



Polarized information ecosystems can reorganize social networks via information cascades

Christopher K. Tokita^{a,1} , Andrew M. Guess^{b,c,2} , and Corina E. Tarnita^{a,2} 

^aDepartment of Ecology and Evolutionary Biology, Princeton University, Princeton, NJ 08544; ^bDepartment of Politics, Princeton University, Princeton, NJ 08544; and ^cSchool of Public and International Affairs, Princeton University, Princeton, NJ 08544

Edited by Robert Axelrod, University of Michigan, Ann Arbor, MI, and approved August 5, 2021 (received for review March 1, 2021)

The precise mechanisms by which the information ecosystem polarizes society remain elusive. Focusing on political sorting in networks, we develop a computational model that examines how social network structure changes when individuals participate in information cascades, evaluate their behavior, and potentially rewire their connections to others as a result. Individuals follow proattitudinal information sources but are more likely to first hear and react to news shared by their social ties and only later evaluate these reactions by direct reference to the coverage of their preferred source. Reactions to news spread through the network via a complex contagion. Following a cascade, individuals who determine that their participation was driven by a subjectively “unimportant” story adjust their social ties to avoid being misled in the future. In our model, this dynamic leads social networks to politically sort when news outlets differentially report on the same topic, even when individuals do not know others’ political identities. Observational follow network data collected on Twitter support this prediction: We find that individuals in more polarized information ecosystems lose cross-ideology social ties at a rate that is higher than predicted by chance. Importantly, our model reveals that these emergent polarized networks are less efficient at diffusing information: Individuals avoid what they believe to be “unimportant” news at the expense of missing out on subjectively “important” news far more frequently. This suggests that “echo chambers”—to the extent that they exist—may not echo so much as silence.

echo chambers | social contagion | political polarization | news media | social media

By standard measures, political polarization in the American mass public is at its highest point in nearly 50 y (1). The consequences of this fundamental and growing societal divide are potentially severe: High levels of polarization reduce policy responsiveness and have been associated with decreased social trust (2), acceptance of and dissemination of misinformation (3), democratic erosion (4), and in extreme cases even violence (5). While policy divides have traditionally been thought to drive political polarization, recent research suggests that political identity may play a stronger role (6, 7). Yet people’s political identities may be increasingly less visible to those around them: Many Americans avoid discussing and engaging with politics and profess disdain for partisanship (8), and identification as “independent” from the two major political parties is higher than at any point since the 1950s (9). Taken together, these conflicting patterns complicate simple narratives about the mechanisms underlying polarization. Indeed, how macrolevel divisions relate to the preferences, perceptions, and interpersonal interactions of individuals remains a significant puzzle.

A solution to this puzzle is particularly elusive given that many Americans, increasingly wary of political disagreement, avoid signaling their politics in discussions and self-presentation and thus lack direct information about the political identities of their social connections (10). However, regardless of individuals’ perceptions about each other, the information ecosystem around them—the collection of news sources available to society—

reflects, at least to some degree, the structural divides of the political and economic system (11, 12). Traditional accounts of media-driven polarization have emphasized a direct mechanism: Individuals are influenced by the news they consume (13) but also tend to consume news from outlets that align with their politics (14, 15), thereby reinforcing their views and shifting them toward the extremes (16, 17). However, large-scale behavioral studies have offered mixed evidence of these mechanisms (18, 19), including evidence that many people encounter a significant amount of counter-attitudinal information online (20–22). Furthermore, instead of directly tuning into news sources, individuals often look to their immediate social networks to guide their attention to the most important issues (23–27). Therefore, it is warranted to investigate how the information ecosystem may impact society beyond direct influence on individual opinions.

Here, we examine media-driven polarization as a social process (28) and propose a mechanism—information cascades—by which a polarized information ecosystem can indirectly polarize society by causing individuals to self-sort into emergent homogeneous social networks even when they do not know others’ political identities. Information cascades, in which individuals observe and adopt the behavior of others, allow the actions of a few individuals to quickly propagate through a social network (29, 30). Found in social systems ranging from fish schools (31) and insect swarms (32) to economic markets (33) and popular culture (29),

Significance

Many argue that partisan media coverage creates political polarization by pushing people’s opinions to extremes, but evidence is mixed. We instead propose that partisan media coverage can cause polarization by altering people’s social connections and reorganizing social networks along political lines. Using computational modeling and social media data, we explore how people may adjust their social ties to avoid the sharing behavior of friends who might be engaging with news from nonpreferred information sources. Our model suggests that polarization is driven to a large extent by unfollowing, which can gradually—and inadvertently—produce homogeneous online networks, known to reduce exposure to challenging information and encourage outgroup hostility. In this way, institutional polarization can reverberate through the networked mass public.

Author contributions: C.K.T., A.M.G., and C.E.T. designed research; C.K.T. and A.M.G. performed research; C.K.T. contributed new reagents/analytic tools; C.K.T. analyzed data; and C.K.T., A.M.G., and C.E.T. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the [PNAS license](#).

¹To whom correspondence may be addressed. Email: christopher.tokita@gmail.com.

²A.M.G. and C.E.T. contributed equally to this work.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2102147118/-DCSupplemental>.

Published December 6, 2021.

information cascades are a widespread social phenomenon that can greatly impact collective behavior such as decision making (34). Online social media platforms are especially prone to information cascades since the primary affordances of these services involve social networking and information sharing (35–38): For example, users often see and share posts of social connections without ever reading the source material (e.g., a shared news article) (39). In addition to altering beliefs and behavior, information cascades can also affect social organization: For instance, retweet cascades on Twitter lead to bursts of unfollowing and following activity (40) that indicate sudden shifts in social connections as a direct result of information spreading through the social network. While research so far has been agnostic as to the content of the information shared during a cascade, it is plausible that information from partisan news outlets could create substantial changes in networks of individuals.

We therefore propose that the interplay between network-altering cascades and an increasingly polarized information ecosystem could result in politically sorted social networks, even in the absence of partisan cues. While we do not argue that this mechanism is the only driver of political polarization—a complex phenomenon likely influenced by several factors—we do argue that the interplay between information and social organization could be one driver that is currently overlooked in discussions of political polarization. We explore this proposition by developing a general theoretical model. After presenting the model, we use Twitter data to probe some of its predictions. Finally, we use the model to explore how the emergence of politically sorted networks might alter information diffusion.

Model Description

Population. We model a population of N individuals living on an undirected network $A = [a_{ij}]$ of average degree $k \ll N$, where $a_{ij} = a_{ji} \in \{0, 1\}$ represents the absence or presence of a social tie between individuals i and j . Each individual has one of two possible fixed political identities, L or R . Importantly, however, individuals do not know the identities of others. Since people in the model are more likely to pay attention to news sources that match their political identity (14, 41), we assume that an individual's identity determines which of two exogenous information sources (e.g., media outlets), M_L or M_R , the individual cares about: L (respectively, R) individuals pay attention only to information source M_L (respectively, M_R). However, we assume that, for various reasons (e.g., inattention, browsing a social media newsfeed), individuals do not always tune directly into the media source, but rather they obtain information (and decide how to react to it) via their social contacts.

