

University of Groningen

Small Worlds and Cultural Polarization

Flache, Andreas; Macy, Michael W.

Published in:
Journal of Mathematical Sociology

DOI:
[10.1080/0022250X.2010.532261](https://doi.org/10.1080/0022250X.2010.532261)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2011

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Flache, A., & Macy, M. W. (2011). Small Worlds and Cultural Polarization. *Journal of Mathematical Sociology*, 35(1-3), 146-176. <https://doi.org/10.1080/0022250X.2010.532261>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Small Worlds and Cultural Polarization

Andreas Flache

Department of Sociology–ICS, University of Groningen, Groningen,
The Netherlands

Michael W. Macy

Department of Sociology, Cornell University, Ithaca, New York, USA

Building on Granovetter’s theory of the “strength of weak ties,” research on “small-world” networks suggests that bridges between clusters in a social network (long-range ties) promote cultural diffusion, homogeneity, and integration. We show that this macro-level implication of network structure depends on hidden micro-level assumptions. Using a computational model similar to earlier studies, we find that ties between clusters facilitate cultural convergence under the micro-level assumptions of assimilation and attraction to similar others. However, these assumptions also have negative counterparts—differentiation and xenophobia. We found that when these negative possibilities are no longer assumed away, the effect of long-range ties reverses: Even very small amounts of contact between highly clustered communities sharply increased polarization at the population level.

[An appendix to this article is featured as an online supplement at the publisher’s website.]

Keywords: agent-based models, cultural dynamics, micro-macro link, polarization, small-world networks, social influence

The research of the first author has been supported by the Netherlands Organization for Scientific Research, NWO (VIDI Grant 452-04-351). We thank the special issue editors and two anonymous reviewers for valuable advice on and constructive criticism of earlier versions of this work. Our work has further benefited from stimulating discussions with Michael Mäs, Tobias Stark, Károly Takács, and other members of the discussion group on norms and networks at the Department of Sociology of the University of Groningen.

Address correspondence to Andreas Flache, Faculty of Behavioral and Social Sciences, Department of Sociology–ICS, University of Groningen, Grote Rozenstraat 31, 9712 TG Groningen, The Netherlands. E-mail: a.flache@rug.nl

1. INTRODUCTION

Recent decades have seen a dramatic increase in the possibilities for cultural contact between people located in widely disparate social milieus and geographical locations around the world. In work organizations, employees increasingly work in multinational teams, with culturally diverse and geographically dispersed members, often collaborating via the Internet (Early and Mosakowski, 2000; Kirkman, Rosen, Gibson, Tesluk, and McPherson, 2002). We live in an ever smaller world due to increased social and spatial mobility and more powerful communication technologies, particularly those based on the Internet. This has led observers to ask whether social and cultural differences at the global scale will survive increasing exposure to outside influences (see Greig, 2002; Friedman, 2005). At the same time, studies also found indications for persistent dissensus and lack of cooperation despite more intercultural interaction, for example, in multinational (Early and Mosakowski, 2000) and so-called “virtual” teamwork (Kirkman et al., 2002).

Despite the march toward an ever smaller world, studies of social segregation (Massey and Denton, 1998) and homophily (McPherson, Smith-Lovin, and Cook, 2001) suggest that, within social networks, ties remain clustered in closely knit local networks based on geographic, social, or organizational proximity. Watts and Strogatz (1998) showed how clustering can co-exist with greater connectedness. In “small-world” networks, only a small proportion of people need to have “long-range” ties in order to dramatically decrease network distances while preserving local clustering. The range of a tie is the length of the shortest path between the adjacent nodes after removal of the tie (Centola and Macy, 2007). Whether the range of a tie is “long” or “short” depends on the overall configuration of the network. Generally, a “long-range tie” connects local clusters in the network that are not directly linked otherwise.¹ Only a few long-range ties between otherwise distant clusters can greatly accelerate the spread of information.

Early formal models of social influence suggest that long-range ties may foster global consensus. Building on French (1956), Abelson (1964) and Harary (1959) proved that convergence to unanimity

¹For example, in the “connected caveman graph” shown in Figure 8a, the local clusters are formed by so-called “caves,” complete subnetworks with five nodes each. Each cave is directly linked to its two neighboring caves in the circular spatial arrangement by exactly one tie but not to any other cave. Thus, in this network a long-range tie links nodes from two caves that are not directly linked in the original graph. In other words, a long-range tie has a range of at least three in this graph.

was inevitable in a “connected” network in which there is a path between every pair of nodes. Later, Friedkin and Johnsen (1990, 1999) added that agents under outside influence always retain some residue of their original traits. However, these models still imply that influence greatly reduces differences, typically resulting in little diversity in connected networks (see Friedkin, 2001). Granovetter’s (1973) theory of the “strength of weak ties” similarly suggests that bridges between clusters in a social network promote homogeneity and integration.

