



# OPEN The functional aspects of selective exposure for collective decision-making under social influence

Poong Oh<sup>1✉</sup>, Jia Wang Peh<sup>2</sup> & Andrew Schauf<sup>3</sup>

Opinion diversity is crucial for collective decision-making, but maintaining it becomes challenging in the face of social influence. We propose selective exposure as an endogenous mechanism that preserves opinion diversity by forming exclusive subgroups of like-minded individuals, or echo chambers, which have been often perceived as an obstacle to achieving collective intelligence. We consider situations where a group of agents collectively make decisions about the true state of nature with the assumption that agents *update* their opinions by adopting the aggregated opinions of their information sources (i.e., naïve learning), or alternatively, *replace* incongruent sources with more like-minded others without adjusting their opinions (i.e., selective exposure). Individual opinions at steady states reached under these dynamics are then aggregated to form collective decisions, and their quality is assessed. The results suggest that the diversity-reducing effects of social influence are effectively confined within subgroups formed by selective exposure. More importantly, strong propensities for selective exposure maintain the quality of collective decisions at a level as high as that achieved in the absence of social influence. In contrast, naïve learning allows groups to reach consensus, which are more accurate than initial individual opinions, but significantly undermines the quality of collective decisions.

The independence of individual opinions is crucial for collective decision-making because it ensures opinion diversity within groups<sup>1–5</sup>. However, maintaining opinion diversity is challenging because social influence among group members tends to reduce it<sup>6–8</sup>. Empirical studies have suggested the diversity-reducing effects of social influence in various settings since Asch's seminal study on social conformity, showing that groups under social influence tend to arrive at biased collective decisions<sup>9–14</sup>. Nevertheless, it has been documented that real-world groups perform well in decision-making across a wide range of domains, for example, from crowd voting for product designs<sup>15</sup> and stock price prediction<sup>16</sup> to COVID-19 detection<sup>17</sup> more recently. It is important to note that all these successful collective decisions were made in highly interactive communication environments where social influence was not just possible but facilitated, and where, therefore, the independence of opinions was hardly expected. The question then becomes, "How can groups perform well in decision-making *despite* social influence?" To address this question, we propose *selective exposure* as a plausible endogenous mechanism that can effectively manage the diversity-reducing effects of social influence to preserve opinion diversity at the global level, thereby maintaining the quality of collective decisions.

We consider situations where a group of agents collectively make decisions about the true state of nature by utilizing information distributed in the form of opinions held by individual members. Assuming that the true state remains *unknown* to individual agents, we exclude from our model self-correction, incentives, and payoffs based on the accuracy of opinions, whereas these factors are considered in previous studies<sup>6,8,18–22</sup>. Further, we assume that agents' initial opinions are *independent* and distributed around the true state with non-trivial errors. The errors in initial opinions, however, cancel each other out when they are aggregated, for example, by the majority rule<sup>23</sup>, median aggregation<sup>24</sup>, or mean aggregation<sup>25</sup>. In our model, therefore, a simple aggregation of initial opinions held by group members can yield sufficiently accurate collective decisions, with the collective accuracy increasing as group size grows, thereby achieving "collective intelligence". The accuracy achieved by a

<sup>1</sup>Wee Kim Wee School of Communication and Information, Nanyang Technological University, Singapore 639798, Singapore. <sup>2</sup>Information Systems Technology and Design Pillar, Singapore University of Technology and Design, Singapore 487372, Singapore. <sup>3</sup>NUS Cities, College of Design and Engineering, National University of Singapore, Singapore 117566, Singapore. ✉email: poongoh@ntu.edu.sg

simple aggregation of initial opinions is considered the *null model* (i.e., the quality of collective decisions that would have been achieved *without* social influence), against which we assess the qualities of both individual and collective decisions formed *under* social influence.

We formulate *social influence* following DeGroot's model of opinion dynamics<sup>26</sup>, where agents naïvely update their opinions by repeatedly adopting the majority or average opinion of their information sources (i.e., naïve learning). In some previous studies, naïve learning is considered as a means for agents to minimize the discrepancies with their information sources (i.e., normative social influence)<sup>27,28</sup>. At the same time, naïve learning can also be seen as a useful heuristic for agents to improve their opinions when access to information is limited. This is because an opinion updated through naïve learning is a “collective” decision in the sense that it is made by aggregating the opinions in one's information pool at a small and local scale (i.e., informational social influence)<sup>29,30</sup>. However, the repetition of naïve learning necessarily reduces opinion diversity within groups over time, regardless of whether it exerts normative or informational social influence, illustrating the diversity-reducing effects of social influence. In each round, the agent with the most extreme opinion is exposed to less extreme opinions than its own, and its opinion becomes less extreme through naïve learning in the following round. Especially when the patterns of influence from information sources to recipients are represented by a directed network, and the network is strongly connected (i.e., every agent is connected to everyone else by at least one directed path), all agents end up adopting a common opinion, and so the group reaches a consensus, as DeGroot's theorem states. We consider the convergence of opinions at steady states reached through naïve learning as the state where there is a lack of opinion diversity within groups.

A group consensus reached through naïve learning is generally not equal to the collective decision formed by a simple aggregation of initial opinions in the group<sup>31</sup>. Instead, it is determined by the *weighted* mean of initial opinions, where the weights are proportional to the eigenvector centralities of corresponding agents in the network<sup>32</sup> and represent the agents' relative influences on the consensus. Golub and Jackson<sup>33</sup> and Acemoglu et al.<sup>34</sup> show that naïve learning can still produce accurate group consensuses in the absence of “excessively” influential agents, modeled as hubs in wheel-shaped networks. However, the degree distributions and the distributions of eigenvector centralities of real-world social networks typically follow power laws<sup>35</sup>. Further, for social networks that are constantly growing, central nodes tend to become more central through preferential attachment processes<sup>36,37</sup> in contrast to the assumption made by Acemoglu et al.<sup>34</sup> Therefore, a well-connected large group is expected to reach a consensus that significantly deviates from the simple aggregation of initial opinions. In addition, those studies (also Buechel et al.<sup>30</sup> and Rauhut and Lorenz<sup>38</sup>) adopt rather “generous” assessment criteria (i.e., the deviation of group consensuses from the true state) without considering that the marginal improvement of collective accuracy necessarily decreases as group size increases<sup>39</sup>. When the effects of group size on collective accuracy are not accounted for, the effects of influential agents in large groups tend to be underestimated.

Other studies focus on individual attributes of agents, specifically stubbornness, as another potential source of the deviation of group consensuses from the true state. For instance, Acemoglu et al.<sup>34</sup> show that the presence of “forceful” agents—those who influence others but do not change their own opinions—results in a consensus biased toward the initial opinions of the forceful agents, increasing the deviation. Similarly, Buechel et al.<sup>30</sup> suggest that agents with low levels of conformity (i.e., high self-confidence) have high impacts on group consensus (i.e., enhanced opinion leadership), and Anufriev et al.<sup>27</sup> show that agents with lower sensitivity to disagreement with others exert more influence on group consensuses. Similarly, Zafeiris and Vicsek<sup>40</sup> demonstrate that collective decisions are optimal when the pliancy and competence of group members are negatively correlated, implying that the stubbornness of incompetent members contributes to further deviation. The stubbornness of agents in those studies functions in a similar way to the centralities of the agents in the sense that stubborn agents exert disproportionate influence on group consensuses, increasing the deviation. Stubborn agents can be viewed as those having no incoming edges at all (i.e., not being influenced by others) or much fewer incoming edges than outgoing edges (i.e., influencing others more than they are influenced by others) in communication networks, thus exhibiting higher centralities and exerting greater influence on group consensuses than others. To examine the isolated impacts of naïve learning in comparison to those of selective exposure, we exclude confounding factors that could undermine the quality of collective decisions. Specifically, we assume that edges from information sources to recipients are formed *uniformly* at random as in the Erdős-Rényi random graphs<sup>41</sup> and that all agents are *equally* responsive to the discrepancies with their information sources<sup>42</sup> (i.e., the absence of both “excessively” influential and “stubborn” agents). Under this assumption, naïve learning should result in a group consensus that is as accurate as the collective decision made by a simple aggregation of opinions.

We then incorporate *selective exposure* into our model as an alternative means for agents to minimize the discrepancies between their own opinions and those of their information sources. In previous literature, “selective exposure” refers to the tendency of people to seek out information that supports their existing views and opinions while avoiding information that challenges their own perspectives<sup>43</sup>. This tendency is commonly observed and has been studied in a wide range of domains, including communication<sup>44</sup>, information processing<sup>45</sup>, attitude and belief formation<sup>46</sup>, as well as decision making<sup>47</sup>. Further, selective exposure has been viewed as a major driving force in the formation and evolution of social networks<sup>48–51</sup>. In the context of collective opinion dynamics, some previous studies formulate selective exposure by assigning differential weights to edges in proportion to the opinion similarity between agents (e.g., interaction frequency<sup>52</sup>, relational intensity<sup>53</sup>, allocated attention<sup>54</sup>, and incentives for homophilic edges<sup>19</sup>), while assuming the connectedness of networks to be exogenously determined and constant over time. However, when the network of agents remains strongly connected, the effects of selective exposure on opinion dynamics are highly limited. Although disagreement among agents is observed temporarily and the convergence of opinions is significantly delayed, groups *always* reach consensus regardless of the level of selective exposure<sup>52,53,55</sup>. In our model, in contrast, connections between agents are assumed to be binary rather than weighted and can be either formed or removed by agents themselves, allowing for endogenous changes in network structure. Specifically, agents can reconstruct their information pools by *removing* the connections

to the information sources whose opinions differ most from their own and *adding* new connections to others whose opinions are more congruent with their own than the removed ones. By doing so, agents can reduce the discrepancies with their information sources without adjusting their own opinions.

When endogenous changes in network structure are permitted, an immediate consequence of selective exposure is the fragmentation of a group into small and exclusive subgroups of like-minded agents<sup>54,56,57</sup>, or *echo chambers*, each of which converges to its own unique consensus<sup>27,53</sup>. Whereas previous studies focus on the conditions for fragmentation<sup>57</sup>, the stability of the resulting fragmented structures<sup>19</sup>, the impacts of fragmented structures on the speed of convergence<sup>53</sup>, and the opinion diversity within groups<sup>53,58</sup>, there has been limited research that systematically investigates the effects of a fragmented structure on the quality of collective decisions. Instead, many studies view fragmented or polarized structures as inherently harmful and so as something to be “overcome” through various intervention measures<sup>19,56,57</sup>. Many political researchers, in particular, perceive selective exposure as the primary culprit of political polarization, social distrust and conflict, and intergroup hostility<sup>59–62</sup>. Furthermore, they argue that it reduces the chance for individuals to learn from diverse views, resulting in biased and poor collective decisions<sup>63,64</sup>. From this perspective, political researchers have suggested “cross-cutting exposure” (i.e., the exposure to and learning from diverse views) as a potential solution to these problems<sup>61,65,66</sup>.