In response to information, whether it comes directly from the news source or from a social contact, an individual i can either react and become activated ($x_i = 1$), e.g., express concern or discontent, protest, etc., or not ($x_i = 0$). To decide how to respond, an individual uses an internal, fixed response threshold θ_i . Fixed response thresholds are commonly used to model collective social behavior (42, 43), including cascades (30). Thresholds in our model are drawn at random from the interval $(0, 1)$, to account for heterogeneity among individuals in their propensity to react to information.

Media. Every round, each of the two media sources reports on the same story but with possibly different significance or intensity $s \in (0, 1)$, with values near 0 signaling low importance and values near 1 signaling high importance. Thus, we assume that media sources are reporting on the same topic at the same time but may vary in their respective coverage. This assumption is consistent both with public opinion studies that suggest considerable partisan convergence on the overall importance of top issues

(44) and with empirical evidence of US news coverage suggesting that news sources' topic selections tend to be fairly similar (45). While these counterintuitive findings describe general tendencies, exceptions—for example, lopsided coverage of partisan scandals (46)—illustrate that when intensity of coverage does differ, it can (but does not always) reflect underlying differences in political slant.

The two media sources are correlated by $\gamma \in [-1, 1]$, which gives the average degree of similarity in their coverage, ranging from highly similar ($\gamma \approx 1.0$) to highly dissimilar ($\gamma \approx -1.0$). To simulate the correlated coverage of a story, s_L and s_R are drawn from a multivariate normal distribution $N(0, \Sigma)$, where Σ is the covariance matrix with entries $\Sigma_{12} = \Sigma_{21} = \gamma$ and $\Sigma_{11} = \Sigma_{22} = 1$. We then normalize s_L and s_R on the interval $(0, 1)$ using a cumulative distribution function.

While we discuss the simplest interpretation, in which M_L and M_R each represent an individual news source, research shows that people tend to be news omnivores and consume news from multiple sources (22). Therefore, one could imagine that M_L (respectively, M_R) represents a set of news sources that someone with political identity L (respectively, R) tends to consume and that generally has similar coverage (e.g., *New York Times*, *Washington Post*, and *Los Angeles Times*).

Cascade Dynamics. Every round, a small fraction of randomly chosen individuals sample their corresponding media sources and respond using their thresholds: An individual becomes activated only if the significance of a story exceeds the individual's threshold. Otherwise, the individual remains inactive. Thus, individual behavior in response to media coverage follows the dynamic

$$x_i = \begin{cases} 1 & \text{if } s > \theta_i \\ 0 & \text{if } s \leq \theta_i. \end{cases} \quad [1]$$

Because motivations behind engagements tend to be obscured by social media platforms' aggregation of social cues (47), activation in our model does not assume a particular valence (positive or negative reactions). However, without much loss of generality, one could assume that activation operates via negative emotional reactions, such as anger or outrage, consistent with evidence on negativity bias in news coverage (48, 49) and moral contagion on social media (50, 51).

If individual i is not one of the information samplers that round, individual i responds based on the activity of individual i 's social neighbors: Individual i becomes active only if the fraction of neighbors that are active exceeds individual i 's threshold. This dynamic classically characterizes complex contagions and information cascades (30, 52, 53). Thus, when relying on social information, an individual's behavior can be described by

$$x_i = \begin{cases} 1 & \text{if } \phi_i > \theta_i \\ 0 & \text{if } \phi_i \leq \theta_i, \end{cases} \quad [2]$$

where $\phi_i = \sum a_{ij} x_j / \sum a_{ij}$ is the fraction of i 's neighbors in the active state. Once an individual reacts and becomes active (i.e., $x_i = 1$), the individual remains in the active state for the rest of the round. This follows the method of node behavioral state change commonly used in cascade models (30), whereby individuals who are swept in the cascade remain in that state for the rest of the round. The response cascade(s) initiated by the information samplers is (are) allowed to propagate through the system until it reaches a steady state, i.e., until no additional behavioral changes occur.

Network Adjustment. Once the cascade has reached a steady state, individuals can adjust their social ties. This

assumption—that network changes happen on a much slower timescale than information diffusion—is consistent with empirical data on social media dynamics (40, 54).

First, one active individual—i.e., an individual who was involved in a cascade, including possibly one of the original information samplers—is selected at random to check the individual's reaction against the individual's own news source: If the individual's behavior was inconsistent with the significance level assigned by the individual's own news source (i.e., $s \leq \theta_i$ but the individual overreacted due to social information and became active), the individual randomly breaks one social tie with a neighbor in the active state. The reason for randomly choosing an active neighbor and breaking a tie is that individuals have no other cue about their neighbors: From the perspective of an individual who was caught up in an undesirable cascade, all neighbors who are in the active state had an undesirable response. If the focal individual's reaction was consistent with the news source, nothing happens. In particular, if the focal individual was one of the round's information samplers, then the reaction is correct by default.

Finally, to prevent networks from irreversibly fragmenting as a consequence of tie breaks, we keep the number of links in the network constant: Specifically, if a tie is broken at the end of a

round, a new social tie is added between two randomly picked, unconnected individuals. While this operation keeps the total number of social ties in the network constant, any given individual's social ties (and the individual's position within the network) will change dynamically over time.

Simulations. Each simulation of the model lasted $T = 3.0 \times 10^6$ rounds. We ran simulations across a range of possible information ecosystems γ , and for each value of γ we ran 100 replicate simulations. After each simulation, we assessed how group-level information spread and individual-level information use changed as a result of the network-breaking dynamics resulting from cascades (see *Materials and Methods* for details and see *SI Appendix, Table S1* for parameter settings).

Results and Discussion

The Information Ecosystem Sorts Social Networks along Political Lines.

We find that information cascades cause social networks to become increasingly politically sorted as the information ecosystem becomes less correlated (Fig. 1A), i.e., as news sources diverge in the significance they assign to the same story. Only very highly correlated news coverage—that is, when news sources almost always assign identical importance to the same story—prevents sorting and keeps networks well mixed. Importantly,




Fig. 1. Social networks become politically sorted when news sources are dissimilar in their coverage. Points in graphs represent the average of 100 simulations. (A) Mean political assortativity (\pm SD) of the final $t = T$ social networks as a function of the information ecosystem. *Insets* show example final networks from a single simulation. (B) The average net change in social ties to individuals of the same and the opposite ideology. (C) The average number of ties that were present in the initial $t = 0$ social network and were broken over the course of the simulation.

this political sorting occurs despite the fact that individuals do not know each other's political identity; instead, it is driven entirely by individuals reacting to the behavior of their social neighbors relative to the information that they can directly access from their preferred news source. Since individuals tend to break ties with others they see as acting out of sync with the reality presented by their preferred news source, uncorrelated news coverage quickly causes individuals to experience a net loss of social ties with individuals of the different ideology (Fig. 1B), who are reacting to a diverging news source. The resulting increase in assortativity is thus driven by an increase in tie breaks between individuals of different ideology (SI Appendix, Fig. S1) and an increase in the persistence of ties between individuals of the same ideology (SI Appendix, Fig. S1). Thus, although new ties form randomly in our model, polarized information ecosystems cause ties between individuals of the same political identity to be more likely to persist, creating the appearance of choice homophily—when people choose to connect to similar individuals (55).