Other studies point in the opposite direction. Centola and Macy (2007) showed that, holding density constant, long-range ties can actually inhibit and even prevent the spread of social contagions that require reinforcement from multiple adopting neighbors. Baldassarri and Bearman (2007) point to evidence that segregation by income, social class, and race in the United States may even have become stronger in recent decades. Mark (1998, 2003) observes clustering of cultural preferences in age and education, and Glaeser and Ward (2006) find that political and social views in the United States are strongly segregated in geographical space. Empirical research on the use of the Internet by extremist minorities (Adams and Roscigno, 2007) resonates theoretical predictions (e.g., Greig, 2002; Shibanaï, Yasuno, and Ishiguro, 2001) that a global range of interaction may allow local minorities to find others who support resistance against pressures to conform to local majorities. When this happens, local minorities may no longer assimilate to the cultural profile of the local majority. In this view, long-range ties that connect otherwise socially distant clusters may help to preserve social diversity.

Thus, it remains an open question whether and under what conditions local social influences in a small-world network can sustain cultural diversity against pressures toward assimilation to global influences that are mediated by long-range ties. In this article, we explore theoretically effects of long-range ties on cultural diversity in an otherwise locally clustered social network. Our formal models may be seen as abstract representation of increasing connectedness in realms such as intercultural contact on a global scale (see further Greig, 2002; Friedman, 2005) or interaction between employees within multinational or virtual organizations (Early and Mosakowski, 2000; Kirkman et al., 2002).

To explain persistent diversity, Axelrod (1997) proposed a model that combines social influence with another elementary social mechanism, social selection. Influence is defined as the tendency to alter one’s opinions, attitudes, beliefs, customs, or other cultural traits to more closely resemble those of a cultural majority, a high-status

minority, or a network neighbor with whom one regularly interacts (Festinger, Schachter, and Back, 1950; see Axelrod, 1997, for an overview). Selection is defined as the choice of interaction partners. Simply put, whereas influence refers to the ability to change the beliefs and opinions of interaction partners, selection refers to the ability to alter the interaction partner with whom one interacts. Selection is often theorized as *homophily*, the attraction toward similar others, or the idea that “birds of a feather flock together” (Lazarsfeld and Merton, 1954; McPherson et al., 2001; for similar models, see also Mark, 1998, 2003; Kitts, Macy, and Flache, 1999; Latané and Nowak, 1997). Axelrod showed how homophily can preserve diversity despite the convergent tendencies created by cultural influence. A preference for interaction with those who are similar can preclude influence once the differences between cultural groups become too large. Buskens, Corten, and Weesie (2008) obtained a similar result using a game-theoretical model that combines a coordination game with social selection based on homophily. However, Axelrod also found that a larger range of interaction renders coordination on unique local cultures increasingly unlikely (see Parisi, Cecconi, and Natale, 2003).

In sum, empirical and theoretical research suggests two very different effects of long-range ties in clustered social networks. An earlier generation of studies predicted the inevitable collapse of diversity into an emergent monoculture unless clusters are entirely disconnected, while more recent studies suggest the preservation and even the enhancement of diversity in an increasingly connected world.

There is also a third possibility overlooked by most previous research: long-range ties can reduce cultural diversity and at the same time deepen cultural divisions, leading not to consensus or cultural homogeneity but to polarization. By “polarization” we mean that a population divides into a small number of factions with high internal consensus and sharp disagreement between them. A perfectly polarized population contains two opposing factions whose members agree on everything with each other and fully disagree on everything with the out-group.

We use an agent-based computational model to show that the macro-level effects of long-range ties on cultural consensus and polarization depend on a key micro-level assumption about the valence of social interaction. Interaction encompasses both social influence and selection. “Valence” refers to whether social influence and selection are positive or negative. Positive influence and selection correspond to assimilation with and attraction to those who

are similar. However, assimilation and homophily also have negative counterparts—differentiation from and xenophobia toward those who are different. Almost all research on cultural dissemination, the diffusion of innovation, the assimilation of immigrant minorities, the evolution of norms, and even the spread of social protest have assumed away the possibility that social interaction can have a negative valence. A small but growing number of studies, in contrast, have explored the polarizing effects of bivalent interaction (Mark, 2003; Macy, Kitts, Flache, and Benard, 2003; Jager and Amblard, 2005; Kitts, 2006; Baldassarri and Bearman, 2007; and Flache and Mäs, 2008a, 2008b).

