Perspectives on Psychological Science
1–14

DOI: 10.1177/17456916231198238
www.psychologicalscience.org/PPS

The Spread of Beliefs in Partially Modularized Communities

Robert L. Goldstone[1,2], Marina Dubova[2], Rachith Aiyappa[3], and Andy Edinger[2,3]

[1]Department of Psychological and Brain Sciences, Indiana University; [2]Program in Cognitive Science, Indiana University; and [3]Center for Complex Networks and Systems, Luddy School of Informatics, Computing, and Engineering, Indiana University

Abstract

Many life-influencing social networks are characterized by considerable informational isolation. People within a community are far more likely to share beliefs than people who are part of different communities. The spread of useful information across communities is impeded by echo chambers (far greater connectivity within than between communities) and filter bubbles (more influence of beliefs by connected neighbors within than between communities). We apply the tools of network analysis to organize our understanding of the spread of beliefs across modularized communities and to predict the effect of individual and group parameters on the dynamics and distribution of beliefs. In our Spread of Beliefs in Modularized Communities (SBMC) framework, a stochastic block model generates social networks with variable degrees of modularity, beliefs have different observable utilities, individuals change their beliefs on the basis of summed or average evidence (or intermediate decision rules), and parameterized stochasticity introduces randomness into decisions. SBMC simulations show surprising patterns; for example, increasing out-group connectivity does not always improve group performance, adding randomness to decisions can promote performance, and decision rules that sum rather than average evidence can improve group performance, as measured by the average utility of beliefs that the agents adopt. Overall, the results suggest that intermediate degrees of belief exploration are beneficial for the spread of useful beliefs in a community, and so parameters that pull in opposite directions on an explore–exploit continuum are usefully paired.

Keywords

misinformation, social networks, belief spread, opinions, communities, communication, polarization, echo chambers, filter bubbles

One of the most conspicuous patterns of human organization is the "clumpy" manner in which behaviors and beliefs are organized. People from different cultures organize their world in strikingly different ways (Medin & Atran, 2004). Populations are often divided into partially modularized communities (Girvan & Newman, 2002), with the result that people within a community are much more likely to share the same beliefs than people belonging to different communities. Even within an ostensibly single culture, members belonging to different political (Andris et al., 2015) or racial (DiPrete et al., 2011) groups have dramatically dissimilar, often diametrically opposed, beliefs. The communities along which beliefs and behaviors clump are often geographically distributed, but increasingly often they transcend geography and are instead based on widely distributed ideologies and self-constructed identities (Wellman, 2001). The societal importance of the spread of (mis)information in social networks has drawn attention to how people living a short distance from each other may nonetheless disagree even on matters of verifiable fact (Lazer et al., 2018). Sizable proportions of Americans, for example, believe that COVID-19 vaccines contain digital trackers, the Earth is flat, or that the 2012 fatal mass shooting at Sandy

Corresponding Author:
Robert L. Goldstone, Department of Psychological and Brain Sciences, Indiana University
Email: rgoldsto@indiana.edu

Hook Elementary School was staged by actors and nobody was, in fact, killed.

The spread of falsehoods such as these is a major obstacle to social progress. Our having greater access to information than ever before in history does not prevent large swathes of our population from holding beliefs that are refuted by science and factual records (Hornsey, 2020). There are several reasons why people believe falsehoods (Pennycook & Rand, 2021), including vested interests, susceptibility to deceit (Lewandowsky et al., 2012), desires to fit in a particular group (Del Vicario et al., 2017; Kaba & Beran, 2016), ideologies (Kahan et al., 2012), and establishing defenses against fears (Hornsey et al., 2018). These are powerful psychological factors that fit broadly within the premise that people are flawed, biased, and frequently irrational (Festinger, 1957).

Although not denying that beliefs often result from flawed and irrational psychological processes or have been adapted for different information environments than we currently reside (Mann, 2022), our present goal is to explore individual and social processes that are arguably even more fundamental in spreading beliefs and behaviors in a population. In particular, even if people are trying to maximize the utility of their beliefs, rely on unbiased rather than self-serving utility calculations, and do not engage in deceitful messaging, we argue that there will still be pockets of a population having suboptimal beliefs even when most people do not. In our analysis, believing falsehoods, such as that the childhood measles, mumps, and rubella vaccine causes autism, is a special case of a broader social pattern in which the utility of different opinions or beliefs is unevenly distributed, with some communities holding beliefs with far higher utility than others. The notion of opinion utility discussed here may be considered a measure of opinion *quality*. Low-quality opinions—beliefs in objective falsehoods—may indeed provide utility for belief holders in other ways, such as the entertainment value of hypothetical conspiracy theories or the social value of holding beliefs that allow one to better fit within one's community. We argue that the structure of social networks and the dynamics of belief spread on those networks are strong determinants of the eventual imbalances in the utility of beliefs found across communities.

Computational models of the spread of beliefs in modularized communities play a valuable role in understanding and predicting unfolding patterns of beliefs in a population. It is often hard to predict how beliefs will spread in a community because of the complex interactions among individuals. There are frequently "rich-get-richer" effects such that the more people who adopt a belief the more likely it is to be adopted by others (Nadeau et al., 1993). More subtle dynamics can arise in which noise in the human decision-making process can prevent a markedly superior belief from catching on, or changing one situational factor effectively compensates for the potentially adverse effects of changing another. Our modeling framework is intended to capture interactions such as these, which can help to organize broad empirical patterns. Rather than providing detailed model fits to specific data sets, our modeling goal is to describe some of the fundamental dynamics likely to be at play in real-world situations involving belief spread (Anderson & Ye, 2019; Dalege & van der Does, 2022; Galesic et al., 2021; Smaldino, 2017).

Broad Empirical Patterns in the Spread of Belief

Generalizing over the many political, ideological, religious, and cultural contexts in which people within a population vary dramatically in their beliefs, there are some major empirical patterns that are frequently found:

- *Clumpy distribution of beliefs*. Most obviously, beliefs in a population tend to be distributed in an uneven, clustered fashion such that people in a particular geospatial region are much more likely to share the same beliefs compared with people in different regions. For example, beliefs regarding health risks associated with global climate change strongly cluster at state, county, and tract levels in the United States (Howe et al., 2019). Whereas geospatial clusters are prominent, beliefs clump according not only to ideological groups but also to regional groups, such as the cluster of correlated beliefs that ideological conservatives in the United States had in 2020 regarding personal vulnerability to COVID, the severity of the COVID virus, and the exaggeration of COVID risks by the media (Calvillo et al., 2020).
- *Echo chambers*. Echo chambers are social structures that feature far more interaction among people that share the same compared with different beliefs. For example, people are far more likely to connect to others on social media sites who share their beliefs and are more likely to spread information to like-minded others (Cinelli et al., 2021). In the case of COVID beliefs, echo chambers have been shown to exacerbate the spread of misinformation and impede corrections to that misinformation (van der Linden, 2022). However, the prevalence of echo chambers, and the extent of their effects, are debated (Guess, 2021; LaCour, 2013).