Choice homophily is often thought to drive political sorting in networks. However, we find that incorporating a form of choice homophily in our model—by allowing tie additions to be made between two individuals who had the same reaction to the news (see SI Appendix, Fig. S2 legend for details) rather than at random—does not significantly change the emergent pattern of political sorting (SI Appendix, Fig. S24). This robustness check confirms that tie breaking, rather than tie addition, is the main driver of sorting in our model.

Another model modification that is worth exploring is the possibility for desensitization—individuals raising their thresholds as a consequence of being swept up in unwanted cascades—in addition to tie breaks. We find that desensitization still leads to political sorting, albeit less pronounced (SI Appendix, Fig. S3). This outcome owes to the fact that cascades become less impactful faster, as individuals are increasingly unmoved by the news and by the behavior of their neighbors.

Political Sorting Is Not Uniform across Individuals. The information ecosystem also reorganizes social networks according to individuals' propensity to react to information.

To first understand how information cascades reorganize networks in the absence of partisan media coverage, we investigated the scenario when news sources are perfectly identical in their coverage ($\gamma = 1$) and therefore political identity does not matter. Under these conditions, networks remain politically well mixed. Yet, highly reactive low-threshold individuals, who are more likely to be caught up in cascades, sever ties until they end up attached mostly to high-threshold individuals, thereby decreasing the chance of erroneously being swept up in a cascade (SI Appendix, Fig. S4). Meanwhile, high-threshold individuals remain attached to individuals with a wide range of threshold values, since high-threshold individuals are less likely to be swept up in a cascade and therefore are more tolerant of neighbors of varying reactivity. As a result of this dynamic, high-threshold individuals end up with more social ties (i.e., a higher degree) and a more central position in the network relative to low-threshold individuals (Fig. 2).

When media coverage is partisan, the polarized information ecosystem further reorganizes social networks according to individuals' propensity to react to information. When the information ecosystem is increasingly polarized, high-threshold individuals increasingly hold the network together by occupying more central positions and retaining more social ties in the emergent social network (Fig. 2). High-threshold individuals also tend to maintain more politically heterogeneous social ties (i.e., less locally assortative). Thus, by virtue of being less reactive to information, high-threshold individuals give fewer reasons for neighbors to break ties with them and hence maintain more politically diverse social ties, even in increasingly assortative networks. However, in maintaining politically diverse social ties, high-threshold individuals also end up attached to social neighbors that have high thresholds and are therefore less reactive to information (SI Appendix, Fig. S4). Conversely, low-threshold individuals end up with locally assortative social networks, meaning that most or all of their social ties are to people of the same political identity as themselves. Thus, under uncorrelated information conditions, we see the emergence of so-called politically homogeneous "echo chambers," populated primarily by highly reactive (i.e., low-threshold) individuals.

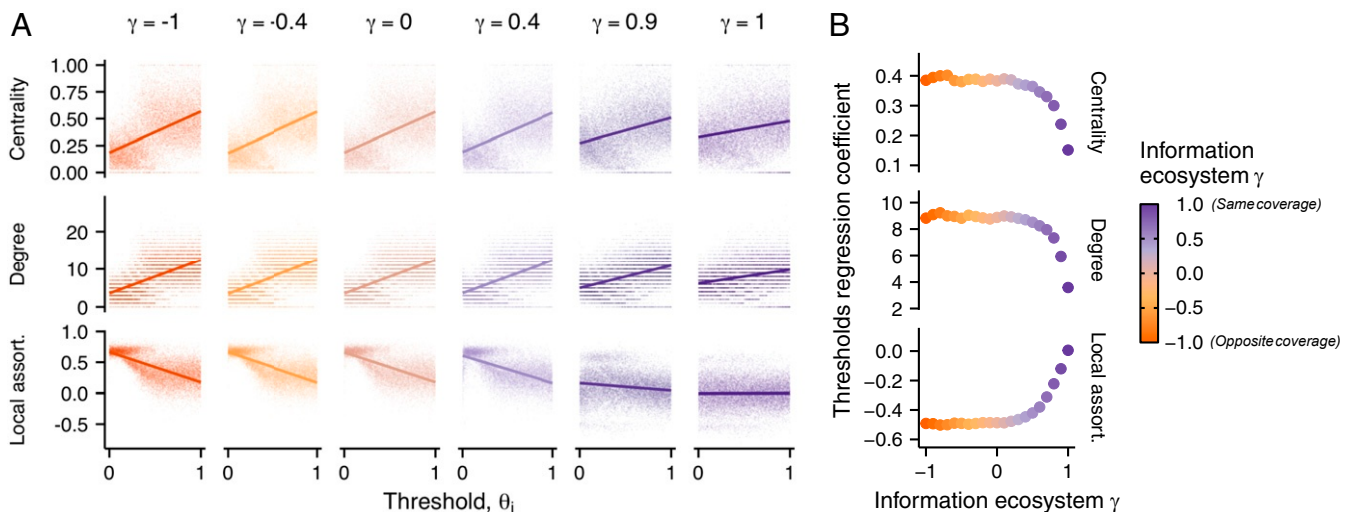


Fig. 2. As the information ecosystem becomes polarized, high-threshold individuals increasingly hold the social network together and bridge the ideological divide, while low-threshold individuals increasingly become isolated in echo chambers. (A) Plots showing the relationship between an individual's threshold θ_i and the individual's position in the final social network under different information ecosystems. Points are individuals across all 100 replicate simulations, and lines are the Bayesian linear regression across all individuals in a particular information ecosystem. (B) Plot showing the regression coefficient for the relationship between threshold and various network metrics. As the information ecosystem becomes polarized, there is an increasingly positive relationship between an individual's threshold value and the individual's centrality or degree in the final network, while conversely there is an increasingly negative relationship between threshold and local assortativity.

Allowing tie additions to occur via choice homophily (details in *SI Appendix*, Fig. S2 legend) rather than at random still leads to the emergence of echo chambers that are primarily populated by low-threshold individuals (*SI Appendix*, Fig. S2B). However, the low-threshold individuals are a bit more well connected and a bit more central than they would otherwise be in our model, albeit still less so than the high-threshold individuals. Thus, choice homophily can prevent very low-threshold individuals from being too socially isolated and might provide insight into how highly reactive individuals can gain a following online.