The present research contributes to this growing body of knowledge by focusing on the effects of long-range ties. We argue that long-range ties can foster cultural polarization. If agents connected via short-range ties are more similar than those connected by long-range ties, differentiation and xenophobia are more likely to occur in interactions between agents within long-range connections, compared to local interactions between agents within short-range ties. The results show how the macro-consequence of these micro-processes can be polarization.

In two computational experiments, we manipulate both the micro-level assumptions about positive and negative valence of interaction and structural assumptions about the proportion of long-range ties in the access network. The access network limits opportunities for interaction and is modeled as an exogenously imposed, static neighborhood structure that defines the structural constraints within which the dynamics of social influence and selection operate. The access network may result from geographical, cultural or social constraints that are not under control of the agents. Selection within these structural constraints then generates a dynamic network that is a subset of the possibilities for interaction in the access graph.

In both experiments, we start with an ideal-typical representation of a small-world access network in which we assume a large number of small clusters, with maximal network density² at the local level, that is, within the clusters. We then test the robustness of our results by manipulating independently the number and density of clusters into which the access network is segmented.

²Following the standard in the literature, we define network density of a cluster as the number of actual neighborhood ties between members of the cluster divided by the number of theoretically possible ties.

A second robustness test focuses on the number of salient cultural dimensions that actors attend to. This can be seen as an indicator of the “broad mindedness” of social actors (Macy et al., 2003) or of a society’s “cultural complexity” (Axelrod, 1997). Intuitively, it may reduce polarization, because the larger number of dimensions in the opinion space generates more possibilities for overlap and thus for positive influence between neighbors. This intuition resonates with results obtained by Axelrod (1997; see further Klemm, Eguiluz, Toral, and San Miguel, 2003a, 2003b) and Macy et al. (2003) who found that more opinion dimensions reduce diversity and polarization. If more issues foster consensus and reduce polarization, we should thus see that the effect of long-range ties diminishes under both sets of micro-assumptions, when the number of issues increases.

In contrast with earlier models of social influence dynamics (e.g., Abelson, 1964; Harary, 1959; Friedkin and Johnsen, 1990, 1999), studies that include bivalent interaction (Mark, 2003; Macy et al., 2003; Jager and Amblard, 2005; Kitts, 2006; Baldassarri and Bearman, 2007; Flache and Mäs, 2008a, 2008b) did not obtain analytical results but instead employed agent-based computational modeling (see Macy and Flache, 2009). An important reason is that the combination of assimilation and differentiation implies an inherently nonlinear dynamic, while previous analytical models assumed linear social influence functions. This makes analytical solutions difficult to obtain (see Hegselmann and Krause, 2002, for similar arguments), a problem that is exacerbated by the additional complication of heterogeneous network structures that we introduce in this research. Accordingly, we employ computational experiments to explore model implications.

To preview the results, we find that the macro-level effect of long-range ties on polarization or consensus depends decisively on the micro-level assumptions about the positive and negative valence of interaction. When negative valence of interaction is assumed away, long-range ties have an intuitively obvious integrative effect. When this restriction is relaxed, we find that long-range ties deepen cultural divisions. This turns out to be robust against variations in the number and local density of network clusters and the number of salient cultural issues.

2. THE MODEL

2.1. Modeling Influence, Selection, and Access

The model of influence and selection is an extension of Hopfield’s attractor network (Hopfield, 1982; Nowak and Vallacher, 1997).

Each node, or agent, has one or more continuous *states*, each corresponding to the agent's position on a cultural dimension. The state (or opinion) of agent i on issue k , ($1 \leq k \leq K$), s_{ik} , ranges between -1 and $+1$, ($-1 \leq s_{ik} \leq 1$). Following Axelrod's broad definition of "culture," we view states as an abstract representation of "opinions" or "what social influence influences" (Axelrod, 1997). States represent, for example, political views, beliefs, religious and moral values, or artistic tastes. Agents are linked to other nodes through endogenous *weights* in a dynamic social network. Weights are only defined between neighbors, that is, agents who are linked to each other in the exogenously given static access network. The weight dynamic implements social selection. The more similar nodes i and j are in their states at time t , the higher is the value of the weight $w_{ij,t}$. Weights represent the valence (positive or negative) and strength of directed social ties between agents.