- *Filter bubbles*. Filter bubbles are social structures that lead to people filtering out beliefs that contradict their own (West & Bergstrom, 2021). That is, even if, despite someone's biased, echo chambered social network, they happen to be exposed to a different belief than their own, they still may not be affected. This "filtering out" may result from discrediting the source or ignoring messages on the basis of their presumed content despite being shown opposing beliefs (Ekström et al., 2022; Flaxman et al., 2016). Filter bubbles may also be fueled by computational or algorithmic artifacts wherein some information may never be shown to individuals because the algorithm may not deem it to be relevant to them on the basis of their previous activity (e.g., their likes; Groshek & Koc-Michalska, 2017; Pariser, 2011).
- *Increasing within-community homogeneity*. A common pattern for newly formed groups is for their members to converge in their beliefs over time (Asch, 1956; Flache et al., 2017). This convergence may be due to individuals who possess minority beliefs changing their beliefs because they believe that the majority knows better than they do (informational conformity), or to curry the favor of the majority group (normative conformity).
- *Persisting minority beliefs*. If the beliefs of group members converge over time, then one might assume that with sufficient time, everyone in the population will share the same beliefs, and simple models of belief spread often predict exactly that dynamic. In fact, this is often not found (Axelrod, 1997; Kelly et al., 2006; Lawson, 1997). Rather than gradually disappearing, minority opinions often persist. Some accounts for the surprising resilience of minority opinions are that people are far more influenced by in-group compared with out-group members (Spears, 2021), people may assert their individuality by differentiating themselves from the majority (Brewer & Roccas, 2001), and group members only being able to access beliefs in their local neighborhood (Latané et al., 1995).
- *Polarization*. Opinions in different communities within a population frequently become increasingly divided over time (Koudenburg & Kashima, 2022). A community may systematically move its opinions away from an opposed community. Although a diversity of opinions in a population promotes resilience and flexibility, the risk of polarization is that communication across groups may become severely inhibited if no common ground can be established (McCoy et al., 2018).

The Spread of Beliefs in Modularized Communities Framework

Given the complex dependencies among the patterns described above, a modeling framework that explicitly incorporates degrees of modularity, in-group biases, and alternative decision rules is useful for organizing and explaining the spread of opinions in modularized communities (Smaldino, 2017). Our Spread of Beliefs in Modularized Communities (SBMC) model[1] begins by creating a global population of N agents that are divided into N_C equally sized communities. We use a stochastic block model (Karrer & Newman, 2011) to generate social networks in which agents within a community have a probability P_{in} of being connected and agents across different communities have a probability P_{out} of being connected. If communities are highly modularized, then $P_{in} \gg P_{out}$, as would be expected for echo chambers.

N agents are randomly assigned one of N_O opinions, with $N = 200$ and $N_O = 10$ for all of our simulations. In many agent-based models, all opinions are equally good, but in the SBMC model the opinions are rank-ordered, and their utility is an exponential function of their rank. Once initialized with random opinions and social connections, agents exchange opinions with their neighbors for 20 rounds. When exchanging opinions, agents have access to the exact utilities of their neighbors' opinions (there is no communication noise). An agent will tend to adopt an opinion if it is shared by many of its neighbors and if its utility is high. For a description of the important parameters of the SBMC model and their default values, see Table 1.

Each agent integrates evidence from its in-group and out-group neighbors from the current time step to determine whether it should change its opinion at the next time step according to

$$E_{i_o} = \sum_{x \in \{N \mid x_o = i_o\}} \left(W_{in} \cdot \delta\left(x^c, i^c\right) + W_{out} \cdot \left(1 - \delta\left(x^c, i^c\right)\right)\right) \cdot e^{-dU_o}, \quad (1)$$

where the raw evidence for Opinion o for Agent i, E_{i_o}, is computed by summing the evidence for o across all agents that are neighbors of i (denoted by x) holding Opinion o. If the neighbor belongs to the same community as i, $i^c = x^c$, then its opinion is given weight W_{in}; otherwise it is given weight W_{out}. This is captured by the delta function ($\delta(x^c, i^c)$). Filter bubbles that may arise because of individuals discrediting the source or ignoring messages on the basis of their presumed content can be modeled by $W_{in} \gg W_{out}$. The utility of an opinion is an exponential function of its rank, U_o. The rank of the best opinion is 1, the second best opinion's rank is 2, and so on. The drop-off in utility with rank

Table 1. Parameters Varied in the Spread of Beliefs in Modularized Communities Model and Their Default Values

Parameter	Interpretation	Default value
N	Number of agents	200
N_O	Number of opinions	10
W_{in}	Weight given to in-group members	1
W_{out}	Weight given to out-group members	0.1
γ	Choice determinism	1
g	Evidence-integration strategy	0 (summation)
N_C	Number of communities	4
P_{in}	Probability that two in-group members are neighbors	0.12
P_{out}	Probability that two out-group members are neighbors	0.004
R	Number of rounds of opinion exchange	20

is controlled by the parameter d, which is set at 0.1. An exponential function is used to transform rank into utility to capture the common pattern that the practical difference between the best and second-best options is greater than the difference between, say, the eighth and ninth best options (Kable & Glimcher, 2007). The integrated evidence for Opinion o for Agent i is its raw evidence normalized by an integration function that generalizes different evidence accumulation rules, as governed by

$$I_{i_o} = \frac{E_{i_o}}{g \cdot (\sum_{x \in \{N | x_o = i_o\}} (W_{in}\delta(x^c, i^c) + W_{out}(1 - \delta(x^c, i^c)) - 1)) + 1}. \quad (2)$$

The two summations in the denominator normalize the evidence by the total amount of possible evidence given the number of in-group and out-group neighbors of i. The parameter g controls whether an agent uses a summation, averaging, or blended decision rule. If $g = 0$, then the denominator resolves to 1, meaning that the integrated evidence is simply the sum of the evidence for o. If $g = 1$, then the denominator resolves to the number of neighbors possessing o, meaning that the integrated evidence is the average evidence for o. As g increases within the $0 < g < 1$ range, the integrated evidence becomes increasingly closer to an averaging than summation rule. By varying g our agents shift from choosing popular opinions ($g = 0$) to choosing ostensibly better ($g = 1$) opinions. Whenever $g < 1$, averaging is not equivalent to taking the sum of neighbors' evidence and dividing by the number of neighbors because the total amount of evidence has an influence beyond the average. For example, imagine that an agent has nine neighbors adopting Opinion X, which has a value of 1, and three neighbors adopting Opinion Y, which has a value of 2. If $g = 0$, then the integrated evidence for Opinion X (9*1/1 = 9) will be greater than the integrated evidence for Opinion Y (2*3/1 = 6). This ordering reverses if $g = 1$ because the integrated evidence of Opinion Y (2*3/(1*(3 − 1) + 1) = 2) will be greater than the integrated evidence for Opinion X (9*1/(1*(9 − 1) + 1) = 1). Finally, if $g = 0.1$, then the integrated evidence of Opinions X (9*1/(0.1*(9 − 1) + 1) = 5) and Y (2*3/(0.1*(3 − 1) + 1) = 5) will be exactly the same, with agents compromising between valuing popular and high-utility opinions.

Using a softmax decision rule, the probability of Agent i adopting Opinion o in the next round of opinion exchange is an exponential function of the integrated evidence for o, I_{i_o}, divided by the integrated evidence that i has for every one of the N_o possible opinions:

$$P(i_o) = \frac{e^{\gamma I_{i_o}}}{\sum_{n=1}^{N_o} e^{\gamma I_{i_n}}}. \quad (3)$$

The parameter γ controls the determinism of the selection of opinions. As γ increases, Agent i will increasingly choose the opinion that has the greatest integrated evidence for it. When γ is low, then opinion selection will be more random, with opinions having relatively little evidential support being more likely to be selected.