Polarized Networks Hinder Information Diffusion. Our model allows us to explore how information diffusion—how far information travels and whom it reaches—is impacted by the politically sorted network structure that emerges in polarized information ecosystems. To this end, at the beginning and end of every model simulation, we held the network structure constant and initiated 10,000 cascades (see *Materials and Methods* for details) to determine the difference between the initial and the final network in 1) average cascade size (i.e., average number of individuals that become active from one initially active individual) and 2) cascade bias (i.e., the degree to which a given cascade is concentrated within one political identity grouping).

Using this method, we found two main dynamics of information diffusion. First, cascades become smaller with time, regardless of the level of polarization of the information ecosystem: As individuals adjust their ties to avoid overreacting (which, as discussed above, happens even in a completely unpolarized information landscape), fewer individuals end up reacting to news stories and cascades shrink in size (Fig. 3A). Second, we found that the polarization of the information ecosystem causes cascades to become more biased (i.e., more concentrated within one political identity; Fig. 3B). Thus, in polarized information ecosystems, circulation of news stories becomes confined within an increasingly politically homogeneous segment of society.

At the level of the individual, we found that cascade-driven network adjustments lead individuals to avoid more news that they would deem unimportant (i.e., false positives declined to nearly zero) at the expense of missing out on important news compared to the unadjusted network (Fig. 4A). Thus, echo chambers may limit access to all sources of information, including one's preferred news source. However, in our model, the pattern of decreased information diffusion is not uniform across individuals, with the most reactive, low-threshold individuals seeing the largest decrease in their false positive rates and the largest increase in their false negative rates (Fig. 4B). Meanwhile, high-threshold individuals see only modest changes in their behavior and information access (Fig. 4C).

Twitter Users Lose More Cross-Ideology Ties When Following Polarized News Outlets. To probe the prediction from our model that a polarized information ecosystem results in a higher rate of cross-ideology tie breaks (Fig. 1C), we conducted an observational study on the social media platform Twitter (*SI Appendix*, *SI Methods*). We first chose four news outlets, two of which would represent a high-correlation information ecosystem and two of which would represent a low-correlation information ecosystem. The high-correlation outlets—CBS News and *USA Today*—are large, mainstream outlets known for factual news reporting and thus likely to more closely represent balanced coverage of a story. The low-correlation outlets—Vox and the *Washington Examiner*—are outlets that offer more slant in their content and thus would represent more noticeable deviations from balanced coverage of a story, even as their ideological perspectives are not always explicit. Ideology estimates of 3,000




Fig. 3. Cascades become smaller and more biased in all transformed networks, but uncorrelated news coverage causes cascades to be increasingly concentrated within one ideological group. Points represent the average of 100 replicate simulations means, each calculated as the average value over the 10,000 cascades in that network. (A) Average number of individuals that participate in a cascade originating from one active individual, calculated as the number of active individuals at the end of the round divided by the number of initially active information samplers. (B) Average cascade bias, with higher values indicating that cascades are increasingly composed of mostly one political identity.

random followers of each of the news sources reveal that CBS and *USA Today* have similar and ideologically balanced follower networks, while Vox and the *Washington Examiner* skew more liberal and conservative, respectively (Fig. 5A).

We then monitored the social ties of 1,000 followers of each of the four news sources, focusing exclusively (per the model setup) on liberal followers of CBS News and Vox and conservative followers of *USA Today* and the *Washington Examiner* (*SI Appendix*, *SI Methods*). We pulled the complete follower network of each of our 4,000 monitored users at the beginning and end of a 6-wk period from August to September 2020, allowing us to assess who unfollowed these users over this period of time. Finally, using the initial follower networks of each monitored user, we estimated the ideology of up to 50 random followers to create a baseline for the ideological composition of each user's follower network (Fig. 5B). This allowed us to set a baseline expectation for unfollows: If unfollows were random, then the proportion of unfollows by opposite-ideology users should match the proportion of followers that were of that ideology.

The observational data from Twitter matched our model's prediction that users in low-correlation information ecosystems will lose cross-ideology social ties at a higher rate (Fig. 5C and D and *SI Appendix*, Fig. S5). Followers of low-correlation news outlets lost cross-ideology ties at a higher rate than one would expect with random unfollowing ($BF = 850.56$, $p(\mu > 0) = 1.0$; $t = 3.028$, $P = 0.002$), while followers of high-correlation news outlets lost cross-ideology ties at a rate that was not notably different from random ($BF = 3.4$, $p(\mu > 0) = 0.78$; $t = 0.751$,