Contrary to this near-unanimous belief prevailing in political research and other areas, fragmented structures—more generally *modular* structures—in which individuals are tightly connected to others with similar traits and form cohesive local clusters, are known to be beneficial for maintaining diversity and enhancing groups’ ability to survive and adapt in hostile and uncertain environments<sup>7,58,67–72</sup>. Along these lines, Kao and Couzin<sup>39</sup> investigate the effects of modular structures on collective decisions, although neither opinion updates nor endogenous changes in network structure are taken into account. The authors stress that the marginal improvement of collective accuracy diminishes, because the amount of information available within a group becomes saturated as its size increases. Therefore, the group size required to achieve a certain level of collective accuracy (i.e., “effective group size”) is typically smaller than the actual group size. Especially when individual opinions are correlated with one another, a modular structure enables a large group to behave as effectively as smaller groups, while retaining other benefits of its large group size. In addition, minority opinions, which might otherwise be overlooked when opinions are simply aggregated, can be better represented within a modular structure, thereby upholding opinion diversity at the global level. Pescetelli et al.<sup>73</sup> provide empirical evidence directly supporting this view, showing that modular groups composed of small independent subgroups consistently outperform non-modular groups in real-world forecast problems. Along with this line, Mann and Woolley-Meza<sup>7</sup> underscore the importance of modular structures in retaining intellectual diversity and fostering the generation of new knowledge within academic communities. This can be achieved by maintaining a basic level of isolation and independence between subfields and thereby preventing the dominance of specific models and methods. Adopting Kao and Couzin’s<sup>39</sup> concept of “effective group size,” we assess the quality of collective decisions achieved under different dynamics against the collective accuracy that a group of the *same size* would have achieved in the absence of social influence among group members.

We perform a series of numerical experiments focusing on the comparative statics of different levels of propensities for selective exposure with respect to group structure, opinion diversity, and the qualities of individual and collective decisions. The numerical experiments yield qualitatively consistent results across all the different settings examined. First, selective exposure fosters exclusive subgroups of like-minded individuals, leading to the emergence of modular and homophilic structures, confirming the findings of previous studies. Also, under strong propensities of selective exposure, opinion diversity *within* subgroups is completely eradicated as each subgroup converges to a local consensus that differs from those of other subgroups, as shown in previous studies. However, opinion diversity at the *global level* is perfectly preserved, suggesting that the diversity-reducing effects of social influence are effectively confined within each subgroup, with no further spread beyond it. Most importantly, the preserved opinion diversity, in turn, maintains the quality of collective decisions at a level as high as that achieved by a simple aggregation of initial opinions. In contrast, naïve learning allows a group to stay connected and thereby reach a consensus that turns out to be more accurate than initial individual opinions, as predicted by previous studies. However, when the effects of group size on the marginal improvement in collective accuracy are taken into consideration, the group consensus is significantly less accurate than both collective decisions made by a simple aggregation and those made under strong propensities for selective exposure across all the simulation settings. This is because naïve learning eradicates opinion diversity at the global level at a faster rate than that of the improvement of individual opinions. That is, the improvement of individual opinions through naïve learning comes at the expense of opinion diversity and collective accuracy, illustrating the trade-off between individual and collective performances.

A unique contribution of the current study is that it examines the causal path from *selective exposure* to the *quality of collective decisions* through the changes in *group structure* and *opinion diversity* in comparison to an alternative mechanism, naïve learning. The causal path is identified based upon a framework that integrates previously disparate concepts and findings, although none of them were completely unprecedented in previous literature. Despite the simplicity of our models, numerical experiments successfully replicate previous findings, indicating that the results we present do not contradict the findings of previous studies. Instead, the current results complement previous studies, providing a new insight into the understanding of the role that selective exposure plays in collective decision-making, which has been often perceived as a frustrating obstacle to achieving collective intelligence. The implication of the current results is straightforward: A fragmented group structure resulting from selective exposure prevents *both* information and influence from flowing across subgroups, while naïve learning allows *both* to flow. This is because communication channels between agents cannot separate information from the influence it carries. In that sense, the current findings simply suggest that blocking both information and influence (i.e., preserving diversity at the expense of additional information) is more beneficial

for collective decisions than allowing both to flow (i.e., gaining additional information at the cost of opinion diversity), reconfirming the crucial role of opinion diversity in collective decision-making.

The rest of this paper is organized as follows. The next section introduces the agent-based models examined in this study, with detailed descriptions of simulation settings and measures. The results from numerical experiments are then summarized in the subsequent section. The paper concludes with a discussion of the main findings and proposals for future research.

## Methods

### Classical models of collective decision-making

The conditions for achieving collective intelligence have been extensively studied. One of the earliest models is Condorcet's jury theorem<sup>23</sup>. The theorem considers a group of  $n$  members tasked with deciding between two alternatives, one of which is objectively better than the other but unknown to the members. Each member is assumed to choose the better decision with a probability of  $p > .5$ , which is at least slightly better than a random guess. Further, all decisions are assumed to be made independently and then aggregated into the group's decision by the majority rule. Under these assumptions, the probability of a group arriving at the correct decision  $P_G$  is always greater than individual competence  $p$ . Moreover,  $P_G$  increases monotonically and approaches 1 in the limit as group size increases.

Galton's field experiment<sup>24</sup> provides similar results. In his experiment, participants were asked to guess the weight of a live ox, and the median estimate of valid guesses (i.e., "valid" in the sense of being uninfluenced by others' guesses) was accurate within a 1% margin of the ox's true weight. Later, he found the mean estimate even more accurate<sup>25</sup>. Based on these observations, he argued that a simple aggregation of independent judgments can produce a sufficiently accurate collective outcome because the errors in individual judgments tend to cancel each other out when aggregated. Also, it can be intuitively conceived that the errors in collective outcomes decrease as group size increases, as indicated by the central limit theorem.

Even though Condorcet and Galton consider different types of opinions (i.e., binary or continuous opinions) and different methods for aggregating individual decisions (i.e., the majority rule or the median and mean aggregation), they both reach the same conclusion: Groups outperform individuals in decision-making, and the quality of collective decisions monotonically increases as a function of group size (see Section S1 in *Supplementary Materials* for more details). However, both share the unrealistic assumption that individuals in groups make their decisions *independently* without influencing or being influenced by others.

To formulate a model of opinion dynamics under social influence, DeGroot's model<sup>26</sup> assumes that individuals update their opinions by repeatedly adopting the average of others' opinions, which is formulated as a Markov process as follows. Individual opinions  $\mathbf{x}(t)$  at a time point  $t \in \{1, 2, 3, \dots\}$  are determined as the weighted means of others' opinions at  $t - 1$ :

$$\mathbf{x}(t) = \mathbf{W}\mathbf{x}(t - 1) \text{ or } \mathbf{x}(t) = \mathbf{W}^t\mathbf{x}(0), \text{ where } \mathbf{W} = [w_{ij}],$$

where  $w_{ij}$  denotes the relative influence of  $j$ 's opinion on the update of  $i$ 's opinion (i.e.,  $0 \leq w_{ij} \leq 1$ ), and  $\mathbf{W}$  is a row-stochastic transition matrix (i.e.,  $\sum_j w_{ij} = 1$  for all  $i$ ) that can be interpreted as the adjacency matrix of the network that describes the patterns of social influence among individuals in a group. Under these dynamics, individual opinions  $\mathbf{x}(t)$  converge to a steady-state distribution  $\mathbf{x}^*$  in the limit:

$$\mathbf{x}^* = \lim_{t \rightarrow \infty} \mathbf{x}(t) = \lim_{t \rightarrow \infty} \mathbf{W}^t\mathbf{x}(0) = \mathbf{v}_1\mathbf{x}(0),$$

where  $\mathbf{v}_1$  is the eigenvector corresponding to the leading eigenvalue of  $\mathbf{W}$ .

The theorem identifies the two conditions under which a group reaches a consensus (i.e., all the entries of  $\mathbf{x}^*$  become equal). First, the Markov process is capable of transitioning from any state to any other state with a non-zero probability, having no isolated subset of states (i.e., irreducibility). Second, the process does not exhibit a regular, repeating pattern in its transition (i.e., aperiodicity). Put simply, the group *always* reaches a consensus, so far as the influence network remains strongly connected<sup>33</sup>. Further, the entries of  $\mathbf{v}_1$  are equivalent to the eigenvector centralities of corresponding agents in the network<sup>32</sup>. Thus, the group consensus  $\mathbf{v}_1\mathbf{x}(0)$  is not necessarily the "middlemost" opinion that evenly reflects the diverse opinions initially held by individual members, and thus the errors in initial opinions cannot be effectively canceled out when aggregated. Instead, the group consensus is determined—or at least, highly influenced—by a few individuals who occupy the most central positions in the group, which necessarily results in a deviation from the collective decision formed by a simple aggregation and the true state of nature.

However, the introduction of selective exposure into opinion dynamics will prevent a group from reaching a global consensus. This is because selective exposure results in the fragmentation of the group into smaller subgroups (i.e., strongly connected components). In that case, the conditions for reaching a *global* consensus (i.e., irreducibility and aperiodicity) will no longer be met. Yet, each subgroup independently reaches a distinct *local* consensus. In this way, opinion diversity can be preserved at the global level.

### Agent-based models

We extend DeGroot's model by allowing endogenous changes in networks through selective exposure and then examine its impacts on group structure, opinion diversity, and the quality of collective decisions in comparison to naïve learning. To this end, we consider four agent-based models, including *Condorcet* and *Galton Models*, which simulate the settings depicted in Condorcet's theorem and Galton's experiment, respectively, and *Bimodal* and *Exponential Models*, which simulate initially polarized and skewed opinion distributions, respectively. The aspects shared by all these models are described first and followed by the details specific to each model, including



the distributions of initial opinions, the rules used to aggregate individual decisions into collective decisions, and the expected qualities of the collective decisions formed in the absence of social interactions, which are to be compared against the simulation results from each model.

#### Common settings

All four models consider a group of  $n$  agents. Each agent starts with its own opinion  $x_i$ . The initial opinion  $x_i$  is an *i.i.d.* random variable that follows a probability distribution specified in each model. This ensures that every agent begins with an independent opinion, and so the group starts with sufficient opinion diversity. Also, each agent is initially connected to  $k$  other agents that serve as its information sources. The connections between agents are directed and not necessarily reciprocal. The number of connections  $k$  is constant and finite, following the assumption that individuals have a limited capacity of interactions<sup>18,19,52,54</sup>. The initial connections between agents are formed uniformly at random and independent of their initial opinions, ensuring that agents are exposed to diverse opinions at the beginning. The communication network of agents  $G$  is a random directed graph with a fixed in-degree, and its adjacency matrix  $A_{ij}$  represents the directions of information flow from agents  $i$  to  $j$ . However, out-degrees follow a binomial distribution as in the Erdős–Rényi random graph model<sup>41</sup>. The set of agent  $i$ 's information sources, or its “information pool,” is denoted by  $G(i) = \{j | a_{ij} = 1\}$ .