Simulation Patterns

Figure 1 shows characteristic networks that arise when the SBMC model is run using the default parameters in Table 1, with 50 replications run for each set of parameters. Throughout all of the simulation figures, agents

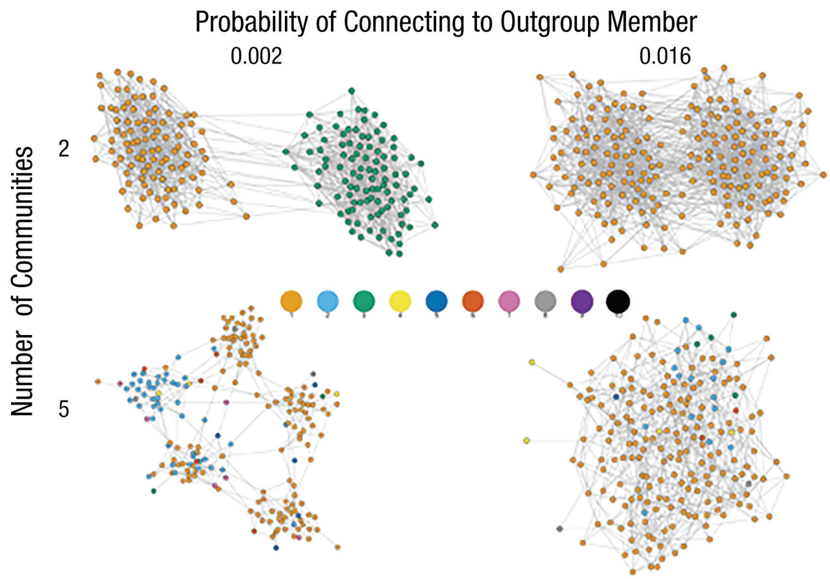

Fig. 1. Sample graphs when the number of communities and probability of an individual connecting to an out-group member are independently manipulated. The assignment of opinion ranks to colors is shown in the center and used for all subsequent figures. Accordingly, agents selecting the opinion with the highest utility are shown in orange, second highest in blue, third in green, and so on. This same rank-to-color conversion scale is used for all figures.

adopting the opinions with the best, second-best, and third-best utilities are shown in orange, blue, and green, respectively. Accordingly, all of the agents are selecting relatively good opinions, often the best opinion. Nonetheless, the distribution of opinions shows many of the broad empirical patterns previously described: Opinions are distributed in a highly clustered manner that echoes the network's community structure, minority opinions persist within particular communities, and opinions within a community become increasingly homogeneous over time. When the probability of connecting to out-group members is low, then different communities within the larger population often adopt different opinions, consistent with echo chambers and prior models (Flache & Macy, 2011). In these cases, all of the agents in a community may adopt a suboptimal opinion, and even if they are exposed to a better opinion, they will not adopt it because the evidence supporting it does not surpass the sum (default; $g = 0$) of the evidence from neighbors adopting suboptimal opinions. This figure also shows that opinions (both optimal and suboptimal) are no longer tightly restricted within communities when the probability of out-group members is high. If one's goal is simply to have equally high-utility opinions spread throughout the population, then increasing the probability of out-group members connecting is beneficial, but further simulations are needed to determine whether and when that increase raises overall opinion quality.

Number of Communities × Evidence Integration

Moving on to other systematic parameter sweeps, Figure 2 shows results from factorially combining four numbers of communities (2, 5, 8, and 12) and four levels of the evidence-integration parameter g (0 = pure summation of evidence; 0.25, 0.75, and 1 = pure averaging of evidence).[2] The value plotted on the vertical axis is the average utility of opinions adopted by agents after 20 rounds of opinion exchange.[3] The highest ranked opinion is given a utility of 5, the second highest a utility of 4, and so on down to the 10th best opinion, which is given a utility of −5. Overall, agents find higher utility opinions when they sum the evidence of their neighbors compared with using the average value for each opinion possessed by one of the neighbors. This might seem surprising. Why is sensitivity to the actual utility of an opinion not the best cue to whether it should be adopted? The reason why summing is beneficial is that each opinion possessed by an agent is the result of an inherently probabilistic process because of the softmax decision rule. Although any single neighbor's opinion will be noisy, if many neighbors are all adopting the same opinion, then it is likely to have high utility. Summing evidence from imperfectly formed opinions is an effective way to amplify and spread subtle utility differences among opinions.