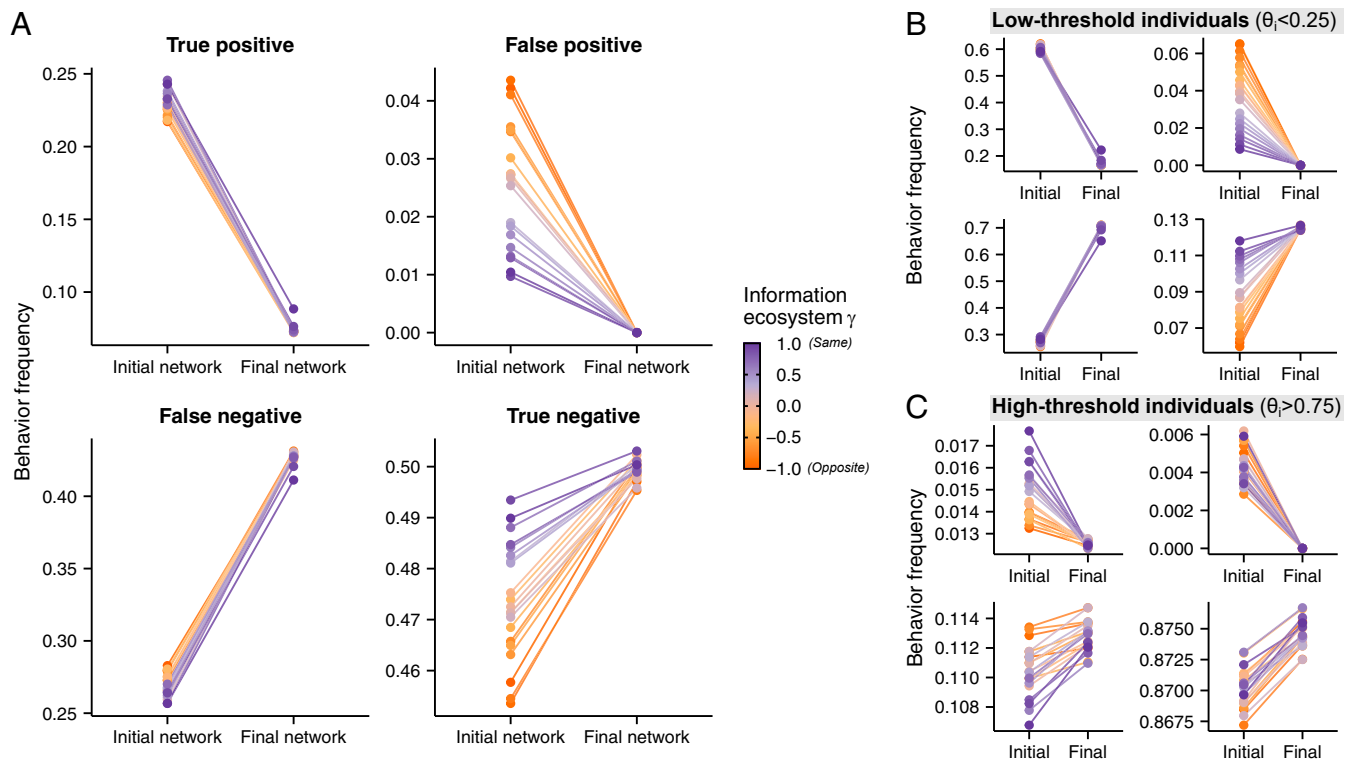


Fig. 4. Transformed social networks allow individuals to avoid reacting to unimportant news at the expense of missing out more often on important news. Points represent the mean of 100 replicate simulations. (A) Relative to initially well-mixed random networks, the final social networks allow individuals to reduce false positive reactions and increase true negative reactions, while also decreasing true positive reactions and increasing false negative reactions. (B and C) Closer analysis reveals that most of this behavioral change occurs among (B) low-threshold individuals, while (C) high-threshold individuals see far less drastic changes in behavior in the transformed social networks.

$P = 0.453$). We estimate that there is strong evidence, specifically a 94.0% chance ($BF = 15.37$; *SI Appendix, Fig. S6*), that low-correlation news followers lost cross-ideology ties at a higher rate than high-correlation news followers; however, this difference is not statistically significant according to conventional frequentist standards (one-tailed t test, $t = 1.55$, $P = 0.06$).

Since the model predicts notable shifts in network structure when comparing highly correlated information ecosystems against even slightly less correlated ones, we estimated the exact position along the information correlation spectrum (i.e., the γ value) for each of our four news sources using natural language processing and the Associated Press as a correlation baseline (*SI Appendix, SI Methods*). We confirm that *USA Today* and CBS News are high-correlation news sources with estimated γ values of 0.923 and 0.884, respectively, while the *Washington Examiner* and Vox are less correlated, with γ values of 0.605 and 0.406, respectively (*SI Appendix, Fig. S7*). Interestingly, the ordinal ranking of estimated γ among these news sources matches the ordinal ranking of cross-ideology unfollows (Fig. 5D and *SI Appendix, Fig. S7*). For example, the most uncorrelated news source is Vox, and its followers experienced the highest rate of cross-ideology unfollows. The most correlated news source is *USA Today*, and its followers experienced the lowest rate of cross-ideology unfollows. The pattern similarly holds for the *Washington Examiner* and CBS News. Moreover, when we consider the broader news diet of each of our monitored users—that is, the mean ideological slant of all news sources that each user follows on Twitter (*SI Appendix, SI Methods*)—we again capture the same ordinal ranking: The followers of Vox and the *Washington Examiner*, who both had the highest rate of cross-ideology unfollows, had the most left-leaning and right-leaning media diets, respectively (*SI Appendix, Fig. S8*). Taken together,

the results of our observational study are consistent with our model's predictions that polarized information ecosystems can lead social networks to sort along political lines.

Conclusion

Politically sorted networks are typically thought to emerge from deliberate action in which users actively create ties with users they know to share the same ideology (56–58) or actively avoid ties with users they know to have opposing ideologies (59). While observing each others' political identities and intentionally forming ties according to similarity are likely one of the drivers of homophilous network formation patterns (60, 61), our results show that knowledge of ideology is not a necessary condition for politically sorted networks to emerge. Instead, the sorting can be driven by individual reactions to a polarized information ecosystem: Individuals who do not want to overreact to news events that they would not deem important will unwittingly self-sort to create politically homogeneous social environments. Twitter data are broadly consistent with this main model prediction: Users in our data who followed more polarized news outlets lost cross-ideology social ties at a rate higher than chance, suggesting that they are self-sorting into more politically homogenous social networks.* Our results, therefore, complement alternative explanations for social network polarization and are especially relevant when the ability to observe identity signals does not guarantee accurate perceptions (62). More broadly, by demonstrating that the structure of the information ecosystem can shape society in its image, our results show how deep-seated

*Although our estimated effect size is small, this is consistent with a low base rate of unfollowing, which research shows is far less common than social tie formation (40).

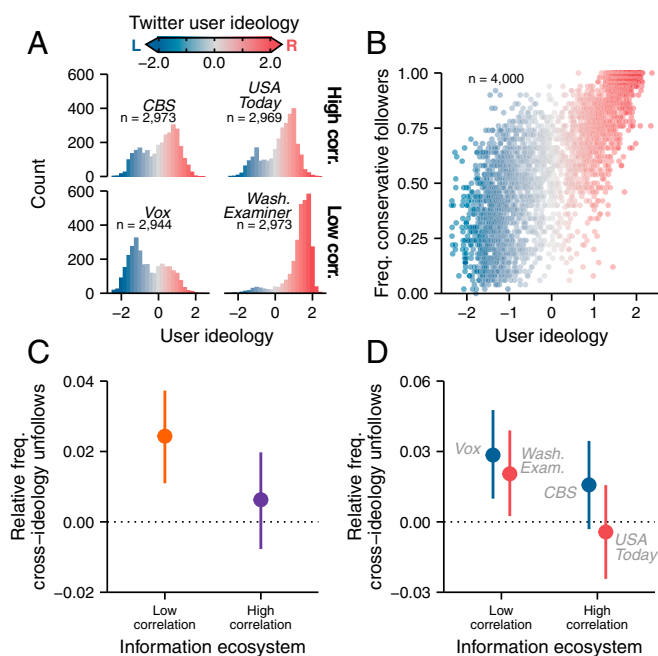


Fig. 5. Observational study on Twitter of networks of news followers reveals that individuals who follow low-correlation (i.e., more partisan) news outlets lose cross-ideology followers at a faster rate than expected by random chance. (A) Ideological distribution of 12,000 random followers of the two high-correlation and two low-correlation news outlets. (B) Plot showing the ideology of the 4,000 monitored news followers and the frequency of conservative users in each of their follow networks. The positive relationship shows evidence of ideologically sorted networks on Twitter. (C) The estimated average relative frequency of cross-ideology unfollows ($\pm 90\%$ credible interval) broken out by information ecosystem. Positive values indicate that cross-ideology unfollows are happening at a rate higher than one would expect with random unfollowing. (D) The estimated average relative frequency of cross-ideology unfollows ($\pm 90\%$ credible interval) broken out by news outlet.

institutional polarization can propagate to the mass level, providing a link between patterns of polarization at the micro- and macrolevels.