Weights change over time because states change through social influence or differentiation. The valence of weights is the key micro-level manipulation of our computational experiments. Positive weights entail positive influence and selection; negative weights impose negative influence and selection. When only nonnegative weights are assumed ($0 \leq w_{ij} \leq 1$), selection is limited to homophily, generating positive ties to similar others and broken ties (neutrality) to maximally dissimilar others, and influence is limited to assimilation. When weights are allowed to be positive or negative ($-1 \leq w_{ij} \leq 1$), selection can be characterized by xenophobia (negative ties to dissimilar others) or homophily (positive ties to similar others), or neutrality ($w_{ij} = 0$). Similarly, influence can be characterized by differentiation ($w_{ij} < 0$), assimilation ($w_{ij} > 0$), or independence ($w_{ij} = 0$).

More precisely, the weight $w_{ij,t}$ increases or decreases depending on similarity between i and j in the K dimensions. After updating of weights, the weight adopts a level that is proportional to the current level of similarity. The new weight is negative if the average distance between i and j across all dimensions of the opinion space exceeds one, that is, half of the maximal average distance. If this average distance is exactly one, the weight is zero and otherwise it obtains a positive value. That is, we assume that an agent follows a simple focal point heuristic to assess whether the opinion distance to some other agent j is so large that the relationship is valued negatively. We assume throughout that states are initially uniformly distributed in $[-1, +1]$, which in turn implies that the expected initial weight of a relationship

is positive.³ Technically,

$$w_{ij,t+1} = 1 - \frac{\sum_{k=1}^K |s_{jk,t} - s_{ik,t}|}{K}. \quad (1)$$

For $-1 \leq s \leq +1$ this function yields results in the range $-1 \leq w \leq +1$. To constrain weights within $[0,1]$ in the model version that includes only nonnegative weights (i.e., homophily and assimilation but not xenophobia and differentiation), we obtain that version by modification of (1) as follows:

$$w_{ij,t+1} = 1 - \frac{\sum_{k=1}^K |s_{jk,t} - s_{ik,t}|}{2K}. \quad (1a)$$

Similar to previous models of social influence with continuous opinions (Abelson, 1964; French, 1956; Hegselmann and Krause, 2002), we assume that the change of agent i 's state k is an aggregated result of the influences imposed by all other agents who influence i . Time is modeled in discrete periods. The updated state, $s_{ik,t+1}$ is obtained by adding to the existing value a weighted sum of the pressures of all neighbors, where these weights correspond to w_{ij} . If $w_{ij} > 0$, j "pulls" i toward j 's states on all dimensions, but j "pushes" i away from j 's states if the weight is negative. As w_{ij} increases in absolute value from 0 to 1, the strength of j 's influence on i increases.

With only positive weights, the "pull" from j is stronger to the extent that i and j hold more similar opinions. Eq. (2) formalizes how the "raw" state change of agent i 's opinion on issue k , $\Delta s_{ik,t}$, results from these assumptions.

$$\Delta s_{ik,t} = \frac{1}{2N_{l,t}} \sum_{j \neq i} w_{ij,t} (s_{jk,t} - s_{ik,t}), \quad (2)$$

³We initialize states to be uniform randomly distributed in $[-1, +1]$. This implies that the distances between the agents are not uniform randomly distributed and hence neither are the weights. More precisely, the expected weight in the initial condition is one third. The special issue editors raised the concern that our results might be sensitive to this "positive bias" in the initial distribution of weights. We therefore modified the weight function to allow the expected weight when opinions are randomly distributed to be set arbitrarily at any level from 0 to 1/3. Results of the corresponding robustness test are reported in the online appendix. In most conditions, we observed no qualitative differences with the central results reported below. Moreover, the conditions in which we observed substantial differences confirmed the explanatory mechanism proposed above, that is, that long-range ties foster polarization if ties between agents within short range are less likely to be negative compared to long-range ties. At the end of the online appendix, we also discuss the behavioral assumptions that correspond to different specifications of the weight function and conclude that our original model is to be preferred.


where $N_{i,t}$ is the number of agent i 's neighbors in the access network, indexed by j . The factor 2 in the denominator assures that $\Delta s_{ik,t}$ is constrained within the interval $[-1, +1]$. The time index t of $N_{i,t}$ captures the possibility that the structure of the access network changes due to exogenous factors. To assure that $-1 \leq s_{ik,t+1} \leq +1$, we smooth the change of opinions when agents move toward the extreme ends of the opinion scale (Eq. (2a)).

$$s_{ik,t+1} = \begin{cases} s_{ik,t} + \Delta s_{ik,t}(1 - s_{ik,t}), & \text{if } s_{ik,t} > 0 \\ s_{ik,t} + \Delta s_{ik,t}(1 + s_{ik,t}), & \text{if } s_{ik,t} \leq 0 \end{cases}. \quad (2a)$$


The exogenously imposed access network that channels social influence and social selection is operationalized as an undirected graph on the population of agents. Our main interest is what happens when, in this access network, a small number of long-range ties are introduced between otherwise disconnected or distant network clusters. Figure 1 shows two of the access networks that form the backbone of our study.