In each round of simulations, agents attempt to minimize the discrepancies between their own opinions and those of their information sources in a random order, following the models used in previous studies<sup>30,42</sup>. *Discrepancy* is measured as the sum of differences between an agent's opinion and those of its information sources  $\sum_{j \in G(i)} |x_i - x_j|$  (i.e., the sum of absolute deviations). We consider two ways of reducing discrepancy. The first is *naïve learning*, by which an agent simply adopts the opinion  $x'_i$  such that it minimizes the discrepancies with its information sources:

$$x'_i = \operatorname{argmin}_{j \in G(i)} |x_i - x_j|.$$

whereas previous studies assume agents to adopt the “mean” opinion of others, we assume agents to adopt the “median” opinion in their information pool for several reasons. First, when the amount of discrepancy is measured as the sum of “absolute” deviations rather than “squared” deviations, the median minimizes discrepancy<sup>74</sup>. Second, the statistical behavior of sample median is asymptotically equivalent to that of sample mean, regardless of the original distribution from which a sample is drawn<sup>75</sup>. Third, when the distribution of initial opinions is skewed, the median is robust against outliers<sup>76,77</sup>, producing more stable results. Fourth, the “mean” opinion is obtained by the linear combination of others' opinions, which makes mathematical analysis convenient and allows for straightforward calculations and interpretation. However, using the “mean” opinion drastically increases the computational time for a simulation to reach its steady state. In theory, no simulation of DeGroot's model can reach its steady state within a finite amount of time. Even when an arbitrary stopping rule is applied, a simulation can approach “closely enough” to its steady state at best. On the other hand, when the “median” opinion is used, simulations can be completed within a reasonable amount of time, while generating practically identical and stable results.

Alternatively, agents can reduce discrepancy via *selective exposure*, by which an agent removes its connection to the information source whose opinion differs most from its own and adds a new connection to a randomly selected agent whose opinion is more congruent with its own than the removed one. In each round, each agent chooses either selective exposure with a probability of  $\beta$  or naïve learning with its complement of  $1 - \beta$ . A set of 51 evenly spaced values between 0 and 1 is used to quantify the parameter describing agents' propensity for selective exposure:  $\beta \in \{.00, .02, .04, \dots, 1.00\}$ . For each value of  $\beta$ , a total of 30,000 simulations are performed.

#### Model specific settings

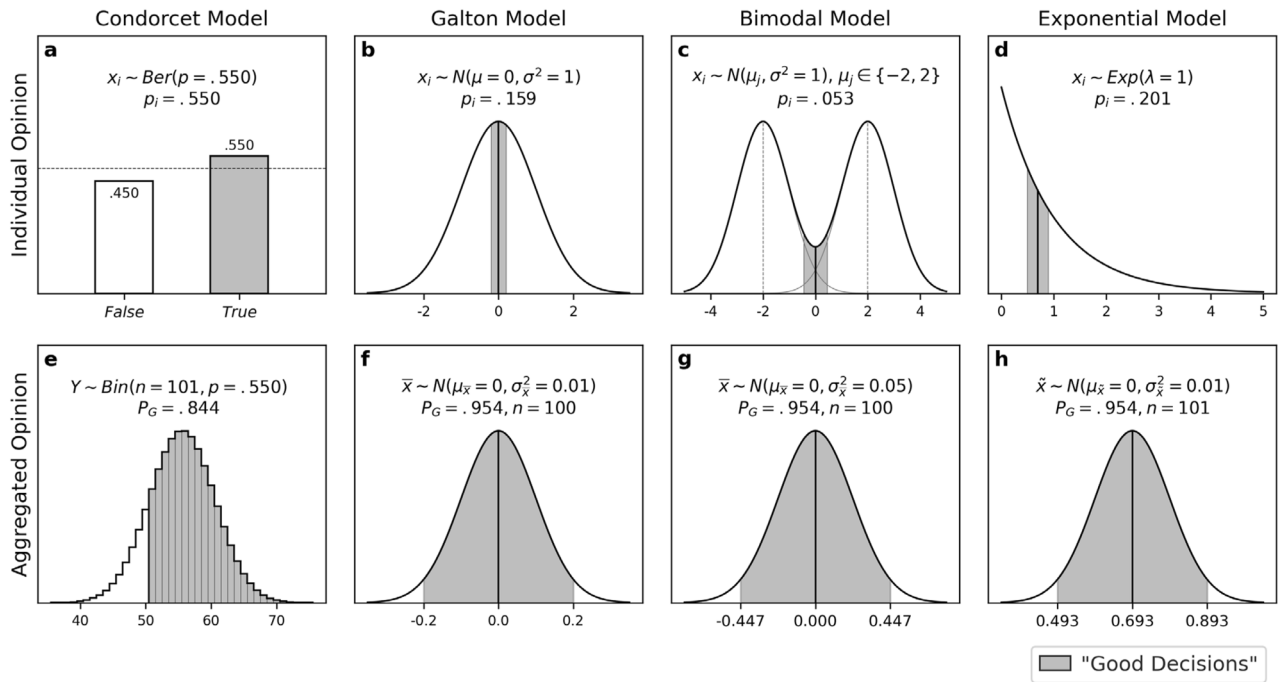
First, group sizes in the *Condorcet Model* are set as an odd number,  $n = 101$ , to avoid the need for tie-breaking rules. Adopting the assumption of Condorcet's theorem, the initial individual opinion  $x_i \in \{True, False\}$  for  $i = 1, 2, 3, \dots, n$  is an *i.i.d.* Bernoulli random variable, and individual competence is set as  $p = .55$  (Fig. 1a). When initial opinions are aggregated, the number of correct decisions  $\sum_i^n x_i$  in a group follows a binomial distribution  $y = \sum_i^n x_i \sim \operatorname{Bin}(n = 101, p = .55)$ . In that case, the probability of the group reaching the correct decision by the majority rule is defined as the sum of probability mass  $P(Y > n/2) = .844$  (Fig. 1e).

Second, group sizes in the *Galton Model* are set as  $n = 100$  and the initial individual opinions as an *i.i.d.* standard normal random variable  $x_i \sim N(\mu = 0, \sigma^2 = 1)$ , where the true state is  $\mu = 0$  (Fig. 1b). The collective decision  $\bar{x}$  is determined as the mean of individual opinions. Thus, collective decisions follow a normal distribution with mean  $\mu = 0$  and variance  $\sigma_{\bar{x}}^2 = 0.01$  by the central limit theorem (Fig. 1f):  $\bar{x} \sim N(\mu = 0, \sigma_{\bar{x}}^2 = 0.01)$ .

As an extension of the Galton Model, we consider the *Bimodal Model* in which initial individual opinions follow a mixture of two normal distributions, both sharing a common variance but having different means, such that the overall distribution has two peaks (Fig. 1c). This model is intended to reflect the polarized opinion distributions commonly observed in political contexts<sup>78,79</sup>. To generate the initial opinion  $x_i$ , a local mean  $\mu_j$  is chosen between  $-2$  and  $2$  with an equal probability, and then, a normal random variable with local mean  $\mu_j$  and common variance  $\sigma_j^2 = 1$  is generated:

$$\mu_j \sim \operatorname{Ber}(q) \text{ and } x_i \sim N(\mu_j, \sigma_j^2), \text{ where } q = .5, \mu_j \in \{-2, 2\}, \text{ and } \sigma_j^2 = 1.$$

The distribution of initial individual opinions is symmetric around the global mean  $\mu = 0$  ( $= q \sum_{j=1}^2 \mu_j$ ) and has two peaks at the local means. By the law of total variance, the variance of the distribution is  $\sigma^2 = 5$ :



**Figure 1.** The distributions of initial individual opinions (a through d) and aggregated initial opinions (e through h). The shades indicate “Good Decisions,” operationally defined as *True* decisions for the Condorcet model and as those that fall within a range of two standard errors from the true state of nature in the other models.

$$\sigma^2 = \sum_{j=1}^2 q(\mu_j - \mu)^2 + \sum_{j=1}^2 q\sigma_j^2.$$

Group size is set as  $n = 100$ , and collective decisions are represented by the mean of individual opinions as in the Galton model. Under the condition of independence, collective decisions are normally distributed with mean  $\mu = 0$  and variance  $\sigma^2 = 0.05$  by the central limit theorem (Fig. 1g):  $\bar{x} \sim N(\mu = 0, \sigma_{\bar{x}}^2 = 0.05)$ .

Lastly, the *Exponential Model* assumes that initial individual opinions are highly skewed, following an exponential distribution with rate parameter  $\lambda = 1$  (Fig. 1d). Unlike the previous models, collective decisions are represented by the *median* of individual opinions (i.e., median aggregation). Because the median is robust against outliers, it is the preferred measure of central tendency for skewed distributions<sup>76,77</sup>. Accordingly, the true state of nature is set as the median of the distribution,  $\tilde{\mu} = \ln 2 / \lambda = 0.693$ .

The statistical behavior of sample median  $\tilde{x}$  is known as follows<sup>75</sup>: When a sample of  $n$  random variables  $\{x_1, x_2, \dots, x_n\}$ , where  $n = 2m + 1$  and  $m$  is a positive integer, are independently and identically drawn from a probability distribution with a density function  $f(x)$ , the distribution of sample medians is approximately normal with mean  $\tilde{\mu}$  (i.e., the median of the parent distribution) and variance  $\frac{1}{8f(\tilde{\mu})^2 m}$  for large  $n$ . To utilize this statistical property of sample median, we set group size as  $n = 101$  (or  $m = 50$ ). In that case, collective decisions approximately follow a normal distribution with mean  $\tilde{\mu} = \ln 2 = 0.6931$  and variance 0.01, as  $f(\tilde{\mu}) = e^{-\ln 2} = 1/2$  and  $8f(\tilde{\mu})^2 m = 8 \cdot (1/2)^2 \cdot 50$  (Fig. 1h).

In addition to the results presented below, we also examined each of the four models under different parametric settings to establish the robustness of our results. These additional settings include different initial network structures (i.e., scale-free and regular networks), numbers of information sources ( $k = 3$  and  $k = 9$ ), and group sizes ( $n = 400$  for the *Galton* and *Bimodal Models* and  $n = 401$  for the *Condorcet* and *Exponential Models*). These results are presented in *Supplementary Materials* (Section S2), where they can be seen to reproduce qualitatively similar outcomes.