Further investigation of how this emergent sorting influences the flow of information reveals an alternative perspective on why politically sorted networks are problematic. The most common existing perspectives view politically sorted networks either as echo chambers—in which individuals are vulnerable to confirmation bias via repeated amplification of the same ideas (59, 63)—or as “epistemic bubbles”—in which individuals are prevented from being exposed to cross-ideological content and to a broader range of viewpoints (64). Unexpectedly, our results stand in contrast to both of these viewpoints. We find that individuals in politically sorted networks are able to avoid the amplification of news that they deem unimportant but at the cost of missing out on news that they deem important, i.e., news from their preferred news source. In other words, we highlight an underappreciated consequence of echo chambers: They might silence more than they echo.

Political sorting was not uniform but was instead concentrated among the lowest-threshold individuals, who were most likely to be caught up in a cascade and therefore have the opportunity to adjust their social networks. If thresholds are a good proxy for sensitivity to or interest in news—as some have argued based on observations that individuals who are highly interested in news and politics are also more likely to share articles online (65)—then our findings are consistent with empirical research showing echo chambers to be made up primarily of individuals who are

highly interested in news and politics (21, 22, 66–68). By sorting the lowest-threshold individuals into echo chambers, a polarized media landscape can indirectly facilitate the further polarization of opinions and entrenchment of political identities that come from interacting with like-minded individuals (56, 69, 70). As a result, low-threshold individuals are the most likely to develop extreme political opinions; but, in light of our finding that echo chambers can be silencers, low-threshold individuals are also the most vulnerable to being less informed than they think they are (71).

Our findings suggest an underexplored consequence of the rise of online misinformation or “fake news”: Rather than creating false beliefs, the power of misinformation may lie more specifically in further isolating consumers of misinformation from the broader society. Since the topics and emphasis of misinformation will often be highly dissimilar from mainstream news coverage, users who heavily consume misinformation will be in an uncorrelated (or perhaps even anticorrelated) information ecosystem relative to individuals who consume only mainstream media. Thus, consistent with analyses of fake news diffusion on Twitter (72), our model predicts a sorting of social networks between those who do and those who do not regularly consume misinformation. If misinformation in the United States is not evenly consumed across the political spectrum but is instead concentrated among right-leaning users (73, 74), misinformation alone—even without obvious political cues—might further exacerbate social separation between liberals and conservatives.

Finally, our results can serve as a benchmark for studying the consequences of platform-specific affordances and algorithms that have been identified as culprits in society’s growing polarization (75). Social-matching algorithms could reinforce the political sorting predicted in our model by connecting similar individuals, for instance (76). The reinforcing potential of algorithms may be especially pronounced in situations of attention scarcity (77): For example, even though limited attention is predicted to decrease cascades (78) and therefore political sorting, technological attempts to maximize this limited attention via preference-driven ranking algorithms—presenting users with tailored, seemingly high-importance information (79)—might have the opposite effect. Such personalized news feeds would decrease the correlation of the media ecosystem overall and, according to our model, exacerbate political sorting. Our approach can be extended in this way to make counterfactual predictions by simulating the effect of additional elements of complexity, such as algorithms, as well as potential policy changes. By employing the tools of computational social science, we demonstrate the potential for interdisciplinary collaboration to shed light on the impacts of modern communication technology on collective behavior (38).

Materials and Methods

Assessing Information Diffusion in Networks. For each simulated network, we ran 10,000 cascades on the initial and final network while holding network ties constant (i.e., we forgo the network-adjusting step of the model). By then comparing the cascades on the initial network against cascades on the final network, we could assess how cascade-driven network adjustments are altering information spread.

To assess how many individuals became active from one individual reacting to the news, average cascade size was then calculated by dividing total cascade activity by the number of initially active individuals. To assess the degree to which news was spread exclusively among one type of individual, cascade bias was calculated as the absolute difference in the proportion of a cascade that was made up of individuals of one type or the other, i.e., $|X_L - X_R|/X$, where X_L and X_R are the number of active individuals of type L and type R , respectively.

To monitor individual-level information use, we compared the behavior of individuals with their threshold and their respective news source. This allowed us to assess which proportion of messages that an individual would want to receive (i.e., $s > \theta_i$) were effectively received and followed (i.e.,

$x_i = 1$). We could call this a true positive case from the perspective of the individual. Similarly, we could ask how many times individuals acted on a false positive; that is, they acted on a message they would not have otherwise deemed important (i.e., $s \leq \theta_i$). We could similarly assess instances of true negatives and false negatives for each individual.

In all of the above metrics, to create an average across all 100 simulated networks of a given information ecosystem γ , we 1) averaged the 10,000 cascades for each simulated network and then 2) calculated the average of these 100 within-network averages.

Measuring Network Structure. To measure the general connectivity of an individual, we calculated the individual's degree, which is simply the number of social ties possessed by an individual. Larger numbers indicate that an individual is attached to more people.

To measure the prominence of an individual, we calculated the individual's centrality using eigenvector centrality, which measures the general position of a node by accounting for the individual's direct and indirect connections in the overall network (80). Higher values indicate that an individual occupies a prominent, central position in the network, meaning that the individual is connected to many different portions of the social network.

To measure the overall degree of political sorting, we calculated assortativity, which is the tendency of nodes to attach to other nodes that are

similar in trait (i.e., political identity) (81). This metric takes values in the range $[-1, 1]$, with positive values indicating that individuals tend to be attached to others of the same ideology and negative values indicating that individuals tend to be attached to others of different ideologies. Typically, assortativity provides a global measurement at the level of the entire social network, but we also calculated local assortativity to measure the degree to which individuals are embedded in politically homogeneous neighborhoods in the social network (82). When calculating the neighborhood size for local assortativity, we used $\alpha = 0.5$, which is midway between calculating local assortativity for only immediate social connections ($\alpha = 0$) and calculating global assortativity ($\alpha = 1$) (see ref. 82 for more details).

Data Availability. Code and simulation data have been deposited in Github (<https://github.com/christokita/information-cascades>) (83). Data from the observational study has been deposited in Zenodo (<https://doi.org/10.5281/zenodo.5277188>) (84).

ACKNOWLEDGMENTS. We are grateful to Richard A. Bonneau and Merlijn Staps for helpful early suggestions on empirical approaches to test the model and network metrics, respectively. We thank Yphthach Lelkes for helpful feedback on an earlier draft of this paper. This work was supported by the NSF Graduate Research Fellowship under Grant DGE1656466 (to C.K.T.) and a research grant from the Princeton Data-Driven Social Science Initiative.