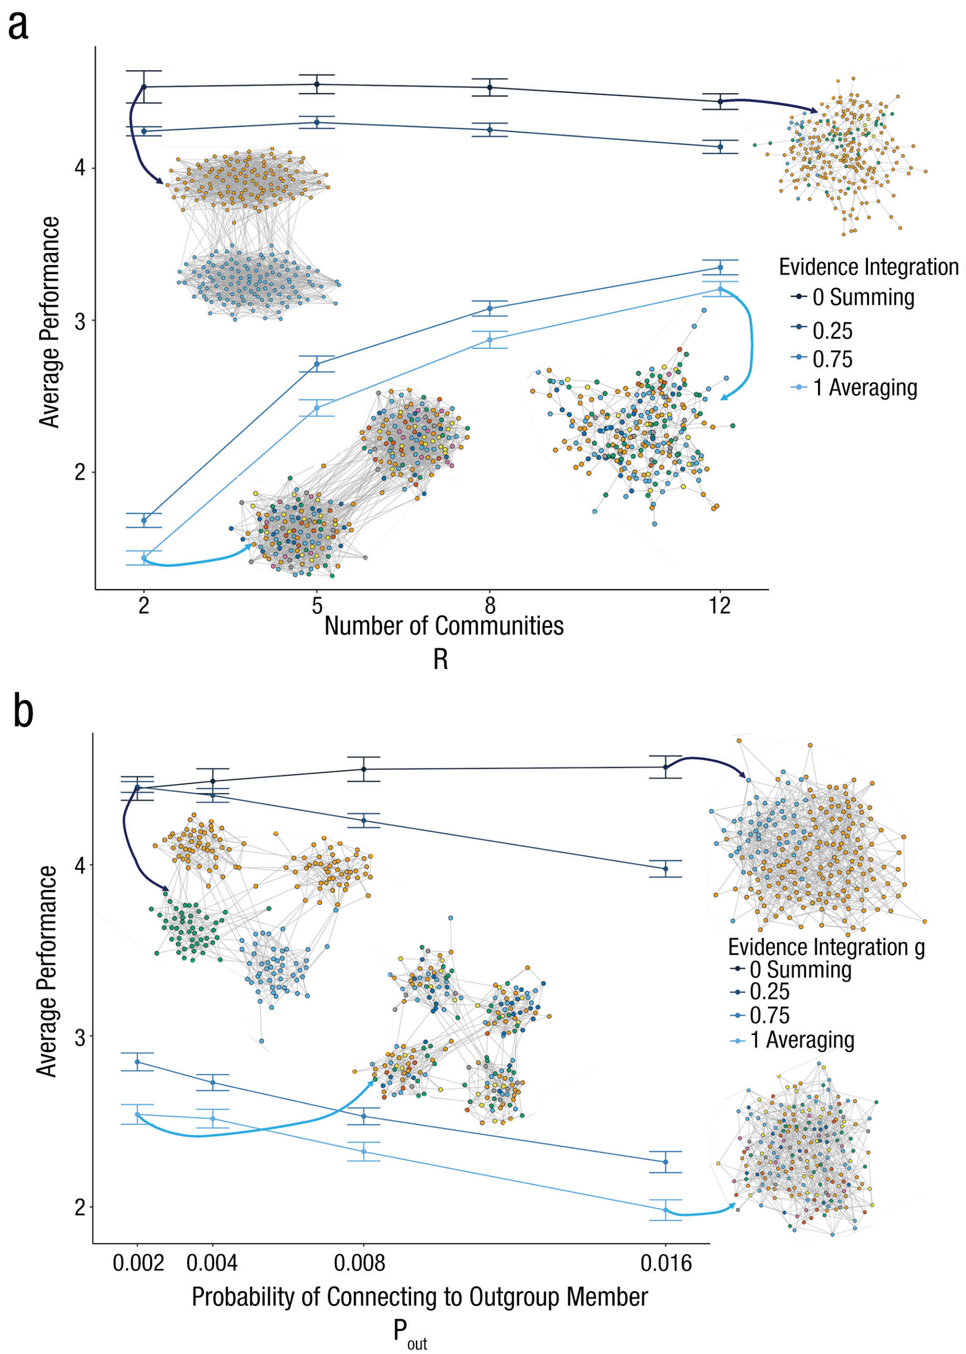

Fig. 2. Two classes of simulations exploring interactions involving evidence integration (g). Panel (a) shows the average performance from simulations that vary two parameters, the number of communities, and the method for integrating evidence across neighbors. Error bars represent 95% confidence intervals. When evidence integration = 0, agents sum all of the evidence from their neighbors. When evidence integration = 1, they average the evidence. Sample networks from the four most extreme parameter combinations are also shown. Panel (b) shows the average performance varying the probability that two agents belonging to different communities are connected and the method for integrating evidence across neighbors. Although summing generally works better than averaging as an evidence-integration strategy given the other model parameters, increasing out-group connectivity is beneficial for summing but impairs performance for averaging.

When choices are made by averaging rather than summing the evidence for each opinion, performance is overall lower. Many more opinions are adopted by agents. No single opinion catches on widely because averaging does not amplify widely held opinions as does summing. The wide variety of opinions in Figure 2 when $g > 0.5$, as shown by the many colors present in the networks, may be beneficial from the perspective of diversity but is ineffective in terms of the average utility of held opinions. Each agent has a tendency to choose opinions with high average utilities, but given the noisy decision process ($\gamma = 1$), agents also frequently make choices that deviate from the assessed utilities for opinions. Collective performance dramatically improves as the number of communities increases. The overall connectivity of the population decreases as the number of communities increases because the probability of out-group members connecting is 30 times less than it is for in-group members (see Table 1). It may seem counterintuitive that decreasing connectivity helps the population find better opinions. The explanation is that exposing agents to too many different opinions prevents the agents from settling down on any opinion, including the relatively good ones. When an agent is part of a large community (e.g., number of communities = 2), the sheer number of opinions that they are exposed to interferes with any opinion spreading widely. This result, that limiting the access of agents to other agents' opinions can improve collective performance, has been found in other simulations (Barkoczi & Galesic, 2016; Galesic et al., 2018; Lazer & Friedman, 2007; Massari et al., 2019; Smaldino et al., 2022; Zollman, 2007) as well as experiments with human groups (Mason et al., 2008). One problem with abundant connectivity is that agents in a population may prematurely converge on a single solution without fully exploring the problem space (Barkoczi & Galesic, 2016; Lazer & Friedman, 2007). The current simulations point to another problem with plentiful neighbors—the inability to extract signal (e.g., high-value beliefs) from noise (Hills, 2019; Roozenbeek et al., 2022).

Out-Group Connectivity × Evidence Integration

Whereas Figure 2a explores connectivity by varying the number of communities into which a population is modularized, Figure 2b explores it by varying the probability of an agent connecting with each out-group member (p_{out}), keeping constant the probability of connecting to an in-group member (p_{in}= 0.12) and the number of communities (N_c= 4). As with the previous simulation, we find that collectives generally perform better when their members sum, rather than average, evidence. However, we also find that this evidence-integration strategy strongly moderates the effect of increasing out-group connectivity. Whereas increasing out-group connectivity modestly improves performance when agents make their judgments on the basis of summed opinion values, it strongly depresses performance when agents average evidence for opinions across their neighbors. These results can be readily integrated with the previous simulation—having fewer, larger communities, as in Figure 2a, and having greater out-group connectivity, as in Figure 2b, both increase connectivity and impair performance when agents use average opinion value. Exposing agents to more numerous and diverse opinions prevents any opinion from widely catching on in either the population as a whole or even a community. The sample graphs show that when averaging agents belong to relatively insular communities, there is still substantial diversity within each community, but there is more rallying of agents around relatively good opinions than when communities intermingle more.

These simulations offer a nuanced perspective on informational isolation in populations. Many people assume that increasing connectivity across communities will always be positive for the population considered as a whole. Indeed, there are benefits for increasing connectedness across members belonging to different communities in terms of decreasing out-group stereotyping (Turner et al., 2007) and increasing empathy for others who belong to different racial, religious, gender, or age groups (Pettigrew & Tropp, 2008). However, our simulations contradict the universality of this assumption when agents choose opinions by assessing average values. Increasing connectivity across communities can sometimes be detrimental to each community's development and maintenance of a coherent viewpoint, consistent with our results. For example, immigrants who place particular value on their own culture and its customs frequently have higher self-esteem (Rumbaut, 1994), and elements of their culture are better preserved (Bloemraad et al., 2008). Although the agents of our simulations have no real culture to speak of, human immigrants and our SBMC agents with high out-group connection probabilities face a common challenge—preserving effective solutions within their community. When out-group connectivity is high, it is difficult for effective solutions to gain stable purchase within a community.

A second common assumption is that increasing communication and permeability across communities will decrease the overall diversity of the population. As different communities interact more, it seems reasonable to expect their differences to reduce and hence for the overall diversity of the population to diminish.

Our simulations contradict this commonsense assumption as well, at least for averaging agents. Increasing out-group connectivity leads to greater opinion diversity when coupled with the averaging decision rule because no opinion becomes widespread.

Evidence Integration × Decision Determinism

An important consideration for decision makers is the degree to which they make the choice that maximizes expected value versus adding some randomness to choices. An agent might want to make a choice that does not maximize immediate expected value to explore a broader range of options. If one always eats at the same good, local restaurant, then one might be missing out on an even better option. Decision makers who face an uncertain world must navigate an *explore–exploit* trade-off, choosing options with unknown value that might offer downstream benefits (explore), options that maximize value according to current estimates (exploit), or a compromise between these poles (Hills et al., 2010; Wu et al., 2020). In the SBMC model, γ captures this compromise, with higher values of γ pushing agents more to the exploit end of the explore-exploit continuum. Independently varying γ and the strategy for integrating evidence reveals that agents generally do better when they choose opinions that maximize value (e.g., high γ value; see Fig. 3a). However, these value-maximizing choices do not always lead to the best average group performance, specifically when agents sum, rather than average, evidence. The problem with exploiting too much for summing agents is that the agents will early on converge on a suboptimal opinion, and no amount of subsequent evidence for a better opinion coming from a few agents will be able to supplant that early-adopted opinion with a large sum supporting it (Sang et al., 2020).