### Measures

Each simulation is terminated and considered to have arrived at its steady state when no further change occurs either in connections between agents or in their opinions. Individual opinions at steady states are considered as the final decisions that individual agents eventually arrive at, denoted by  $x_i^*$ , and then aggregated to form collective decisions, denoted by  $\bar{x}^*$  (or  $\tilde{x}^*$  for the Exponential Model).

In the *Condorcet Model*, individual performance  $p_i$  is measured as the proportion of agents whose final decisions are *True*. Similarly, collective performance  $P_G$  is measured by the proportion of groups whose majority decide *True* at steady states. Given that group size  $n = 101$ , it is expected that  $p_i = .55$  and  $P_G = P(\frac{1}{n} \sum_{i=1}^n x_i^* > .5) = .844$  in the absence of social influence.

To measure individual and collective performances in the other models that assume continuous opinions, we first operationalize “Good Decisions” and then measure individual performance  $p_i$  and collective performance  $P_G$  as the fractions of “good decisions” found at steady states. “Good decisions” are defined as those that fall within a range of two standard errors from the true state of nature, adopting the conventional decision criterion in statistical analysis. That is,

$$\text{Good Individual Decisions } H_i = \{x_i^* \mid |x_i^* - \mu| < 2\sigma_{\bar{x}}\} \text{ and}$$

$$\text{Good Collective Decisions } H_G = \{\bar{x}^* \mid |\bar{x}^* - \mu| < 2\sigma_{\bar{x}}\},$$

where the standard error is  $\sigma_{\bar{x}} = \sigma/\sqrt{n}$ , and  $\mu$  is the true state ( $\tilde{\mu}$  for the Exponential Model) specified by a model. Then,

$$\text{Individual Performance } p_i = \frac{|H_i|}{n} \text{ and}$$

$$\text{Collective Performance } P_G = \frac{|H_G|}{N},$$

where  $n$  is the group size, and  $N$  is the total number of simulations for each value of  $\beta$ .

Previous studies measure collective performance, as alluded to in references such as “wise society”<sup>33</sup>, “wise crowd”<sup>30</sup> and “(mis)information”<sup>34</sup>, by the (squared) deviation of a collective decision from the true state of nature and show that the deviation approaches zero (i.e., achieving collective intelligence) in the absence of “excessively” influential agents as group size approaches infinity. However, it should be noted that the deviation is inversely proportional to the square root of group size  $\sqrt{n}$ , and therefore the infinitely large size of groups is sufficient for the deviation to converge to zero, as indicated by the central limit theorem. Also, the sublinear relationship between the deviation and group size demonstrates the diminishing marginal improvement of collective performance, as pointed out by Kao and Couzin<sup>39</sup>. Therefore, when the effects of group size are not properly controlled for, the deviation alone tends to overestimate collective performance and thereby underestimate the negative effects of influential agents. For this reason, we include standard error  $\sigma_{\bar{x}}$  to define the range of “good decisions,” which allows for direct and more precise comparison between collective performances under different settings across models. This is because the ranges of “good decisions” are all standardized. In the absence of social influence (i.e., a simple aggregation of initial opinions), a group is expected to make a “good” collective decision with a probability of 0.954 [=  $P(|x_i^* - \mu| < 2\sigma_{\bar{x}})$ ], since  $(\bar{x}^* - \mu)/\sigma_{\bar{x}}$  follows the standard normal distribution  $N(0, 1)$ , regardless of the distributions of initial opinions. Also, in the case of an infinitely large group, our definition becomes equivalent to those in previous studies  $\{\bar{x}^* \mid |\bar{x}^* - \mu| = 0\}$ , since  $\sigma_{\bar{x}}$  converges to zero.

In addition, we measure the *opinion diversity* in groups  $\sigma_x$ , the *size of the giant components*  $|C_1|$ , and the *modularity*  $Q$  of networks at steady states to examine the impacts of local dynamics on the global structure. Opinion diversity  $\sigma_x$  is measured by the standard deviation of individual opinions at steady states. The size of the giant component  $|C_1|$  is measured by the proportion of nodes that belong to the largest strongly connected component of a communication network  $G$ , within which every node is reachable from every other node. Modularity  $Q$  is measured as

$$Q = \sum_d \left[ \frac{L_d}{m} - \left( \frac{k_d^{\text{in}} k_d^{\text{out}}}{2m} \right)^2 \right],$$

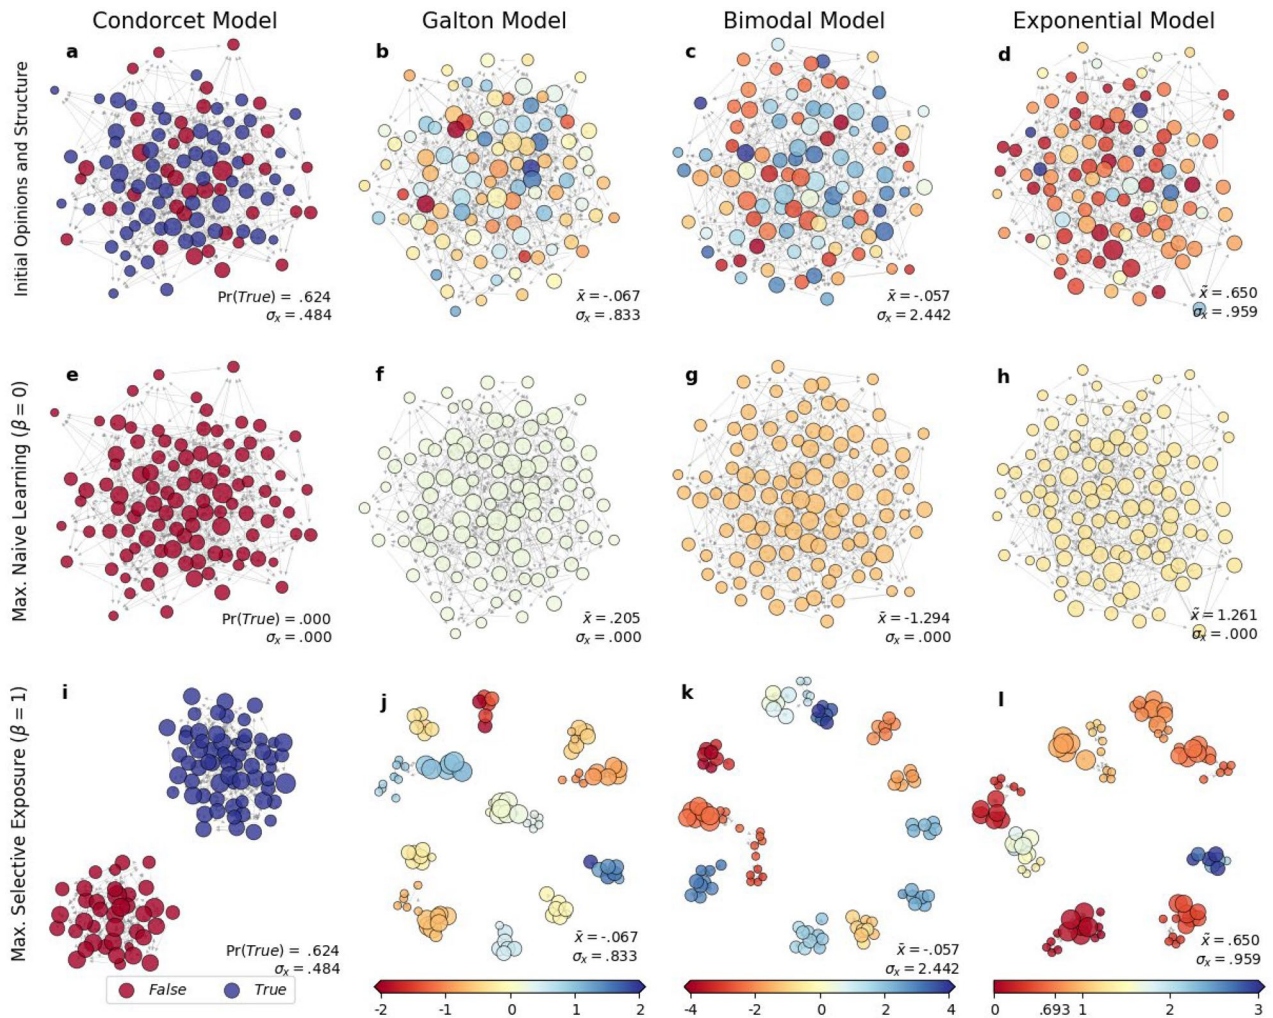
where  $d$  represents the dichotomized categories of individual opinions (i.e., *True* or *False*; *above* or *below* the true state),  $m$  is the total number of edges,  $L_d$  is the total number of edges within category  $d$ , and  $k_d^{\text{in}}$  and  $k_d^{\text{out}}$  are the sums of in-degrees and out-degrees of the nodes in category  $d$ , respectively<sup>80</sup>. Both the size of the giant components  $|C_1|$  and modularity  $Q$  are used to measure the degree to which the communication network is fragmented. Low values of  $|C_1|$  and high values of  $Q$  indicate the presence of echo chambers within groups.

## Results

### Sampled results: replication of previous findings

Figure 2 visualizes the sampled results from numerical experiments and compares opinion distributions and group structures at the initial states with those at steady states under the maximal propensities for naïve learning ( $\beta = 0$ ) and selective exposure ( $\beta = 1$ ). In each plot, nodes represent individual agents, and their colors and sizes indicate the opinions and eigenvector centralities of the corresponding agents, respectively. Nodes are placed closely to one another if they are connected.

Across all four models, groups initially contain diverse opinions, and agents with different opinions are connected to one another (Figs. 2a–d). However, the two different dynamics—*naïve learning* and *selective exposure*—yield drastically different outcomes. Under the maximal propensity for naïve learning ( $\beta = 0$ , Figs. 2e–h), all agents end up adopting the same opinion, and thus, groups reach a consensus, as DeGroot’s model predicts. As a result, the opinion diversity at the global level vanishes completely ( $\sigma_x = 0$ ), demonstrating the diversity-reducing effects of social influence. Under the maximal propensity for selective exposure ( $\beta = 1$ , Figs. 2i–l), agents alter their local networks without adjusting their opinions at all. As a result, agents with similar opinions end up being densely clustered and form exclusive subgroups between which no edges are found. Opinions within each subgroup are nearly identical, indicating the loss of opinion diversity within subgroups. However,



**Figure 2.** Sampled results from the simulations of the four models: Initial opinion distributions and network structure of groups (a through d) and those at steady states under the maximal propensities for *naïve learning* ( $\beta = 0$ , e through h) and *selective exposure* ( $\beta = 1$ , i through l). Node size and color indicate the eigenvector centralities and the opinions of corresponding agents, respectively.

each subgroup reaches its own unique consensus distinguished from those reached by others, and therefore, the global-level opinion diversity  $\sigma_x$  is perfectly preserved. These results confirm the findings of previous studies, validating our approaches.