1. L. Boxell, M. Gentzkow, J. M. Shapiro, Greater Internet use is not associated with faster growth in political polarization among US demographic groups. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 10612–10617 (2017).
2. M. J. Hetherington, T. J. Rudolph, *Why Washington Won't Work: Polarization, Political Trust, and the Governing Crisis* (University of Chicago Press, 2015), vol. 104.
3. M. Osmundsen, A. Bor, P. B. Vahlstrup, A. Bechmann, M. B. Petersen, Partisan polarization is the primary psychological motivation behind "fake news" sharing on Twitter. *Am. Polit. Sci. Rev.* **115**, 999–1015 (2021).
4. S. Levitsky, D. Ziblatt, *How Democracies Die* (Broadway Books, 2018).
5. N. Kalmoe, *With Ballots and Bullets: Partisanship and Violence in the American Civil War* (Cambridge University Press, 2020).
6. L. Mason, *Uncivil Agreement: How Politics Became Our Identity* (University of Chicago Press, 2018).
7. N. Dias, Y. Lelkes, The nature of affective polarization: Disentangling policy disagreement from partisan identity. *Am. J. Pol. Sci.* (2021).
8. S. Klar, Y. Krupnikov, J. B. Ryan, Affective polarization or partisan disdain? Untangling a dislike for the opposing party from a dislike of partisanship. *Public Opin. Q.* **82**, 379–390 (2018).
9. J. M. Jones, Americans continue to embrace political independence. *Gallup*, 7 January 2019. <https://news.gallup.com/poll/245801/americans-continue-embrace-political-independence.aspx>. Accessed 24 February 2021.
10. J. E. Settle, T. N. Carlson, Opting out of political discussions. *Polit. Commun.* **36**, 476–496 (2019).
11. Y. Benkler, R. Faris, H. Roberts, *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics* (Oxford University Press, 2018).
12. N. McCarty, K. T. Poole, H. Rosenthal, *Polarized America: The Dance of Ideology and Unequal Riches* (MIT Press, 2016).
13. C. R. Sunstein, *Republic: Divided Democracy in the Age of Social Media* (Princeton University Press, 2018).
14. S. Iyengar, K. S. Hahn, Red media, blue media: Evidence of ideological selectivity in media use. *J. Commun.* **59**, 19–39 (2009).
15. N. J. Stroud, Polarization and partisan selective exposure. *J. Commun.* **60**, 556–576 (2010).
16. M. Levendusky, *How Partisan Media Polarize America* (University of Chicago Press, 2013).
17. K. Arceneaux, M. Johnson, *Changing Minds or Changing Channels?: Partisan News in an Age of Choice* (University of Chicago Press, 2013).
18. C. A. Bail et al., Exposure to opposing views on social media can increase political polarization. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9216–9221 (2018).
19. R. Levy, Social media, news consumption, and polarization: Evidence from a field experiment. *Am. Econ. Rev.* **111**, pp. 831–870 (2021).
20. M. Gentzkow, J. M. Shapiro, Ideological segregation online and offline. *Q. J. Econ.* **126**, 1799–1839 (2011).
21. G. Eady, J. Nagler, A. Guess, J. Ziliinsky, J. A. Tucker, How many people live in political bubbles on social media? Evidence from linked survey and Twitter data. *SAGE Open* **9**, 215824401983270 (2019).
22. A.M. Guess, (Almost) everything in moderation: New evidence on Americans' online media diets. *Am. J. Pol. Sci.*, 10.1111/ajps.12589 (2021).
23. N. Kligler-Vilenchik, A. Hermidá, S. Valenzuela, M. Villi, Studying incidental news: Antecedents, dynamics and implications. *Journalism* **21**, pp. 1025–1030 (2020).
24. D. C. Mutz, L. Young, Communication and public opinion: Plus ça change? *Public Opin. Q.* **75**, 1018–1044 (2011).
25. K. Thorson, C. Wells, Curated flows: A framework for mapping media exposure in the digital age. *Commun. Theory* **26**, 309–328 (2016).
26. E. Katz, The two-step flow of communication: An up-to-date report on an hypothesis. *Public Opin. Q.* **21**, 61–78 (1957).
27. D. Guilbeault, J. Becker, D. Centola, Social learning and partisan bias in the interpretation of climate trends. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9714–9719 (2018).
28. P. F. Lazarsfeld, R. K. Merton, "Friendship as a social process: A substantive and methodological analysis" in *Freedom and Control in Modern Society*, M. A. T. Berger, C. H. Page, Eds. (Van Nostrand, New York, NY, 1954), vol. 18, pp. 18–66.
29. S. Bikhchandani, D. Hirshleifer, I. Welch, A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* **100**, 992–1026 (1992).
30. D. J. Watts, A simple model of global cascades on random networks. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 5766–5771 (2002).
31. S. B. Rosenthal, C. R. Twomey, A. T. Hartnett, H. S. Wu, I. D. Couzin, Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 4690–4695 (2015).
32. O. L. S. Michel, J. L. Deneubourg, G. Sempo, Information cascade ruling the fleeing behaviour of a gregarious insect. *Anim. Behav.* **85**, 1271–1285 (2013).
33. D. Hirshleifer, S. H. Teoh, Herd behaviour and cascading in capital markets: A review and synthesis. *Eur. Financ. Manag.* **9**, 25–66 (2003).
34. L. Conradt, C. List, Group decisions in humans and animals: A survey. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **364**, 719–742 (2009).
35. S. Goel, D. J. Watts, D. G. Goldstein, The structure of online diffusion networks. *Proc. 13th ACM Conf. Electron. Commer.* **1**, 623–638 (2012).
36. E. Sun, I. Rosenn, C. A. Marlow, T. M. Lento, Gesundheit! Modeling contagion through Facebook news feed mechanics of Facebook page diffusion. *Proc. 3rd Int. ICWSM Conf.* **3**, 146–153 (2009).
37. H. Kwak, C. Lee, H. Park, S. Moon, "What is Twitter, a social network or a news media?" in *Proceedings of the 19th International Conference on World Wide Web, WWW '10* (Association for Computing Machinery, New York, NY, 2010), pp. 591–600.
38. J. B. Bak-Coleman et al., Stewardship of global collective behavior. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2025764118 (2021).
39. M. Gabelkov, A. Ramachandran, A. Chaintreau, A. Legout, "Social clicks: What and who gets read on Twitter?" in *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science* (Association for Computing Machinery, New York, NY, 2016), pp. 179–192.
40. S. Myers, J. Leskovec, "The bursty dynamics of the twitter information network" in *WWW 2014 - Proceedings of the 23rd International Conference on World Wide Web* (Association for Computing Machinery, Inc., New York, NY, 2014), pp. 913–923.
41. E. Peterson, S. Goel, S. Iyengar, Partisan selective exposure in online news consumption: Evidence from the 2016 presidential campaign. *Polit. Sci. Res. Methods* **9**, 242–258 (2019).
42. M. Granovetter, Threshold models of collective behavior. *Am. J. Sociol.* **83**, 1420–1443 (1978).
43. S. N. Beshers, J. H. Fewell, Models of division of labor in social insects. *Annu. Rev. Entomol.* **46**, 413–440 (2001).
44. Election 2020: Voters are highly engaged, but nearly half expect to have difficulties voting (Pew Research Center, 2020). <https://www.pewresearch.org/politics/2020/08/13/important-issues-in-the-2020-election/>. Accessed 10 June 2021.
45. C. Budak, S. Goel, J. M. Rao, Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opin. Q.* **80** (suppl. 1), 250–271 (2016).
46. R. Puglisi, J. M. Snyder Jr., Newspaper coverage of political scandals. *J. Polit.* **73**, 931–950 (2011).
47. S. Messing, S. J. Westwood, Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online. *Commun. Res.* **41**, 1042–1063 (2014).