The “disconnected caveman graph” (Fig. 1a) represents the situation of maximally dense local network components (see Watts, 1999, for the original definition of the “caveman” structure). “Caves” are the technical implementation of network clusters. Each of the 20 caves that are shown in the figure is a complete subgraph containing five members, but there are no ties between caves. We describe the local structure using Watts's (1999) clustering coefficient C , defined as the proportion of pairs of i 's neighbors who are neighbors of each other, averaged across all nodes in the network. In the disconnected caveman graph, local clustering is at its theoretical maximum, indicated by $C = 1$. The global structure of the graph is captured by the mean geodesic L , the average length of the shortest path between two nodes i and j , averaged across all pairs of nodes. In the disconnected caveman graph of Figure 1a, $L = \infty$.

Small worlds combine high local clustering and small average distances between nodes, made possible by a small proportion of long-range ties that connect otherwise distant (or even completely disconnected) communities. Figure 1b shows a small-world access network created by adding to the disconnected caveman graph 50 bridge ties between randomly selected pairs of previously nonadjacent nodes. In the example shown here no disconnected components remain. Density has increased only marginally (from .04–.045) and clustering remains high ($C = 0.81$), but now there are on average only $L = 6.29$ steps in the network between any two randomly chosen nodes (equivalent to the famous “six degrees of separation”). Adding random ties in



(a)



(b)

FIGURE 1 Typical network structures used in experiment 1. $N = 100$, five agents per cave: (a) disconnected caveman graph (density $D = 0.04$, clustering $C = 1.0$, average path length $L = \infty$); (b) disconnected caveman graph with 50 random ties added (density $D = 0.045$, clustering $C = 0.814$, average path length $L = 6.29$).

this way does not necessarily eliminate all disconnected components, but it does increase the number of nodes that are connected and reduces the average path length within components. We are interested in the effects on consensus and polarization when a small number of long-range ties are added to a disconnected caveman graph. In the first computational experiment of our article, we explore this manipulation. In the second experiment, we disentangle effects of connecting components with long-range ties from the effects of reducing average path length with long-range ties. For this, we start in the second experiment from a connected graph, and examine effects of long-range ties that reduce the mean geodesic L .

We use asynchronous updating of states and weights to avoid possible artifacts from synchronous updating (see Huberman and Glance, 1993). Both the K states and the weights of the relationships are endogenous and change in discrete time steps. In every time step,

one agent is selected randomly with equal probability. To further avoid synchronicity, either a randomly selected state or the weights of the focal agent are selected for updating, but not both at the same time. Agents are updated with replacement. That is, the same agent can be selected in two consecutive time steps. The duration of a single realization of the model is expressed in the number of iterations. An iteration corresponds to N time steps, where N refers to the number of individuals in the population. Throughout this article, we assume $N = 100$.

2.2. Outcome Measure for Polarization

Our main interest is the effect of long-range ties on social polarization. Polarization captures the degree to which a population can be separated into two equal-sized factions who are maximally dissimilar in the opinion space and have maximal internal agreement. To compute the population-level measure of *polarization* (P) we first calculate for every pair of agents (ij) in the population the population distance d_{ij} (averaged across all K cultural dimensions), $d_{ij,t} = \frac{1}{K} \sum_{k=1}^K |s_{jk,t} - s_{ik,t}|$. This distance ranges between zero (full agreement) and two (full disagreement). Distances to self are zero by definition and are excluded from the computation of the polarization measure. The level of polarization of the population at time t , P_t is then obtained as the variance of the distribution of $d_{ij,t}$:

$$P_t = \frac{1}{N(N-1)} \sum_{\substack{i=N, j=N \\ i \neq j}} (d_{ij,t} - \bar{d}_t)^2, \quad (3)$$

where the symbol \bar{d}_t denotes the average opinion distance across all pairs (ij) where $i \neq j$. The minimum level of polarization ($P = 0$) obtains when all pairwise distances (excluding self-distances $d_{ii} = 0$) are equal, corresponding to full consensus on all opinion dimensions.