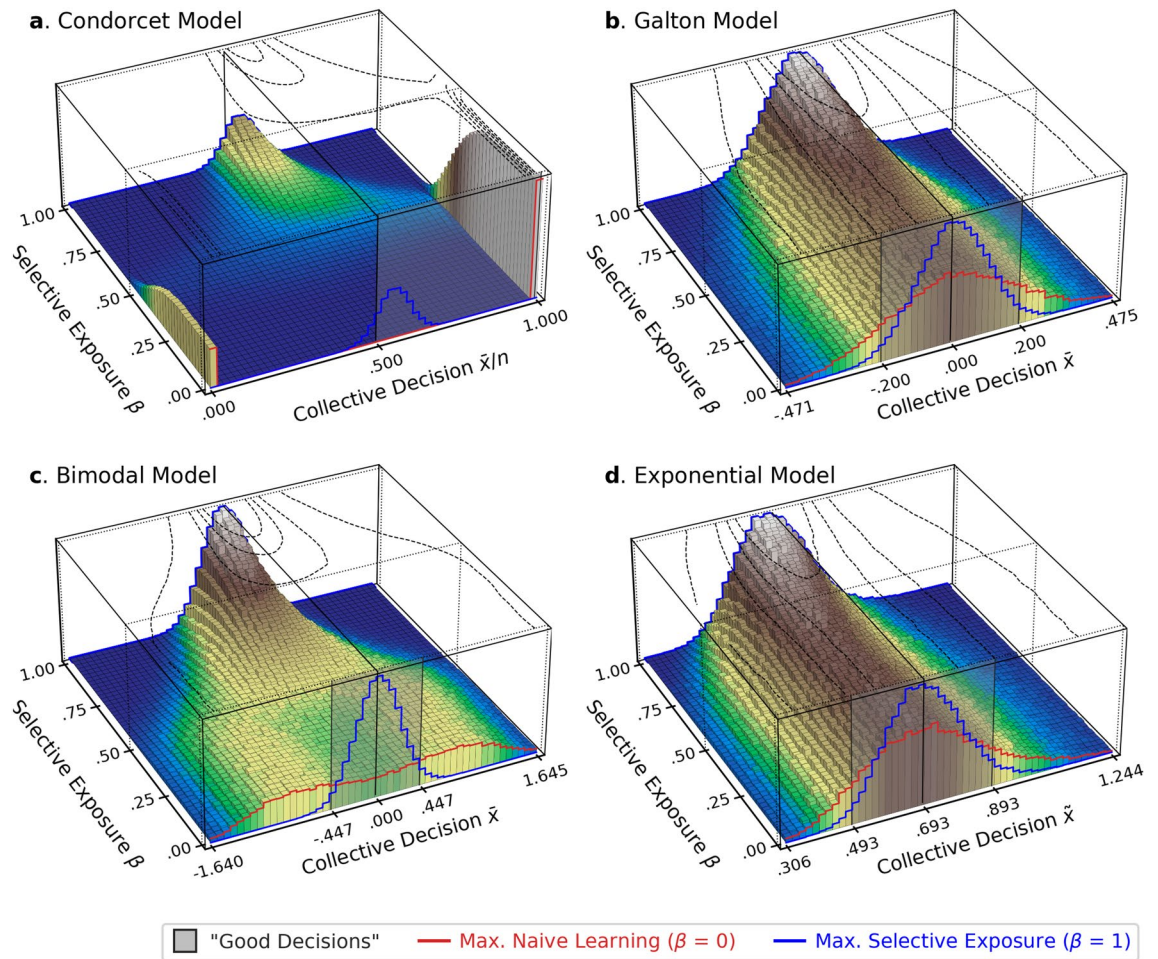
### The effects of selective exposure

Figure 3 presents 3D histograms in which the heights of bars represent the relative frequencies of collective decisions obtained from 30,000 simulations for each of the intermediate levels of propensity for selective exposure  $\beta \in \{.00, .02, .04, \dots, 1.00\}$  across the four models.

Across all the models, the distributions of collective decisions formed under the maximal propensity for *selective exposure* ( $\beta = 1$ ; the blue histograms projected on the front in Fig. 3) are nearly identical to those made by a simple aggregation of initial opinions (Figs. 1e–h). In terms of collective performance, 84.4% of groups arrive at the correct decision in the Condorcet models, and 95.4% of groups make “good” decisions in the other models, suggesting that collective performance is maintained at a level as high as that achieved by a simple aggregation of initial opinions. On the other hand, as *naïve learning* becomes more dominant, the distributions become more dispersed with longer and thicker tails, indicating that collective decisions are more likely to deviate from the true state. These deviations are most salient in the *Condorcet model* due to the binary nature of the opinions it models (i.e., *True* or *False*). When *naïve learning* is dominant over *selective exposure* ( $\beta < .50$ ), groups produce only two collective outcomes: Either all agents make the correct decision, or all make the wrong decision. Effectively, each group is constrained to behave like a single individual.

Although their results are not as extreme as those observed in the Condorcet Model, similar patterns are observed in the other models that assume continuous opinions. In the *Bimodal Model*, when *naïve learning* is dominant, the distribution of collective decisions resembles that of initial individual opinions characterized by the two peaks, implying that collective decisions are effectively determined by a few agents who happen to occupy





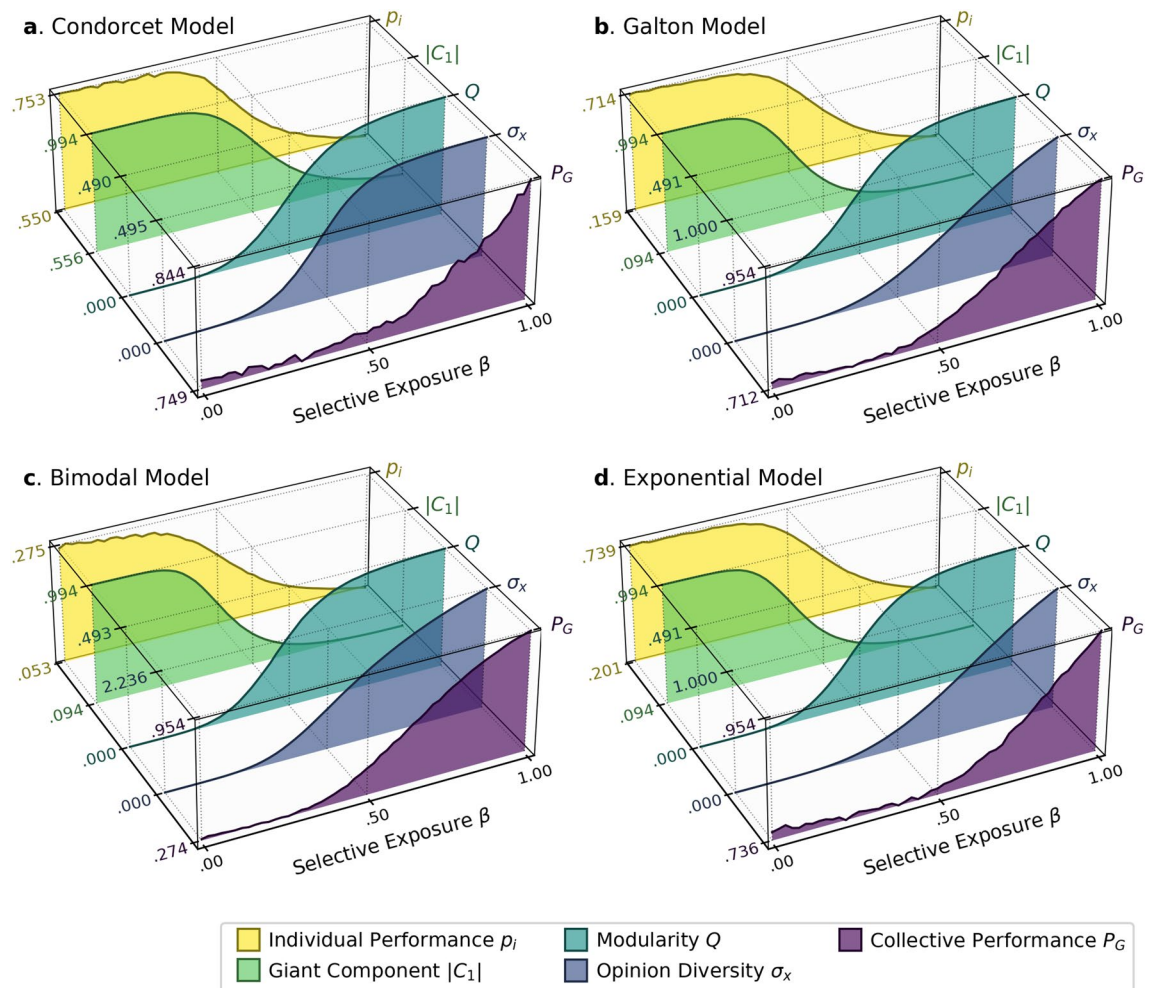
**Figure 3.** 3D histograms of the collective decisions formed at different levels of selective exposure across the four models. The red and blue histograms projected on the front represent the distributions of collective decisions reached under the maximal propensities for *naïve learning* ( $\beta = 0$ ) and *selective exposure* ( $\beta = 1$ ), respectively.

central positions in the communication network and so exert more influence on the collective decisions. In the *Exponential Model*, collective decisions reached under the maximal propensity for *naïve learning* are skewed to the right, reflecting the initial distribution of individual opinions, even though the method of median aggregation is robust against outliers. This also illustrates that collective decisions tend to be determined by a few central agents when *naïve learning* is dominant.

Figure 4 summarizes individual performance  $p_i$ , giant component size  $|C_1|$ , modularity  $Q$ , opinion diversity  $\sigma_x$ , and collective performance  $P_G$  at different levels of *selective exposure*, each of which displays nearly identical patterns across all the four models. First, selective exposure sharply decreases the size of the giant component  $|C_1|$  and increases modularity  $Q$ , which together suggest that selective exposure leads to the emergence of exclusive and disconnected subgroups of like-minded individuals. Second, selective exposure preserves opinion diversity  $\sigma_x$  at the global level, suggesting that the diversity-reducing effect of social influence has been effectively confined within subgroups. Most importantly, selective exposure allows groups to maintain their collective performance at a level as high as that achieved by a simple aggregation of initial opinions, as already shown in Fig. 3.