48. S. Knobloch-Westernwick, C. Mothes, N. Polavin, Confirmation bias, ingroup bias, and negativity bias in selective exposure to political information. *Communic. Res.* **47**, 104–124 (2020).
49. S. Soroka, P. Fournier, L. Nir, Cross-national evidence of a negativity bias in psychophysiological reactions to news. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 18888–18892 (2019).
50. W. J. Brady, J. A. Wills, J. T. Jost, J. A. Tucker, J. J. Van Bavel, Emotion shapes the diffusion of moralized content in social networks. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 7313–7318 (2017).
51. S. Rathje, J. J. Van Bavel, S. van der Linden, Out-group animosity drives engagement on social media. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2024292118 (2021).
52. D. Centola, M. Macy, Complex contagions and the weakness of long ties. *Am. J. Sociol.* **113**, 702–734 (2007).
53. D. Centola, The spread of behavior in an online social network experiment. *Science* **329**, 1194–1197 (2010).
54. S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online. *Science* **359**, 1146–1151 (2018).
55. J. M. McPherson, L. Smith-Lovin, Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. *Am. Sociol. Rev.* **52**, 370–379 (1987).
56. D. Baldassarri, P. Bearman, Dynamics of political polarization. *Am. Sociol. Rev.* **72**, 784–811 (2007).
57. M. Mäs, A. Flache, D. Helbing, Individualization as driving force of clustering phenomena in humans. *PLoS Comput. Biol.* **6**, e1000959 (2010).
58. M. Mäs, A. Flache, Differentiation without distancing. Explaining bi-polarization of opinions without negative influence. *PLoS One* **8**, e74516 (2013).
59. C. R. Sunstein, *Republic.com* (Princeton University Press, 2001).
60. P. Bogdanov, M. Busch, J. Moehlis, A. K. Singh, B. K. Szymanski, Modeling individual topic-specific behavior and influence backbone networks in social media. *Soc. Netw. Anal. Min.* **4**, 204 (2014).
61. X. Lu, B. K. Szymanski, Scalable prediction of global online media news virality. *IEEE Trans. Comput. Soc. Syst.* **5**, 858–870 (2018).
62. S. Goel, W. Mason, D. J. Watts, Real and perceived attitude agreement in social networks. *J. Pers. Soc. Psychol.* **99**, 611–621 (2010).
63. L. A. Adamic, N. Glance, “The political blogosphere and the 2004 U.S. Election: Divided they blog” in *3rd International Workshop on Link Discovery, LinkKDD 2005 - in conjunction with 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Association for Computing Machinery, New York, NY, 2005).
64. E. Pariser, *The Filter Bubble: What the Internet is Hiding From You* (Penguin UK, 2011).
65. A. Kalogeropoulos, S. Negro, I. Picone, R. K. Nielsen, Who shares and comments on news?: A cross-national comparative analysis of online and social media participation. *Soc. Media Soc.* **3**, 205630511773575 (2017).
66. A. Boutyline, R. Willer, The social structure of political echo chambers: Variation in ideological homophily in online networks. *Polit. Psychol.* **38**, 551–569 (2017).
67. P. Barberá, “Social media, echo chambers, and political polarization” in *Social Media and Democracy: The State of the Field, Prospects for Reform*, N. Persily, J. A. Tucker, Eds. (Cambridge University Press, 2020), pp. 34–55.
68. J. Shore, J. Baek, C. Dellarocas, Network structure and patterns of information diversity on Twitter. *MIS Q.* **42**, 849–872 (2018).
69. C. R. Sunstein, The law of group polarization. *J. Polit. Philos.* **10**, 175–195 (2002).
70. C. K. Tokita, C. E. Tarnita, Social influence and interaction bias can drive emergent behavioural specialization and modular social networks across systems. *J. R. Soc. Interface* **17**, 20190564 (2020).
71. B. A. Lyons, J. M. Montgomery, A. M. Guess, B. Nyhan, J. Reifler, Overconfidence in news judgments is associated with false news susceptibility. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2019527118 (2021).
72. A. Bovet, H. A. Makse, Influence of fake news in Twitter during the 2016 US presidential election. *Nat. Commun.* **10**, 7 (2019).
73. N. Grinberg, K. Joseph, L. Friedland, B. Swire-Thompson, D. Lazer, Fake news on Twitter during the 2016 U.S. presidential election. *Science* **363**, 374–378 (2019).
74. A. M. Guess, B. Nyhan, J. Reifler, Exposure to untrustworthy websites in the 2016 US election. *Nat. Hum. Behav.* **4**, 472–480 (2020).
75. Z. Tufekci, Youtube, the great radicalizer. *NY Times*, 10 March 2018. <https://www.nytimes.com/2018/03/10/opinion/sunday/youtube-politics-radical.html>. Accessed 22 June 2021.
76. F. P. Santos, Y. Lelkes, S. A. Levin, Link recommendation algorithms and dynamics of polarization in online social networks. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2102141118 (2021).
77. J. G. Webster, T. B. Ksiazek, The dynamics of audience fragmentation: Public attention in an age of digital media. *J. Commun.* **62**, 39–56 (2012).
78. S. Sreenivasan, K. S. Chan, A. Swami, G. Korniss, B. K. Szymanski, Information cascades in feed-based networks of users with limited attention. *IEEE Trans. Netw. Sci. Eng.* **4**, 120–128 (2017).
79. T. Abdelzaher *et al.*, The paradox of information access: Growing isolation in the age of sharing. arXiv [Preprint] (2020). <https://arxiv.org/abs/2004.01967> (Accessed 7 September 2021).
80. M. E. J. Newman, “The mathematics of networks” in *The New Palgrave Dictionary of Economics*, M. Vernengo, E. Perez Caldentey, B. J. Rosser Jr, Eds. (Springer, 2016).
81. M. E. J. Newman, Mixing patterns in networks. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **67**, 026126 (2003).
82. L. Peel, J. C. Delvenne, R. Lambiotte, Multiscale mixing patterns in networks. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 4057–4062 (2018).
83. C. K. Tokita, A. M. Guess, C. E. Tarnita, Information cascade model of politically polarized social networks. Github. <https://github.com/christokita/information-cascades>. Deposited 26 August 2021.
84. C. K. Tokita, Observational data for: Polarized information ecosystems can reorganize social networks via information cascades. Zenodo. <https://doi.org/10.5281/zenodo.5277188>. Deposited 26 August 2021.