If self-distance were included in computation of \bar{d}_t and P , the maximum level of polarization would be $P = 1$. Maximal polarization is obtained when the population is equally divided on all issues between the opposite ends of the opinion scale at -1 and $+1$, such that two members of different factions never agree on any issue. With self-distances excluded, the maximal value of P slightly differs from one⁴ but approaches one as N increases. For $N = 100$ it is, for example, 0.999898. With self-distance included, i and j belong in the case of a maximally polarized population to the same faction in 50% of all dyads

⁴To be precise, with self-distances excluded the value of P for maximal polarization is $P = 1 - 1/(N-1)^2$.

and thus fully agree, with $d_{ij}=0$. In the other dyads, i and j belong to different factions and fully disagree on all issues, with $d_{ij}=2$. Accordingly, the squared distance between d_{ij} and the average opinion distance of $\bar{d} = 1$ is one in all dyads, yielding a variance of one, that is, polarization $P = 1$. With uniformly distributed opinions (in which every state obtains with equal frequency), the polarization measure yields approximately $P = 0.214$ for $K = 1$ and approximately $P = 0.11$ for $K = 2$.

3. RESULTS

We report results for two computational experiments. The core question addressed by both experiments is how the effect of long-range ties—on cultural integration versus polarization—depends on the micro-level assumption about the valence of influence and selection. In each experiment, we manipulate whether the path weights in the dynamic interaction network are constrained (or not) to nonnegative values. The experiments differ in the structure of the exogenous access networks.

In experiment 1, we take the disconnected caveman graph (Fig. 1a) as an extreme case of a “large world” and test the effect of adding random ties that connect the previously disconnected caves. In experiment 2, we compare the effect on polarization of adding “short-range” versus long-range ties, which differ in the extent to which they reduce the mean geodesic in a connected graph. In both experiments, we submit the effects generated by the central manipulation to three robustness tests: variation in the number and size of caves, variation in the network density of caves, and variation in the number of salient cultural issues, K .

3.1. Experiment 1

The baseline condition of experiment 1 is the “large-world” network structure of the disconnected caveman graph shown in Figure 1a. We manipulated this structure by connecting each pair of previously unconnected agents with a small probability (0.003), resulting in a network similar to the one shown in Figure 1b. Averaged across 50 realizations, 28.75 undirected new network ties were added by this procedure, an average increase of the density of about .006.⁵ The manipulated

⁵Previous work (e.g., Watts and Strogatz, 1998) “rewired” ties in order to keep density constant. However, rewiring changes the local structure of the access graph in some caves such that caves have no longer the same local structure. To avoid this complication, we left the local structure of all caves intact and add long-range ties, which resulted in a marginal increase in density. In experiment 2, we show that effects of this manipulation are not caused by changes of the network density of the access graph.

access network is not necessarily connected; we found that on average about four out of 20 caves remained unconnected. We assume $K=2$ dimensions of the state space. Previous work (Flache and Mäs, 2008a, 2008b; Macy et al., 2003) has shown that for this number of dimensions polarization is possible but not trivial if there is both positive and negative influence. Opinions are initially drawn randomly from a uniform distribution independently across dimensions.

Figure 2 illustrates the effect of adding random ties on the polarization dynamics for a prototypical realization of the model with only non-negative selection and influence. The figure shows the change in the distribution of agents' opinion vectors in the two-dimensional opinion space for snapshots from the first 100 iterations. The initial distribution is held constant across both conditions in Figure 2.

The dynamics shown in Figure 2 reflect the strong tendency toward consensus known for models based on positive influence and selection. In both conditions, agents' opinions move on average closer toward each other over time. However, perfect consensus can not be reached in the disconnected caveman graph; all agents within the same cave reach a local consensus different from those of other caves by iteration

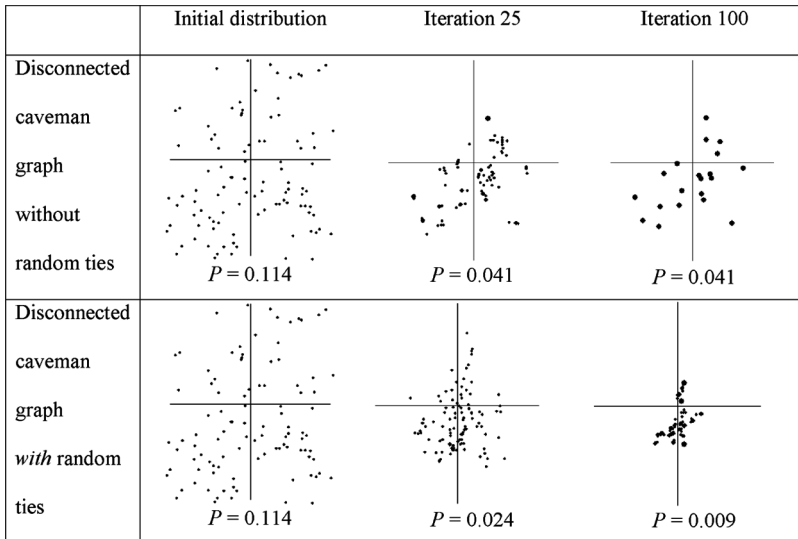


FIGURE 2 Change of opinion distribution for prototypical realizations. Numbers below graphs indicate polarization measure. Model *without* negative valence of interaction. $N=100$, $K=2$, initially 20 isolated caves with five agents per cave. The size of a dot indicates the number of agents with the opinion vector corresponding to the midpoint.