One important observation is that whereas individual performance  $p_i$  remains unchanged under strong propensities for *selective exposure*, it is substantially improved under strong propensities for *naïve learning* across all the models (the yellow layers in Fig. 4), which is consistent with previous findings<sup>33,34</sup>. For instance, when agents naïvely adopt the majority opinion in their information pools, 75.3% of agents end up making the correct decision at the steady states in the Condorcet Model ( $p_i = .753$  when  $\beta = 0$ ). This proportion is substantially higher than the initial individual competence  $p = .550$ , which can be explained as follows. A naïve learner adopts the “majority” or “median” opinion of its information pools. Therefore, the updated opinion is effectively a “collective” decision made by simply aggregating the opinions of its information sources, as in Condorcet’s theorem and Galton’s experiment. Even though the size of this “group” is small ( $k = 5$  in our simulation settings), the updated opinion should be statistically more accurate than any of the individual opinions held by one’s information sources. In that sense, individual opinions are expected to progressively improve through each round of naïve learning.

Contrary to the findings of previous studies<sup>33,34</sup>, however, collective decisions formed through naïve learning significantly deviate from the true state *despite* the improvement in individual decisions (the purple layers



**Figure 4.** The effects of selective exposure on individual performance  $p_i$ , giant component size  $|C_1|$ , modularity  $Q$ , opinion diversity  $\sigma_x$ , and collective performance  $P_G$ .

in Fig. 4). This cannot be attributed to the presence of “excessively” influential agents, which is unlikely in our models because all the connections between agents are formed *uniformly* at random. Instead, the current results suggest that the negative effects of (even “moderately”) influential agents are much greater than predicted by previous studies (see also Section S2.3 in *Supplementary Materials*). Further, the lower collective accuracy, despite improved individual accuracy, suggests the importance of opinion diversity for collective decision-making relative to individual competence. Naïve learning, operating as collective decision-making at a small and local scale, progressively improves individual opinions, but the rate of improvement drastically decreases because the opinions within information pools become homogenous over time at a faster rate than that of improvement. At steady states, the aggregation of opinions in one’s information pool cannot yield any improvements in collective accuracy, simply because not only the information sources but also all other agents adopt a common opinion. That is, the whole group effectively behaves as if it were a single “moderately competent” individual. According to our simulation results, for example, individual agents make the correct decisions with a probability of 0.753 under the maximal propensity for naïve learning in the Condorcet Model ( $p_i = .753$ ), while collectively making the correct decisions with a probability of 0.749 (Fig. 4a). The minor gap between the two probabilities is due to the failure of groups to remain strongly connected and thus to reach consensus [i.e.,  $P(|C_1| < 1) = 1 - .994$ ]. Either probability is significantly lower than the probability of a group making the correct decision by a simple aggregation of initial opinions (0.844 in Fig. 1e). Not surprisingly, similar patterns are observed in the other models.

The results presented above are all successfully reproduced in extended numerical experiments performed over a wider range of simulation settings, confirming the robustness of the results above (see *Supplementary Materials* Section S2 for more details). In addition, the results from the extended numerical experiments allow us to identify the conditions under which naïve learning can also function to maintain the quality of collective decisions despite the utter absence of opinion diversity within groups. First, when each agent has access to a *large pool of information*, both individual and collective decisions improve (Section S2.2). Ideally, if every agent can be exposed to and has a chance to learn from everyone else’s opinions (i.e.,  $k = n - 1$ ; complete networks), collective decisions under naïve learning are expected to be as good as those made under the maximal propensity for selective exposure or in the absence of social influence. Second, the group structure should be perfectly



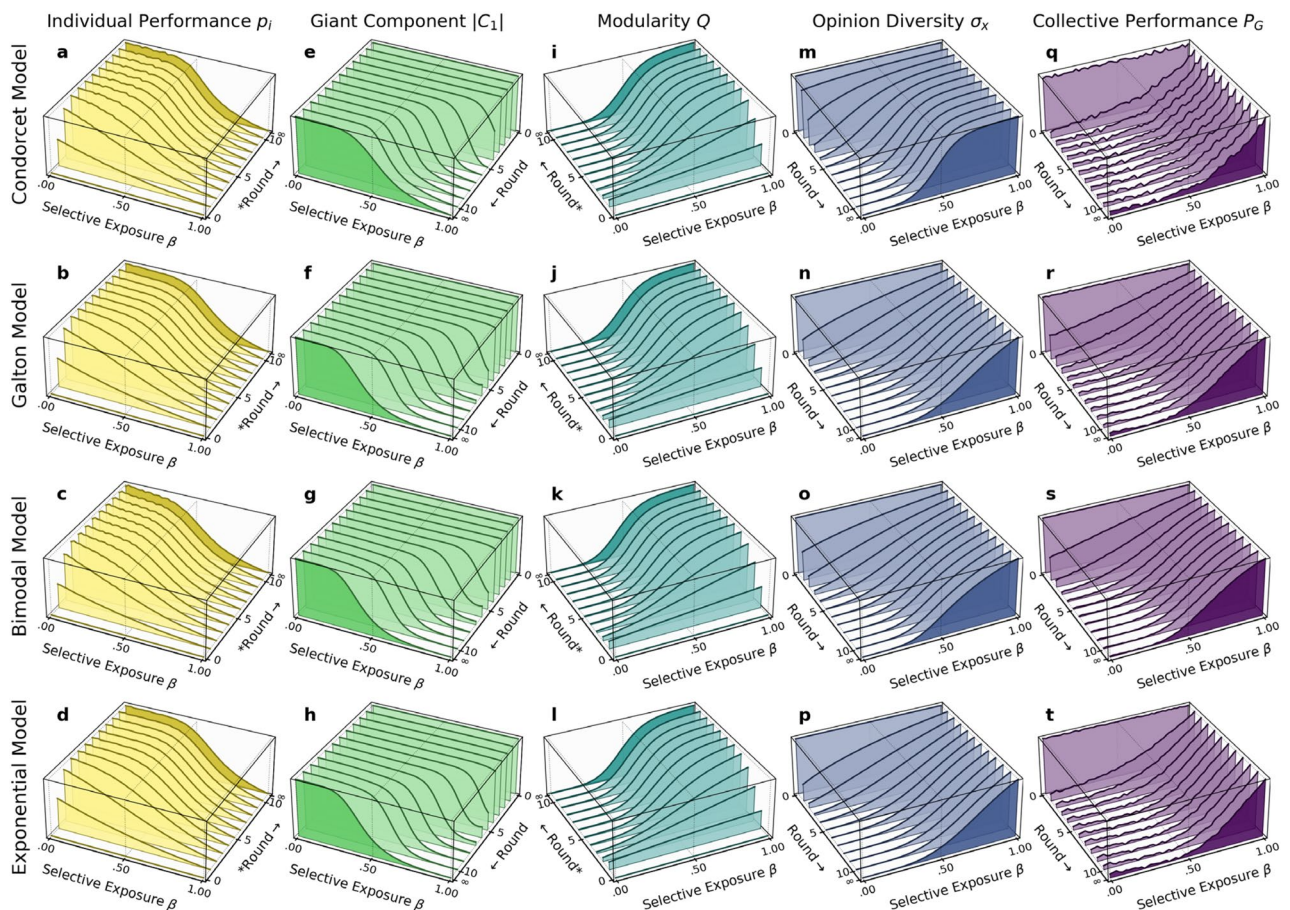
*decentralized* (i.e., regular networks). That is, every group member must exert exactly an equal amount of influence on everyone else's decision, and there must be neither "leaders" nor "followers" in the group (Section S2.3). Although these conditions can be implemented in our agent-based models, they are hardly feasible or expected to be met in real-world collective decision-making.

### Collective decision-making under time constraints

In real-world decision-making, both individual and collective decisions are typically subject to time constraints. Rather than having unlimited time to deliberate, individuals and groups are often forced to make decisions *before* arriving at a steady state. In that sense, transient opinion dynamics may be more relevant. From this perspective, we examine the effects of selective exposure and naïve learning during the early stages of dynamics (Fig. 5).

Figure 5 visualizes the changes in performance and group structure for the first ten rounds, during which most major changes occur. Even after the very first round, selective exposure (as well as naïve learning) already shows non-trivial impacts on both performance and group structure. Subsequently, all the measures remain nearly the same until steady states are reached (the darkened layer in each plot). These results suggest that the effects of selective exposure become observable almost immediately, further reinforcing the robustness of our findings. Even if decisions must be made over truncated time scales, selective exposure still affects both individual and collective performances and group structure in a manner consistent with the steady-state results reported above. It should be noted that these "immediate" effects observed in the early stage of simulations are "relative" to the time scale, and thus they can be seen as an artifact of the parameter settings specified in our numerical experiments.

Furthermore, the changes observed over consecutive rounds help explicate the causal relationships among the local behavior of agents, global structural properties, and individual and collective performances, allowing for a better understanding of the steady-state results reported above. Naïve learning substantially improves individual performance for the first few rounds (the first column in Fig. 5), as agents update their opinions by aggregating diverse opinions in their information pools (i.e., collective decision-making at small and local scales), but the rate of improvement sharply declines. This is because the repetition of naïve learning eradicates opinion diversity (the fourth column in Fig. 5). On the other hand, selective exposure fragments groups into subgroups, which become more homogeneous (the third column in Fig. 5) and smaller (the second column in Fig. 5).



**Figure 5.** Changes in performance and group structure for the first ten rounds at different levels of propensity of selective exposure. The darkened layer in each plot represents the corresponding measures at steady states. *Note:* "Rounds" for individual performance  $p_i$  and modularity  $Q$  are arranged in reverse order for a clearer visual presentation.

However, selective exposure leads to no changes in opinion diversity at the global level (the fourth column in Fig. 5), thereby maintaining the quality of collective decisions over time (the fifth column in Fig. 5). It is worth noting that, under strong propensities for selective exposure, the decrease in giant component size  $|C_1|$  is delayed compared to the changes in other measures. This is because, during the early stage, the whole network remains robust against the modification of local structures due to the high initial connectivity among agents ( $k = 5$ )<sup>81</sup>.

## Discussion

In this study, we have examined the impacts of selective exposure on the quality of collective decisions through an extension of DeGroot's model<sup>26</sup> by allowing for endogenous changes in the communication networks of agents. A series of numerical experiments suggest that our models successfully replicate previous findings. The results confirm that naïve learning leads to the convergence of opinions in well-connected groups, and further, collective decisions formed under naïve learning are more accurate than initial individual opinions<sup>30,33,34</sup>. Also, our results are consistent with previous findings, showing that selective exposure leads to the fragmentation of a group into exclusive and homogenous subgroups<sup>53,54,56,57</sup>, each of which converges to a unique consensus<sup>27</sup>. These successful replications of previous findings by our models validate our approach and indicate that the results we present are not in contradiction with the findings of previous studies.

