics of polarization in online social networks. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2102141118 (2021).
16. M. Sherif, C.I. Hovland, Social judgment: Assimilation and contrast effects in communication and attitude change (1961).
17. A. Flache *et al.*, Models of social influence: Towards the next frontiers. *J. Artif. Soc. Soc. Simul.* **20**, 2 (2017).
18. G. Deffuant, D. Neau, F. Amblard, G. Weisbuch, Mixing beliefs among interacting agents. *Adv. Complex Syst.* **3**, 87–98 (2000).
19. M. H. DeGroot, Reaching a consensus. *J. Am. Stat. Assoc.* **69**, 118–121 (1974).
20. R. Hegselmann *et al.*, Opinion dynamics and bounded confidence models, analysis, and simulation. *J. Artif. Soc. Soc. Simul.* **5**, 3 (2002).
21. L. Assaad *et al.*, A Bayesian agent-based framework for argument exchange across networks. *arXiv [Preprint]* (2023). <http://arxiv.org/abs/2311.09254> (Accessed 20 February 2026).
22. U. Hahn, C. Merdes, M. von Sydow, Knowledge through social networks: Accuracy, error, and polarisation. *PLoS One* **19**, e0294815 (2024).
23. J. Stein, V. Frey, A. van de Rijt, Realtime user ratings as a strategy for combatting misinformation: An experimental study. *Sci. Rep.* **13**, 1626 (2023).
24. D. Guilbeault, J. Becker, D. Centola, Social learning and partisan bias in the interpretation of climate trends. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9714–9719 (2018).
25. J. Becker, E. Porter, D. Centola, The wisdom of partisan crowds. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 10717–10722 (2019).
26. C. A. Bail *et al.*, Exposure to opposing views on social media can increase political polarization. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9216–9221 (2018).
27. J. A. Becker, D. Guilbeault, E. B. Smith, The crowd classification problem: Social dynamics of binary choice accuracy. *Manag. Sci.* **68**, 3949–3965 (2022).
28. S. Bikhchandani, D. Hirshleifer, I. Welch, A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* **100**, 992–1026 (1992).
29. A. V. Banerjee, Quarterly journal. *Q. J. Econ.* **107**, 797–817 (1992).
30. D. Lazer, A. Friedman, The network structure of exploration and exploitation. *Adm. Sci. Q.* **52**, 667–694 (2007).
31. B. Peleg, S. Zamir, Extending the condorcet jury theorem to a general dependent jury. *Soc. Choice Welf.* **39**, 91–125 (2012).
32. M. Pivato, Epistemic democracy with correlated voters. *J. Math. Econ.* **72**, 51–69 (2017).
33. U. Wilensky, Netlog. Evanston, IL: Northwestern university. Center for Connected Learning and Computer-Based Modeling (1999). <https://ccl.northwestern.edu/netlog/>.
34. E. Elkind, A. Slinko, “Rationalizations of voting rules” in *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, A. D. Procaccia, Eds. (Cambridge University Press, 2015).
35. S. Cruz., D. Grossi, J. Stein, M. Testori., Free information disrupts even bayesian crowds. Open Science Framework. <https://osf.io/vg3y2/overview>. Deposited 3 March 2026.