100 (indicated by 20 different opinion vectors). In all caves, the local consensus reflects roughly the average of initial opinions, yielding outcomes close to the midpoint of the opinion distribution due to the uniform distribution of initial opinions. The remaining differences between caves imply very low polarization ($P = 0.041$). Figure 2 shows how in the graph with random ties, social influence occurs across the boundaries between caves (lower row of snapshots). Starting from the same initial distribution of opinions, a higher level of consensus has developed by iteration 100, indicated by an even lower level of polarization of $P = 0.009$. Replication of the experiment with a larger number of iterations confirmed that in the long run perfect consensus develops within each of the caves in the network.

Figure 3 charts the average dynamics of polarization based on 50 independent realizations. We added for this figure a third condition with random ties added after 50 iterations, when initial opinion differences within caves had largely disappeared. Random ties may accelerate the formation of opinion consensus more when the opinion differences are larger at the moment that the random ties are added to the network.

Figure 3 shows that polarization is low and declines over time in all three access network conditions. The results confirm the pattern illustrated by Figure 2. In the disconnected caveman graph, social influence has largely eliminated differences within caves by iteration 20 and polarization stabilizes at about $P = 0.03$. With random ties added, polarization declines further, down to about $P = 0.01$ after 100 iterations. When random ties are added in iteration 50, opinion differences between the caves are quickly reduced so that about the same level of consensus is reached in iteration 100 as in the condition where random ties were added at the outset.

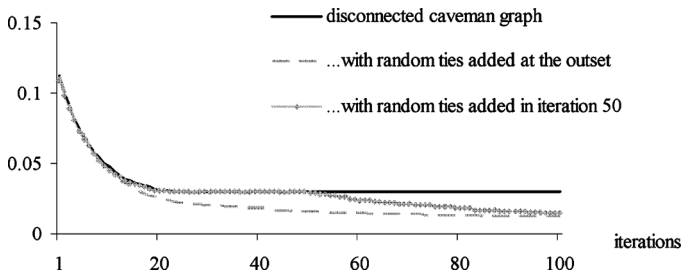


FIGURE 3 Effect of adding random ties to disconnected caveman graph on dynamics of polarization measure. Averages based on 50 independent replications per condition. Model *without* negative valence of interaction. $N = 100$, $K = 2$, initially 20 isolated caves with five agents per cave.

This result reflects Abelson’s (1964) and Harary’s (1959) findings that consensus is inevitable in connected graphs. In a model with no negative weights, consensus forms within each of the network components, and opinion differences remain only between components. In the baseline graph these components are the caves; in the graph with random ties the number of components varies randomly but is typically very small.

In sum, the micro-theory assuming positive influence and selection implies that long-range ties increase consensus and reduce polarization. Next, we repeat the experiment with the alternative micro-theory that introduces the possibility for xenophobia and differentiation. Figure 4 shows the dynamics that result from the same initial distribution as that used for Figure 2.




FIGURE 4 Change of opinion distribution for prototypical realizations. Numbers below graphs indicate polarization measure. Model *with* negative valence of interaction. $N = 100$, $K = 2$, initially 20 isolated caves with five agents per cave. The size of a dot indicates the number of agents with the opinion vector corresponding to the midpoint.

Opinion dynamics shown in Figure 4 are characterized by both assimilation and differentiation. Without random ties, most agents have by iteration 25 moved toward the midpoint of the opinion distribution in their cave, but in some caves agents also increase their distance to others and shift toward extreme opinions. The self-reinforcing dynamics of polarization are constrained by the boundaries of the caves. By iteration 100, a local consensus has established in 14 of the 20 caves. In those caves, all agents adopt “moderate” opinions, close to the midpoint of the opinion scale. However, in six caves, all agents have adopted extreme positions on both dimensions. The prevalence of local consensus is due to the fact that initially weights are more likely to be positive than negative, such that there is more assimilation than differentiation within ties. In those caves that did polarize, initial opinion differences created relatively more negative ties by random chance. To differentiate from their “enemies,” agents shift their opinions toward the extreme opposite end of the spectrum, pulling their friends with them, and turning positive ties into hostile relationships in the process. These dynamics split each of the polarized caves into two factions with identical states within and opposite states between. However, not all caves split along the same dividing line. In some, one faction moves towards the pole $(+1, +1)$ and its counterpart to $(-1, -1)$, in others the split is $(-1, +1)$ vs. $(+1, -1)$. The overall polarization level of $P = 0.27$ reflects the resulting distribution of about one third of the population on one of the corners of the opinion space and about two thirds on a local consensus close to its midpoint.