In addition, the present study offers important insights into the understanding of the role selective exposure plays in collective decision-making in comparison to naïve learning, complementing previous findings. The major findings of the current study can be summarized as follows. A group consensus reached through naïve learning is more accurate than initial individual opinions but significantly deviates from the true state of nature, even in the absence of “excessively” influential individuals. Naïve learning progressively improves individual opinions, because the learning process is conceptually equivalent to a collective decision-making process at a small and local scale. However, the rate of improvement decreases, as naïve learning eradicates opinion diversity at both local and global levels at a faster rate than that of improvement. This suggests that the improvement in individual opinion is achieved at the cost of opinion diversity, and that learning from others' opinions can be beneficial only if those opinions are sufficiently diverse. On the other hand, the diversity-reducing effects of social influence are effectively confined within exclusive subgroups formed by selective exposure, while global-level opinion diversity is well preserved. As a result, the collective decisions made under strong propensities for selective exposure are as accurate as a simple aggregation of independent opinions initially held by individual members. This suggests that selective exposure can function to manage the diversity-reducing effects of social influence in highly interactive communication environments, providing a plausible answer to our motivating question, “How can groups perform well in decision-making *despite* social influence?”

However, it should be noted that our models examined the effects of selective exposure only in terms of the *accuracy* of collective decisions without considering any potential *costs* associated with coordinating and integrating individuals with diverse or even opposing opinions as well as other negative consequences of structural fragmentation, such as intergroup hostility and conflicts. Perhaps, those costs might exceed the benefits of producing more accurate and unbiased decisions in real-world contexts. Therefore, the current findings should not be interpreted to imply that selective exposure is necessarily superior to other alternative means of reducing disagreement among group members, utilizing local information, or managing the diversity-reducing effects of social influence. Likewise, it is equally undesirable to blindly believe that learning from diverse views or the equivalent would be a “panacea” for a wide range of social and political problems because it leads to poorer and more biased collective decisions than groups could have possibly produced otherwise. This trade-off between the “melting pot” ideal of harmonious consensus and the “salad bowl” concept of preserving diversity deserves further examination. Therefore, future research is required to present more balanced perspectives and theoretical frameworks on opinion dynamics and collective decision-making.

It is important to acknowledge the limitations of our models and discuss the directions of future research to extend the current findings. First, although naïve learning has been widely adopted as a model of opinion dynamics, a problem with its fundamental assumption needs to be carefully considered. Naïve learning is a rational opinion update rule when agents reside in a sparse network, as in such cases, their neighbors' opinions can be assumed to be effectively independent. However, when the connectivity among neighbors exceeds a certain level, this independence assumption cannot be held. Thus, the failure to recognize and adjust for the excessive connectivity leads to suboptimal social responses and poorer decisions<sup>5</sup>. Nonetheless, the adoption of naïve learning can be defended by assuming that agents lack the cognitive ability to handle the complexity inherent in the densely interconnected networks, they do not have access to the information about the correlation structure among neighbors, and/or they are overly accustomed to situations in which the connectivity among neighbors can be ignored. Alternatively, other formulations of opinion updating can be considered. For example, agents can be assumed to be (Bayesian) rational decision-makers who update their (private) prior opinions based on the observed decisions previously made by others (i.e., sequential decision-making problems). In that case, a high connectivity of network structures leads to the phenomenon known as *information cascades*, where the rational decision-making process is simplified to merely tallying each option chosen by others, and the decision made by a few initial agents takes over the entire group<sup>82–87</sup>. In this scenario, we expect that selective exposure might still function by isolating cascades within the subgroups from which they originate. Also, when selective exposure is viewed as a response to normative social influence by adjusting one's local structure rather than opinions, its impacts need to be compared to those of alternative ways by which agents avoid disagreement with neighbors without adjusting their own opinions, such as remaining silent<sup>88,89</sup> or even falsely stating their opinions<sup>27,30</sup>. Buechel et al.,<sup>30</sup> for instance, consider the possibility that agents intentionally misrepresent their opinions either by conforming or counter-conforming with their neighbors. They show that the misrepresentation of opinions does not necessarily undermine collective accuracy but may even enhance it. According to this



framework, “dishonest” agents can uphold their initial opinions despite the normative social influence of their information sources, specifically to the extent that their true opinions differ from their stated opinions. Accordingly, the aggregate of these differences can be seen as the total amount of preserved diversity at the global level. However, the stated opinion of an agent can still influence the opinion formation of other agents who rely on the agent as their information sources. That is, normative social influence is simply “passed along” to other agents instead of being “stopped.” Therefore, it would be interesting to see how the introduction of “dishonest” agents affects the current results.

Second, our models assume that agents initially hold *independent* opinions, the aggregation of which is statistically equal to the true state of nature. However, this independence assumption is not applicable to real-world groups, because individual opinions tend to be correlated with each other even before any social interaction starts. Previous studies show that collective decisions made by a simple aggregation are more accurate when the *average* correlation among agents’ opinions is *negative* than when it is *positive*<sup>90,91</sup>. Although it is impossible for every agent to be mutually contrarian to each other within a group of three or more agents, it is still conceivable that *some* agents’ opinions are *perfectly* negatively correlated with some others, or that *every* agent’s opinion is *moderately* negatively correlated with every other. More specifically, the lower bound of the *average* correlation among  $n$  agents’ opinions is  $\bar{\rho} \geq -1/(n - 1)$ <sup>90</sup>. Further, when agents’ opinions are maximally negatively correlated, the generalized variance of the joint distribution of  $n$  opinions (i.e., multivariate dispersion) is accordingly maximized, maintaining opinion diversity within groups. Therefore, collective accuracy can be significantly improved by incentivizing individuals to be in the minority<sup>6</sup>. Further, whereas our model assumes no changes in opinions *after* selective exposure, the reinforcement model<sup>55,92</sup>—a framework proposed to account for the personalization-polarization hypothesis<sup>93</sup>—predicts that individuals tend to develop more extreme views after interacting with like-minded others. If that is the case, selective exposure widens the gaps between opinions that are initially leaning toward different ends, and the correlation between them grows in the negative direction, *enhancing* the quality of collective decisions rather than merely maintaining it. Perhaps, a *positive* correlation among individual opinions is more common to real-world groups, even though their effects on collective accuracy become negligible as group size increases<sup>94</sup>. In particular, the *external* factors of correlated opinions (i.e., “induced homophily” as opposed to preference-based “choice homophily”<sup>95</sup>), such as spatial constraints<sup>39</sup> and the reliance on common information sources<sup>45</sup>, should be duly incorporated into our models. One possible way is to include those factors as a set of pseudo-agents with no incoming edges but with a disproportionately large number of out-going edges, forming a bipartite network in which one mode is composed of the pseudo-agents and the other consists of the “ordinary” agents. These relatively straightforward extensions could help clarify how these additional factors affect the causal path from selective exposure to the quality of collective decisions through group structure and opinion diversity.

## Data availability

The datasets generated by numerical experiments and Python scripts used for simulations are available from the corresponding author on reasonable request.

Received: 10 October 2023; Accepted: 12 March 2024

Published online: 17 March 2024

## References

- Davis-Stober, C. P., Budescu, D. V., Dana, J. & Broomell, S. B. When is a crowd wise?. *Decision* **1**, 79–101 (2014).
- Lu, J., Liu, S., Cui, L., Pan, L. & Wu, L. Crowd wisdom drives intelligent manufacturing. *Int. J. Crowd Sci.* **1**, 39–47 (2017).
- Keck, S. & Tang, W. Enhancing the wisdom of the crowd with cognitive-process diversity: The benefits of aggregating intuitive and analytical judgments. *Psychol. Sci.* **31**, 1272–1282 (2020).
- Sobel, J. Economists’ models of learning. *J. Econ. Theory* **94**, 241–261 (2000).
- Mann, R. P. Collective decision-making under changing social environments among agents adapted to sparse connectivity. *Collect. Intell.* **1**, 263391372211213 (2022).
- Mann, R. P. & Helbing, D. Optimal incentives for collective intelligence. *Proc. Natl. Acad. Sci.* **114**, 5077–5082 (2017).
- Mann, R. P. & Woolley-Meza, O. Maintaining intellectual diversity in data science. *Data Sci.* **1**, 85–94 (2017).
- Hong, L., Page, S. E. & Riolo, M. Incentives, information, and emergent collective accuracy. *Manag. Decis. Econ.* **33**, 323–334 (2012).
- Asch, S. E. Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychol. Monogr. Gen. Appl.* **70**, 1–70 (1956).
- Deutsch, M. & Gerard, H. B. A study of normative and informational social influences upon individual judgment. *J. Abnorm. Soc. Psychol.* **51**, 629–636 (1955).
- Anderson, L. R. & Holt, C. A. Information cascades in the laboratory. *Am. Econ. Rev.* **87**, 847–862 (1997).
- Lorenz, J., Rauhut, H., Schweitzer, F. & Helbing, D. How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci.* **108**, 9020–9025 (2011).
- Simoiu, C., Sumanth, C., Mysore, A. & Goel, S. Studying the “wisdom of crowds” at scale. *Proc. AAAI Conf. Hum. Comput. Crowdsourcing* **7**, 171–179 (2019).
- Frey, V. & van de Rijt, A. Social influence undermines the wisdom of the crowd in sequential decision making. *Manag. Sci.* **67**, 4273–4286 (2021).
- Brabham, D. C. Moving the crowd at threadless: Motivations for participation in a crowdsourcing application. *Inf. Commun. Soc.* **13**, 1122–1145 (2010).
- Hill, S. & Ready-Campbell, N. Expert stock picker: The wisdom of (experts in) crowds. *Int. J. Electron. Commer.* **15**, 73–102 (2011).
- Turiel, J., Fernandez-Reyes, D. & Aste, T. Wisdom of crowds detects covid-19 severity ahead of officially available data. *Sci. Rep.* **11**, 13678 (2021).
- Baumann, L. A model of weighted network formation. *Theor. Econ.* **16**, 1–23 (2021).
- Jiménez-Martínez, A. & Melguizo-López, I. Making friends: The role of assortative interests and capacity constraints. *J. Econ. Behav. Organ.* **203**, 431–465 (2022).