Figure 3a also shows a strong interaction between γ and the evidence-integration strategy. The previously reported pattern of agents performing better when they sum rather than average is, once again, found for agents using a relatively exploratory choice rule (e.g., low γ). Choice rules with considerable randomness pair well with summing because the randomness allows agents to occasionally make choices that go against a suboptimal but strong majority opinion that may have been adopted early on on the basis of limited, local evidence. However, for highly deterministic choice rules, averaging is now advantageous over summing. Averaging opinion values is a noisier strategy than summing in the sense that it does not amplify high-value, popular opinions. Averaging is usefully paired with a deterministic rule that is sensitive to small differences between opinion values. Overall, the results from these simulations suggest that good decision rules adopt an intermediate position on the explore–exploit continuum. Exploration leads to a diversity of opinions, which is beneficial for displacing a suboptimal opinion with a more optimal one. But diversity can also prevent optimal opinions from spreading. Averaging and low determinism both increase exploration, whereas summing and high determinism both increase exploitation. Advantageous outcomes are obtained when an exploration-biasing factor is combined with an exploitation-biasing one. Other collective-behavior paradigms have found benefits for combining parameters to promote intermediate degrees of diversity (Smaldino et al., 2022), with the added twist that more diversity tends to be beneficial for more difficult search problems (Campbell et al., 2022). Averaging agents directly use true utilities to guide their choices (exploiting), which pairs well with considerable noise added to decisions (exploring).

Outgroup Connectivity × Decision Determinism

As a final simulation, we factorially combined four different levels of choice determinism with four levels of out-group connectivity. Figure 3b shows the results. Even more strikingly than Figure 3a, this simulation shows benefits for an intermediary degree ($\gamma = 0.8$) of determinism. Given the default choice of summing, rather than averaging, agents, lower γ values provide beneficial diversity to the amplifying effect of summing. As with Figure 3a, decreasing out-group connectivity sometimes promotes performance—this time only for agents adopting a relatively random decision rule. By contrast, more deterministic, value-maximizing agents perform modestly better as out-group connectivity increases. We interpret this interaction in terms of the benefits of intermediate degrees of exploration and diversity, consistent with the previous section. Like averaging, increasing out-group connectivity biases agents toward exploration and diversification. If the agents are already exploratory because of a low γ value, then adding even more opinion diversity by increasing out-group connectivity is undesirable. However, if agents are exploitation-biased, then the diversity provided by greater access to out-group opinions is beneficial.

Although both of the factors controlling an agent's position on the explore–exploit continuum are internal to an individual's decision rule in Figure 3a, Figure 3b shows that external and internal factors can also compensate for each other. Depending on the network structure that an agent finds itself in, it would ideally adopt different γ values, all else being equal. Together with

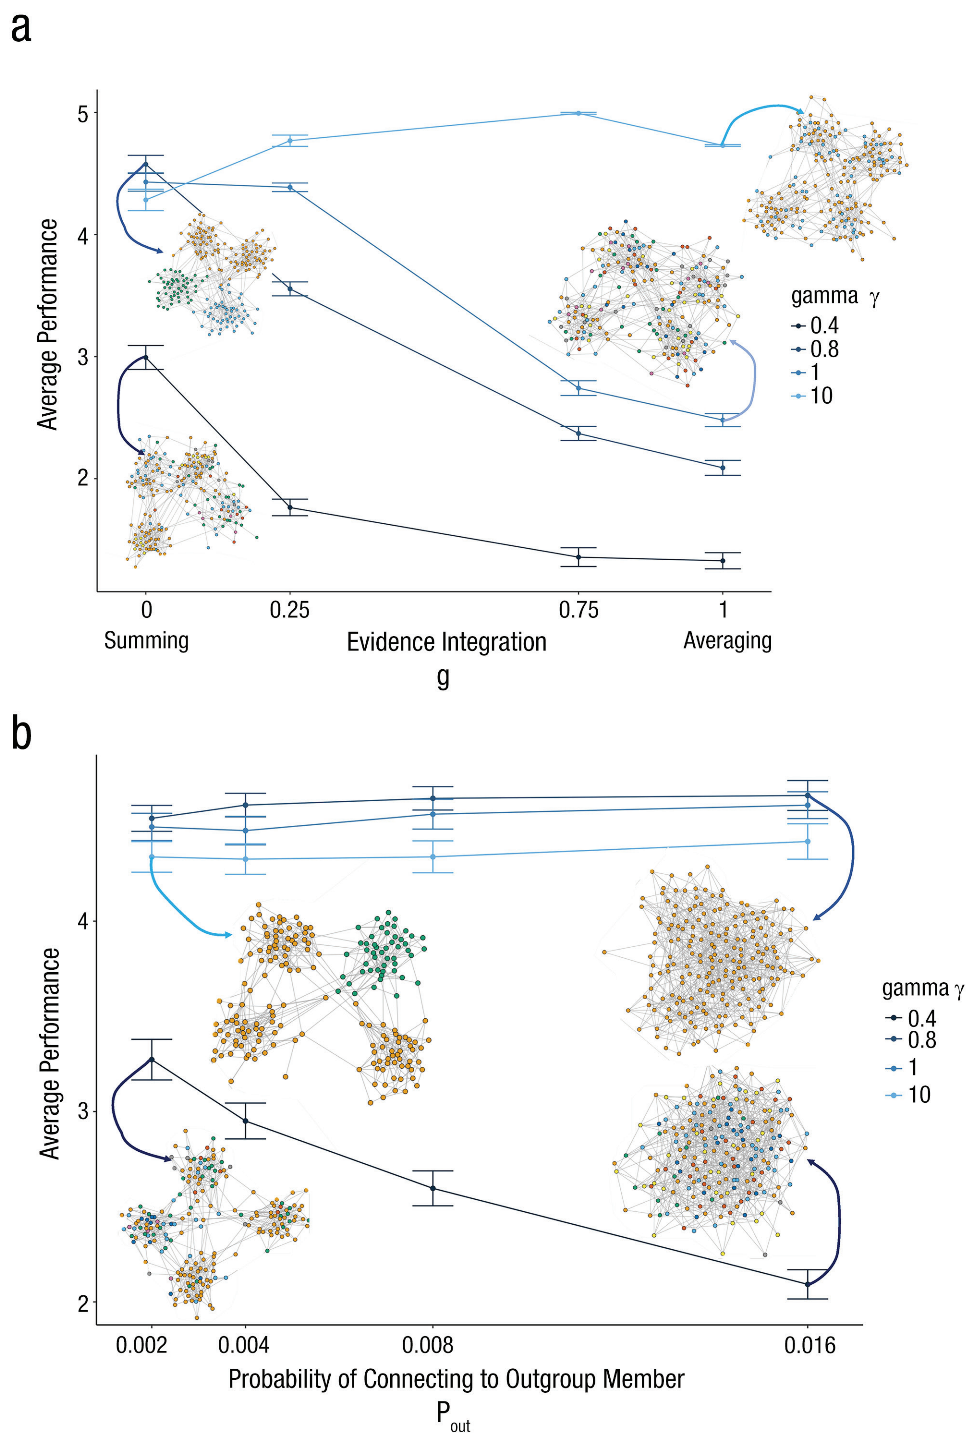

Fig. 3. Two classes of simulations exploring interactions involving decision determinism (γ). Panel (a) shows the average performance varying the agents' evidence-integration rule and their decision determinism. Error bars represent 95% confidence intervals. Whereas summing agents perform better than averaging ones when choices are made with considerable randomness, the reverse is found when agents more deterministically choose the opinion that has the greatest evidence. Panel (b) shows the average performance varying the probability that two agents belonging to different communities are connected and agents' decision determinism. Relatively high levels of out-group connectivity help agents using more deterministic choice rules and impair agents using more random choice rules.

other simulations showing compensatory relations between network structure and individual decision rules (Barkoczi & Galesic, 2016; Goldstone et al., 2013), the current results point to both complexity and flexibility for networked cognitive agents. Agents that find themselves in a particular network configuration can still make decisions that are good for themselves and their group if they can strategically control internal factors in their decision-making (Goldstone & Theiner, 2017), such as their choice determination and strategy for integrating evidence. Conversely, if agents are relatively fixed in their decision-making strategies, then it may be possible to restructure their social network to accommodate these internal strategies (Goldstone et al., 2006).

Discussion

The SBMC framework presented here is currently too simplified to offer compelling recommendations for how humans should make choices between opinions expressed in their social networks or how these social networks should be constructed. It does, however, serve other purposes. It provides examples of collective dynamics that might be expected to be common because they derive from a simple and general model without processes that are tailored to fit a specific scenario. As it turns out, the observed collective dynamic patterns run counter to many commonly held assumptions about belief spread. For example, contradicting some intuitions, the simulations showed the following:

- Performance can be better for collectives made up of agents that sum, rather than average, the evidence/advice coming from their neighbors. Averaging might be assumed to be better than summing because it relies on the actual values of opinions and is not distorted by popularity. However, when there is some noise in the decision process, summing is superior because subtle differences in values are amplified by considering the number of neighbors adopting each opinion. This simulation result is consistent with the advantageous signal-amplifying effect that even uninformed individuals have on their groups' ability to track a resource (Couzin et al., 2011).
- Decreasing connectivity in a population by dividing it into many partially isolated communities can improve the population's overall performance (see also Cantor et al., 2021; Derex & Boyd, 2016). Keeping the total population constant by dividing it into more communities results in less connectivity among agents but often improves the average quality of opinions in the simulations. It is easier for a good opinion to catch on in a small community, and once it has caught on, it is hard for it to be displaced by randomness in choices.
- Decreasing the probability that out-group members will be connected can sometimes improve an entire population's opinions. Although the adage "the more information the better" may sound plausible, our collectives often spread better opinions when communities are relatively isolated from each other because communities exposed to too many opinions have difficulty rallying around any single opinion.
- Better opinions can arise when agents do not always choose the opinion with better evidence for it. Agents that deterministically choose the opinion with the highest value run the risk of missing out on better opinions that have not yet had a chance to become popular within their community. By adopting opinions that do not maximize value among neighbors' current opinions, an agent can sacrifice short-term reward for the possibility of finding better long-term rewards.

A second contribution of the SBMC framework is that it suggests patterns of interactions between the above factors. There is an unfortunate tendency in psychology to look for one-size-fits-all solutions. The real world often resists simple accounts in terms of main effects, revealing important interactions and moderators. The SBMC provides an organizational framework for understanding interactions in terms of two yoked trade-offs—explore versus exploit and diversity versus convergence. These continua are yoked because exploration leads to collective diversity of opinions, whereas exploiting the existing evidence to choose the best opinion among current options leads to collective convergence. Either pole is problematic. Too much exploiting leads to premature convergence of the group on good but not great solutions. Too much exploration leads to a failure of the group to eventually converge on any opinion. The explore–exploit continuum provides a compelling account for why particular pairings of factors either do or do not work well together. If one factor tends to pull agents toward exploration, then it will be well paired with another factor that pulls agents toward exploiting. This generalization provides a coherent synthesis for the patterns of performance shown in Figure 4. In this figure, a pair of factors is categorized as positive (green line) if the pair performs better than expected from each factor's main effect and negative (red line) otherwise. The factor pairs that create relatively high-value opinions all have offsetting pulls toward both exploration and exploitation. The factors that pull toward exploration are high out-group connectivity, noisy

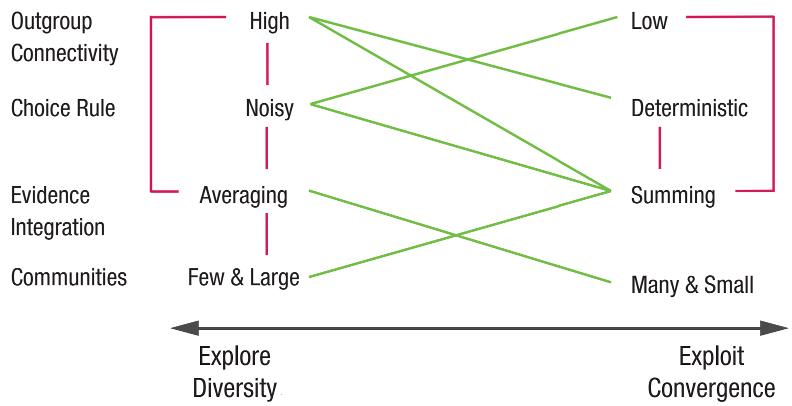

Fig. 4. A summary of the Spread of Beliefs in Modularized Communities simulations. Factor pairs that result in relatively (compared with each factor's main effect) high-value opinions are connected in green, whereas poorly performing factor pairs are connected in red. Factor pairs that pull in opposite directions along the explore-exploit continuum tend to perform well.

choice rule, averaging, and having a few large communities. The factors that pull toward exploiting are low out-group connectivity, deterministic choice rule, summing, and having many small communities. One might think that having many small communities should be construed as exploration-biasing, but the probability of agents being connected across, compared to within, communities is always very small, and so agents will be exposed to far more, possibly diverse, opinions when there are a few large communities.

Although deriving predictions and interventions for specific real-world social networks using the SBMC framework is difficult and prone to error, the framework does make general recommendations. First, it cautions against common assumptions such as that increasing connectivity across groups is always beneficial for the population. In fact, there is evidence that an evidence-resistant minority can retard consensus formation as connectivity between this minority and the general population increases (Lewandowsky et al., 2019). We concur with the general sentiment that polarization is problematic if communities become so separated that they cannot benefit from each other's discoveries. However, our simulations suggest a complex relation between out-group connectivity and average performance, with less out-group connectivity benefiting agents with either significant randomness in their choices or using an averaging strategy. Consistent with simulations showing population-wide benefits when members distrust out-group members (Fazelpour & Steel, 2022), the SBMC framework suggests that populations at risk for homogeneity can benefit from decreasing out-group connectivity.

There are certainly limitations to replacing true beliefs with high-value opinions and community isolation with low out-group connectivity, and future work could incorporate psychological processes related to social identity, emotion, individual differences, racism, and deliberate deceit. Still, an agent-based model approach provides an unusual and fertile perspective on the social problem of misinformation. It shows that even well-meaning and unbiased agents that are sensitive to the quality of different opinions will often live in communities in which opinions are distributed in a decidedly uneven manner. Furthermore, there are advantages for the population in having belief clusters, bearing in mind that groups prosper not only when they find the best opinion from among its members' currently held opinions but also when they build in processes for creating and spreading valuable new opinions.

Transparency

Action Editor: David Garcia
Editor: Interim Editorial Panel
Declaration of Conflicting Interests
The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

ORCID iDs

Robert L. Goldstone https://orcid.org/0000-0001-8357-8358
Andy Edinger https://orcid.org/0000-0001-8159-2121

Acknowledgments

We thank Mirta Galesic, David Garcia, Henrik Olsson, Stephan Lewandowsky, Ben Motz, Paul Smaldino, Peter Todd, Jennifer Trueblood, and Kevin Zollman for helpful comments and discussions that significantly improved this work.