20. Bazazi, S., von Zimmermann, J., Bahrami, B. & Richardson, D. Self-serving incentives impair collective decisions by increasing conformity. *PLOS ONE* **14**, e0224725 (2019).
21. Almaatouq, A. *et al.* Adaptive social networks promote the wisdom of crowds. *Proc. Natl. Acad. Sci.* **117**, 11379–11386 (2020).
22. van de Rijt, A. Self-correcting dynamics in social influence processes. *Am. J. Sociol.* **124**, 1468–1495 (2019).
23. Condorcet, N. D. *Essai Sur l'application de l'analyse à La Probabilité Des Décisions Rendues à La Pluralité Des Voix* (Edward Elgar Publishing, 1994).
24. Galton, F. Vox populi. *Nature* **75**, 450–451 (1907).
25. Galton, F. The ballot-box. *Nature* **75**, 509–510 (1907).
26. DeGroot, M. H. Reaching a consensus. *J. Am. Stat. Assoc.* **69**, 118–121 (1974).
27. Anufriev, M., Borissov, K. & Pakhnin, M. Dissonance minimization and conversation in social networks. *SSRN Scholarly Paper* <https://doi.org/10.2139/ssrn.4213302> (2023).
28. Zafar, B. An experimental investigation of why individuals conform. *Eur. Econ. Rev.* **55**, 774–798 (2011).
29. Levy, G. & Razin, R. Information diffusion in networks with the Bayesian peer influence heuristic. *Games Econ. Behav.* **109**, 262–270 (2018).
30. Buechel, B., Hellmann, T. & Klößner, S. Opinion dynamics and wisdom under conformity. *J. Econ. Dyn. Control* **52**, 240–257 (2015).
31. Berg, S. & Paroush, J. Collective decision making in hierarchies. *Math. Soc. Sci.* **35**, 233–244 (1998).
32. Bonacich, P. Power and centrality: A family of measures. *Am. J. Sociol.* **92**, 1170 (1987).
33. Golub, B. & Jackson, M. O. Naïve learning in social networks and the wisdom of crowds. *Am. Econ. J. Microecon.* **2**, 112–149 (2010).
34. Acemoglu, D., Ozdaglar, A. & ParandehGheibi, A. Spread of (mis)information in social networks. *Games Econ. Behav.* **70**, 194–227 (2010).
35. Clauset, A., Shalizi, C. R. & Newman, M. E. J. Power-law distributions in empirical data. *SIAM Rev.* **51**, 661–703 (2009).
36. Simon, H. A. On a class of skew distribution functions. *Biometrika* **42**, 425–440 (1955).
37. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
38. Rauhut, H. & Lorenz, J. The wisdom of crowds in one mind: How individuals can simulate the knowledge of diverse societies to reach better decisions. *J. Math. Psychol.* **55**, 191–197 (2011).
39. Kao, A. B. & Couzin, I. D. Modular structure within groups causes information loss but can improve decision accuracy. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **374**, 20180378 (2019).
40. Zafeiris, A. & Vicsek, T. Group performance is maximized by hierarchical competence distribution. *nature.com.remotexs.ntu.edu.sgmunications* **4**, 1–8 (2013).
41. Erdős, P. & Rényi, A. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.* 17–61 (1960).
42. Arifovic, J., Eaton, B. C. & Walker, G. The coevolution of beliefs and networks. *J. Econ. Behav. Organ.* **120**, 46–63 (2015).
43. Smith, S. M., Fabrigar, L. R. & Norris, M. E. Reflecting on six decades of selective exposure research: Progress, challenges, and opportunities. *Soc. Personal. Psychol. Compass* **2**, 464–493 (2008).
44. Katz, E., Blumler, J. G. & Gurevitch, M. Uses and gratifications research. *Public Opin. Q.* **37**, 509–523 (1973).
45. Lazarsfeld, P. F., Berelson, B. & Gaudet, H. *The People's choice: How the voter makes up his mind in a presidential campaign* (Columbia University Press, 1968).
46. Kunda, Z. The case for motivated reasoning. *Psychol. Bull.* **108**, 480–498 (1990).
47. Nickerson, R. S. Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* **2**, 175–220 (1998).
48. Heider, F. *The psychology of interpersonal relations* (Wiley, 1958).
49. Cartwright, D. & Harary, F. Structural balance: A generalization of Heider's theory. *Psychol. Rev.* **63**, 277–293 (1956).
50. Davis, J. A. Clustering and structural balance in graphs. *Hum. Relat.* **20**, 181–187 (1967).
51. Doreian, P., Kapuscinski, R., Krackhardt, D. & Szczypula, J. A brief history of balance through time. *J. Math. Sociol.* **21**, 113–131 (1996).
52. Golub, B. & Jackson, M. O. How homophily affects the speed of learning and best-response dynamics. *Q. J. Econ.* **127**, 1287–1338 (2012).
53. Melguizo, I. Homophily and the persistence of disagreement. *Econ. J.* **129**, 1400–1424 (2019).
54. Hu, L., Li, A. & Tan, X. A rational inattention theory of echo chamber. <https://doi.org/10.48550/arXiv.2104.10657> (2023).
55. Dandekar, P., Goel, A. & Lee, D. T. Biased assimilation, homophily, and the dynamics of polarization. *Proc. Natl. Acad. Sci.* **110**, 5791–5796 (2013).
56. Bolletta, U. & Pin, P. Polarization when people choose their peers. *SSRN Scholarly Paper* <https://doi.org/10.2139/ssrn.3245800> (2020).
57. Polanski, A. & Vega-Redondo, F. Homophily and influence. *J. Econ. Theory* **207**, 105576 (2023).
58. Axelrod, R. The dissemination of culture: A model with local convergence and global polarization. *J. Confl. Resolut.* **41**, 203–226 (1997).
59. Sunstein, C. R. *Republic.Com* (Princeton University Press, 2001).
60. Mutz, D. C. *Hearing the Other Side: Deliberative versus Participatory Democracy* (Cambridge University Press, 2006).
61. Huckfeldt, R. R., Johnson, P. E. & Sprague, J. D. *Political disagreement: The survival of diverse opinions within communication networks* (Cambridge University Press, 2004).
62. Iyengar, S. & Hahn, K. S. Red media, blue media: Evidence of ideological selectivity in media use. *J. Commun.* **59**, 19–39 (2009).
63. Tokita, C. K., Guess, A. M. & Tarnita, C. E. Polarized information ecosystems can reorganize social networks via information cascades. *Proc. Natl. Acad. Sci.* **118**, e2102147118 (2021).
64. Levy, G. & Razin, R. Echo chambers and their effects on economic and political outcomes. *Annu. Rev. Econ.* **11**, 303–328 (2019).
65. Mutz, D. C. Cross-cutting social networks: Testing democratic theory in practice. *Am. Polit. Sci. Rev.* **96** (2004).
66. Landmore, H. & Page, S. E. Deliberation and disagreement: Problem solving, prediction, and positive dissensus. *Polit. Philos. Econ.* **14**, 229–254 (2015).
67. Ashby, W. R. Requisite variety and its implications for the control of complex systems. *Cybernetica* **1**, 83 (1958).
68. Marcoux, M. & Lusseau, D. Network modularity promotes cooperation. *J. Theor. Biol.* **324**, 103–108 (2013).
69. Gianetto, D. A. & Heydari, B. Network modularity is essential for evolution of cooperation under uncertainty. *Sci. Rep.* **5**, 9340 (2015).
70. Ikemoto, Y. & Sekiyama, K. Modular network evolution under selection for robustness to noise. *Phys. Rev. E* **89** (2014).
71. Kauffman, S. A. *The origins of order: Self-organization and selection in evolution* (Oxford University Press, 1993).
72. Whitehead, H. & Lusseau, D. Animal social networks as substrate for cultural behavioural diversity. *J. Theor. Biol.* **294**, 19–28 (2012).
73. Pescetelli, N., Rutherford, A. & Rahwan, I. Modularity and composite diversity affect the collective gathering of information online. *Nat. Commun.* **12**, 3195 (2021).
74. Schwertman, N. C., Gilks, A. J. & Cameron, J. A simple noncalculus proof that the median minimizes the sum of the absolute deviations. *Am. Stat.* **44**, 38–39 (1990).
75. Chu, J. T. On the distribution of the sample median. *Ann. Math. Stat.* **26**, 112–116 (1955).
76. Hora, S. C., Franssen, B. R., Hawkins, N. & Susel, I. Median aggregation of distribution functions. *Decis. Anal.* **10**, 279–291 (2013).
77. Nehring, K. & Pivato, M. The median rule in judgement aggregation. *Econ. Theory* **73**, 1051–1100 (2022).

78. Dixit, A. K. & Weibull, J. W. Political polarization. *Proc. Natl. Acad. Sci.* **104**, 7351–7356 (2007).
79. Fiorina, M. P. & Abrams, S. J. Political polarization in the American public. *Annu. Rev. Polit. Sci.* **11**, 563–588 (2008).
80. Clauset, A., Newman, M. E. J. & Moore, C. Finding community structure in very large networks. *Phys. Rev. E* **70** (2004).
81. Newman, M. E. J. *Networks* (Oxford University Press, Oxford, 2018).
82. Bala, V. & Goyal, S. Learning from neighbours. *Rev. Econ. Stud.* **65**, 595–621 (1998).
83. Banerjee, A. V. A simple model of herd behavior. *Q. J. Econ.* **107**, 797–817 (1992).
84. Bikhchandani, S., Hirshleifer, D. & Welch, I. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* **100**, 992–1026 (1992).
85. Banerjee, A. & Fudenberg, D. Word-of-mouth learning. *Games Econ. Behav.* **46**, 1–22 (2004).
86. Pérez-Escudero, A. & De Polavieja, G. G. Collective animal behavior from Bayesian estimation and probability matching. *PLoS Comput. Biol.* **7**, e1002282 (2011).
87. Mann, R. P. Collective decision-making by rational agents with differing preferences. *Proc. Natl. Acad. Sci.* **117**, 10388–10396 (2020).
88. Noelle-Neumann, E. The spiral of silence: A theory of public opinion. *J. Commun.* **24**, 43–51 (1974).
89. Gaisbauer, F., Olbrich, E. & Banisch, S. Dynamics of opinion expression. *Phys. Rev. E* **102**, 042303 (2020).
90. Ladha, K. K. Information pooling through majority-rule voting: Condorcet's jury theorem with correlated votes. *J. Econ. Behav. Organ.* **26**, 353–372 (1995).
91. Kaniowski, S. Aggregation of correlated votes and Condorcet's Jury Theorem. *Theory Decis.* **69**, 453–468 (2010).
92. Mäs, M. & Flache, A. Differentiation without distancing: Explaining bi-polarization of opinions without negative influence. *PLoS ONE* **8**, e74516 (2013).
93. Keijzer, M. A. & Mäs, M. The complex link between filter bubbles and opinion polarization. *Data Sci.* **5**, 139–166 (2022).
94. Kaniowski, S. An invariance result for homogeneous juries with correlated votes. *Math. Soc. Sci.* **57**, 213–222 (2009).
95. McPherson, J. M. & Smith-Lovin, L. Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. *Am. Sociol. Rev.* **52**, 370–379 (1987).

## Acknowledgements

This research was funded by the Ministry of Education of Singapore (Grant Number: 022240-00001). The authors wish to thank Nuri Kim and Xiangting Bernice Lin for their helpful comments on an earlier version of the manuscript.

## Author contributions

P.O. designed the models and derived analytical results. P.O. and J.P. conducted numerical experiments and analyzed the results. P.O., J.P., and A.S. wrote and reviewed the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-56868-8>.

**Correspondence** and requests for materials should be addressed to P.O.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024