Notes

1. The source code for the simulations to be reported can be accessed at https://github.com/rgoldsto/SBSC.
2. For the simulations to follow, we focus on the average performance of agents in a population, but many other measures are relevant. In the Supplemental Material available at https://github.com/rgoldsto/SBSC, we report five measures: average performance (the average quality of opinions of the agents in a population), percentage best (the percentage of agents who have the highest ranked opinion), entropy (as agents adopt a more diverse set of opinions, entropy increases), mutual information (the dependency between the communities to which agents belong and the opinions possessed by agents), and assortativity (the average probability of two network neighbors having the same opinion).
3. Given that Equations 1 through 3 describe a dynamic process, how opinions spread over time is important, and graphs for these dynamics are presented in the Supplemental Material available at https://github.com/rgoldsto/SBSC.

References

Anderson, B. D. O., & Ye, M. (2019). Recent advances in the modelling and analysis of opinion dynamics on influence networks. *International Journal of Automation and Computing*, *16*(2), 129–149. https://doi.org/10.1007/s11633-019-1169-8

Andris, C., Lee, D., Hamilton, M. J., Martino, M., Gunning, C. E., & Selden, J. A. (2015). The rise of partisanship and super-cooperators in the U.S. House of Representatives. *PLOS ONE*, *10*(4), Article e0123507. https://doi.org/10.1371/journal.pone.0123507

Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, *70*(9), 1–70. https://doi.org/10.1037/h0093718

Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, *41*(2), 203–226. https://doi.org/10.1177/0022002797041002001

Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. *Nature Communications*, *7*(1), Article 13109. https://doi.org/10.1038/ncomms13109

Bloemraad, I., Korteweg, A., & Yurdakul, G. (2008). Citizenship and immigration: Multiculturalism, assimilation, and challenges to the nation-state. *Annual Review of Sociology*, *34*(1), 153–179. https://doi.org/10.1146/annurev.soc.34.040507.134608

Brewer, M. B., & Roccas, S. (2001). Individual values, social identity, and optimal distinctiveness. In C. Sedikides & M. B. Brewer (Eds.), *Individual self, relational self, collective self* (pp. 219–237). Psychology Press.

Calvillo, D. P., Ross, B. J., Garcia, R. J. B., Smelter, T. J., & Rutchick, A. M. (2020). Political ideology predicts perceptions of the threat of COVID-19 (and susceptibility to fake news about it). *Social Psychological and Personality Science*, *11*(8), 1119–1128. https://doi.org/10.1177/1948550620940539

Campbell, C. M., Izquierdo, E. J., & Goldstone, R. L. (2022). Partial copying and the role of diversity in social learning performance. *Collective Intelligence*, *1*(1). https://doi.org/10.1177/26339137221081849

Cantor, M., Chimento, M., Smeele, S. Q., He, P., Papageorgiou, D., Aplin, L. M., & Farine, D. R. (2021). Social network architecture and the tempo of cumulative cultural evolution. *Proceedings of the Royal Society B: Biological Sciences*, *288*(1946), Article 20203107. https://royalsocietypublishing.org/doi/full/10.1098/rspb.2020.3107

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences, USA*, *118*(9), Article e2023301118. https://doi.org/10.1073/pnas.2023301118

Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., Conradt, L., Levin, S. A., & Leonard, N. E. (2011). Uninformed individuals promote democratic consensus in animal groups. *Science*, *334*(6062), 1578–1580. https://www.science.org/doi/full/10.1126/science.1210280

Dalege, J., & van der Does, T. (2022). Using a cognitive network model of moral and social beliefs to explain belief change. *Science Advances*, *8*(33), Article eabm0137. https://doi.org/10.1126/sciadv.abm0137

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

Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences, USA*, *113*(11), 2982–2987. https://doi.org/10.1073/pnas.1518798113

DiPrete, T. A., Gelman, A., McCormick, T., Teitler, J., & Zheng, T. (2011). Segregation in social networks based on acquaintanceship and trust. *American Journal of Sociology*, *116*(4), 1234–1283. https://doi.org/10.1086/659100

Ekström, A. G., Niehorster, D. C., & Olsson, E. J. (2022). Self-imposed filter bubbles: Selective attention and exposure in online search. *Computers in Human Behavior Reports*, *7*, Article 100226. https://doi.org/10.1016/j.chbr.2022.100226

Fazelpour, S., & Steel, D. (2022). Diversity, trust, and conformity: A simulation study. *Philosophy of Science*, *89*(2), 209–231.

Festinger, L. (1957). *A theory of cognitive dissonance* (Vol. 2). Stanford University Press.

Flache, A., & Macy, M. W. (2011). Small worlds and cultural polarization. *The Journal of Mathematical Sociology*, *35*(1–3), 146–176. https://doi.org/10.1080/0022250X.2010.532261

Flache, A., Mäs, M., Feliciani, T., Chattoe-Brown, E., Deffuant, G., Huet, S., & Lorenz, J. (2017). Models of social influence: Towards the next frontiers. *Journal of Artificial Societies and Social Simulation*, *20*(4), Article 2. https://doi.org/10.18564/jasss.3521

Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, *80*(S1), 298–320.

Galesic, M., Barkoczi, D., & Katsikopoulos, K. (2018). Smaller crowds outperform larger crowds and individuals in realistic task conditions. *Decision*, *5*, 1–15. https://doi.org/10.1037/dec0000059

Galesic, M., Olsson, H., Dalege, J., van der Does, T., & Stein, D. L. (2021). Integrating social and cognitive aspects of belief dynamics: Towards a unifying framework. *Journal of the Royal Society Interface*, *18*(176), Article 20200857. https://doi.org/10.1098/rsif.2020.0857

Girvan, M., & Newman, M. E. J. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences, USA*, *99*(12), 7821–7826. https://doi.org/10.1073/pnas.122653799

Goldstone, R. L., Jones, A., & Roberts, M. (2006). Group path formation. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, *36*(3), 611–620. https://doi.org/10.1109/TSMCA.2005.855779

Goldstone, R. L., & Theiner, G. (2017). The multiple, interacting levels of cognitive systems (MILCS) perspective on group cognition. *Philosophical Psychology*, *30*(3), 338–372. https://doi.org/10.1080/09515089.2017.1295635

Goldstone, R. L., Wisdom, T. N., Roberts, M. E., & Frey, S. (2013). Learning along with others. In B. H. Ross (Ed.), *Psychology of learning and motivation* (Vol. 58, pp. 1–45). Academic Press. https://doi.org/10.1016/B978-0-12-407237-4.00001-3

Groshek, J., & Koc-Michalska, K. (2017). Helping populism win? Social media use, filter bubbles, and support for populist presidential candidates in the 2016 us election campaign. *Information, Communication & Society*, *20*(9), 1389–1407.

Guess, A. M. (2021). (Almost) everything in moderation: New evidence on Americans' online media diets. *American Journal of Political Science*, *65*(4), 1007–1022.

Hills, T. T. (2019). The dark side of information proliferation. *Perspectives on Psychological Science*, *14*(3), 323–330. https://doi.org/10.1177/1745691618803647

Hills, T. T., Todd, P. M., & Goldstone, R. L. (2010). The central executive as a search process: Priming exploration and exploitation across domains. *Journal of Experimental Psychology: General*, *139*, 590–609. https://doi.org/10.1037/a0020666

Hornsey, M. J. (2020). Why facts are not enough: Understanding and managing the motivated rejection of science. *Current Directions in Psychological Science*, *29*(6), 583–591. https://doi.org/10.1177/0963721420969364

Hornsey, M. J., Harris, E. A., & Fielding, K. S. (2018). The psychological roots of anti-vaccination attitudes: A 24-nation investigation. *Health Psychology*, *37*(4), 307–315. https://doi.org/10.1037/hea0000586

Howe, P. D., Marlon, J. R., Wang, X., & Leiserowitz, A. (2019). Public perceptions of the health risks of extreme heat across US states, counties, and neighborhoods. *Proceedings of the National Academy of Sciences, USA*, *116*(14). https://doi.org/10.1073/pnas.1813145116

Kaba, A., & Beran, T. N. (2016). Impact of peer pressure on accuracy of reporting vital signs: An interprofessional comparison between nursing and medical students. *Journal of Interprofessional Care*, *30*(1), 116–122.

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625–1633. https://doi.org/10.1038/nn2007

Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., & Mandel, G. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, *2*(10), 732–735. https://doi.org/10.1038/nclimate1547

Karrer, B., & Newman, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Physical Review E*, *83*(1), Article 016107. https://doi.org/10.1103/PhysRevE.83.016107

Kelly, J. W., Fisher, D., & Smith, M. (2006). Friends, foes, and fringe: Norms and structure in political discussion networks. In *dg.o '06: Proceedings of the 2006 International Conference on Digital Government Research* (pp. 412–417). Digital Government Society of North America.

Koudenburg, N., & Kashima, Y. (2022). A polarized discourse: Effects of opinion differentiation and structural differentiation on communication. *Personality and Social Psychology Bulletin*, *48*(7), 1068–1086.

LaCour, M. (2013, August 29–September 1). *A balanced information diet, not echo chambers: Evidence from a direct measure of media exposure* [Paper presentation]. American Political Science Association 2013 Annual Meeting, Chicago, IL, United States.

Latané, B., Liu, J. H., Nowak, A., Bonevento, M., & Zheng, L. (1995). Distance matters: Physical space and social impact. *Personality and Social Psychology Bulletin*, *21*(8), 795–805.

Lawson, R. (1997). The persistence of apocalypticism within a denominationalizing sect: The apocalyptic fringe groups of seventh-day adventism. In T. Robbins & S. J. Palmer (Eds.), *Millennium, Messiahs, and Mayhem: Contemporary apocalyptic movements* (pp. 207–228). Routledge.

Lazer, D., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., & Zittrain, J. L. (2018). The science of fake news. *Science*, *359*(6380), 1094–1096.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, *52*(4), 667–694. https://doi.org/10.2189/asqu.52.4.667

Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*(3), 106–131.

Lewandowsky, S., Pilditch, T., Madsen, J. K., Oreskes, N., & Risbey, J. S. (2019). Seepage and influence: An evidence-resistant minority can affect scientific belief formation and public opinion. *Cognition*, *188*, 124–139.

Mann, R. P. (2022). Collective decision-making under changing social environments among agents adapted to sparse connectivity. *Collective Intelligence*, *1*(2). https://doi.org/10.1177/26339137221121347

Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, *137*, 422–433. https://doi.org/10.1037/a0012798

Massari, G. F., Giannoccaro, I., & Carbone, G. (2019). Are distrust relationships beneficial for group performance? The influence of the scope of distrust on the emergence of collective intelligence. *International Journal of Production Economics*, *208*, 343–355. https://doi.org/10.1016/j.ijpe.2018.12.005

McCoy, J., Rahman, T., & Somer, M. (2018). Polarization and the global crisis of democracy: Common patterns, dynamics, and pernicious consequences for democratic polities. *American Behavioral Scientist*, *62*(1), 16–42.

Medin, D. L., & Atran, S. (2004). The native mind: Biological categorization and reasoning in development and across cultures. *Psychological Review*, *111*(4), 960–983. https://doi.org/10.1037/0033-295X.111.4.960

Nadeau, R., Cloutier, E., & Guay, J.-H. (1993). New evidence about the existence of a bandwagon effect in the opinion formation process. *International Political Science Review*, *14*(2), 203–213.

Pariser, E. (2011). *The filter bubble: What the internet is hiding from you*. Penguin.

Pennycook, G., & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, *25*(5), 388–402.

Pettigrew, T. F., & Tropp, L. R. (2008). How does intergroup contact reduce prejudice? Meta-analytic tests of three mediators. *European Journal of Social Psychology*, *38*(6), 922–934. https://doi.org/10.1002/ejsp.504

Roozenbeek, J., van der Linden, S., Goldberg, B., Rathje, S., & Lewandowsky, S. (2022). Psychological inoculation improves resilience against misinformation on social media. *Science Advances*, *8*(34), Article 6254. https://doi.org/10.1126/sciadv.abo6254

Rumbaut, R. G. (1994). The crucible within: Ethnic identity, self-esteem, and segmented assimilation among children of immigrants. *International Migration Review*, *28*(4), 748–794. https://doi.org/10.1177/019791839402800407

Sang, K., Todd, P. M., Goldstone, R. L., & Hills, T. T. (2020). Simple threshold rules solve explore/exploit trade-offs in a resource accumulation search task. *Cognitive Science*, *44*(2), Article 12817. https://doi.org/10.1111/cogs.12817

Smaldino, P. E. (2017). Models are stupid, and we need more of them. In R. R. Vallacher, S. J. Read, & A. Nowak (Eds.), *Computational social psychology* (pp. 311–331). Routledge.

Smaldino, P. E., Moser, C., Velilla, A. P., & Werling, M. (2022). *Maintaining transient diversity is a general principle for improving collective problem solving*. SocArXiv. https://osf.io/preprints/socarxiv/ykrv5

Spears, R. (2021). Social influence and group identity. *Annual Review of Psychology*, *72*(1), 367–390.

Turner, R. N., Hewstone, M., & Voci, A. (2007). Reducing explicit and implicit outgroup prejudice via direct and extended contact: The mediating role of self-disclosure and intergroup anxiety. *Journal of Personality and Social Psychology*, *93*, 369–388. https://doi.org/10.1037/0022-3514.93.3.369

van der Linden, S. (2022). Misinformation: Susceptibility, spread, and interventions to immunize the public. *Nature Medicine*, *28*(3), 460–467. https://doi.org/10.1038/s41591-022-01713-6

Wellman, B. (2001). Computer networks as social networks. *Science*, *293*(5537), 2031–2034. https://doi.org/10.1126/science.1065547

West, J. D., & Bergstrom, C. T. (2021). Misinformation in and about science. *Proceedings of the National Academy of Sciences, USA*, *118*(15), Article e1912444117. https://www.pnas.org/doi/full/10.1073/pnas.1912444117

Wu, C. M., Schulz, E., Garvert, M. M., Meder, B., & Schuck, N. W. (2020). Similarities and differences in spatial and non-spatial cognitive maps. *PLOS Computational Biology*, *16*(9), Article e1008149. https://doi.org/10.1371/journal.pcbi.1008149

Zollman, K. (2007). The communication structure of epistemic communities. *Philosophy of Science*, *74*(5), 574–587.