Random ties leave the local structure intact but remove the constraints for the diffusion of extremism. The lower row of Figure 4 shows the dynamics that evolve from the same initial opinion distribution with random ties added. By iteration 200, initially moderate agents are on the move toward one of those two corners of the opinion space, in which most of the members of the polarized caves ended up when caves were disconnected.

The local dynamics within caves are in the short term similar to those obtained when random ties were not added. Random ties create bridges between agents in caves that polarized in the short term and caves that are initially nonpolarized, with all agents adopting a moderate opinion. These bridges exert positive influence when the moderate shades toward the position of the extremist and otherwise the influence is negative. The effect of negative influence between extremists and moderates is asymmetric: the extremist is reinforced in their “one-sided” view, but the moderate is not reinforced in their ambivalence. Rather, the moderate is pushed toward a less ambivalent position but in the opposite direction from the extremist. The former

moderate then exerts both positive and negative influence on its moderate neighbors, allowing for extremism to spread through the previously moderate cluster. In short, the asymmetric influence when bridges connect extremists and moderates destabilizes the moderate consensus in previously nonpolarized caves, causing agents in those caves to adopt more extreme positions. As a result, the overall population ends up considerably more polarized than without random ties added, with an equilibrium level of polarization of $P=0.721$ attained in iteration 1000. At this stage, only one (disconnected) cave has not polarized. That is, about 95% of the population members adopt extreme opinions on both dimensions such that they are either in complete agreement or complete disagreement with their neighbors. In this case, most agents adopted either the opinion vector $(-1, +1)$ or its counterpart $(+1, -1)$.

For reliability, we conducted 50 independent realizations of the experiment. Figure 5 reports the average dynamics of polarization. As in Figure 3, we inspected also a third condition in which the random ties were added after 2000 iterations, when it was certain that there were no more changes of opinion within the disconnected caves.

Figure 5 confirms that random ties considerably increase polarization under the assumptions of both assimilation and differentiation. Without random ties, polarization reaches a peak level of about $P=0.23$ around iteration 100 and then stabilizes. When long-range ties are added, polarization further increases at that point and eventually reaches about $P=0.5$. This also happens when random ties are added later, in iteration 2000.

Experiment 1 shows that the effects of random ties on social polarization depend on the micro-level mechanism assumed. With only

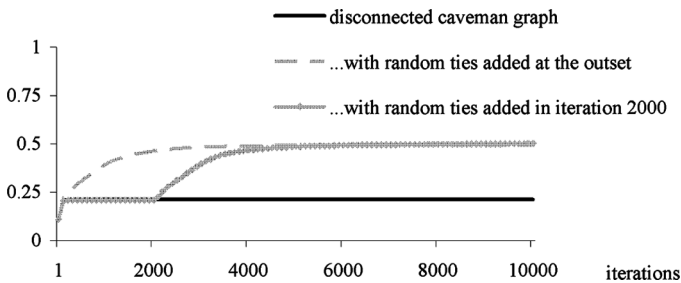


FIGURE 5 Effect of adding random ties to disconnected caveman graph on dynamics of polarization measure. Averages based on 50 independent replications per condition. Model *with* negative valence of interaction. $N=100$, $K=2$, initially 20 isolated caves with five agents per cave.

assimilation (positive influence) and homophily (positive selection), random ties decrease polarization. When differentiation (negative influence) and xenophobia (negative selection) are also allowed, random ties entail higher levels of polarization. This result turns out to be robust when we relax our structural assumptions about the size and local density of caves. Figure 6 reports results for a *ceteris paribus* replication of experiment 1 with cave size varying between $N_c = 3$ and $N_c = 100$. N was chosen as number closest to 100 such that all caves have equal size ($N = 99, 90$, and 120 for $N_c = 3, 30$, and 40 , respectively; otherwise, $N = 100$). At all cave sizes inspected, we observed for the model with only positive influence and selection that polarization was lower when random ties were added, but the difference was negligible for $N_c \geq 20$. Conversely, when both positive and negative

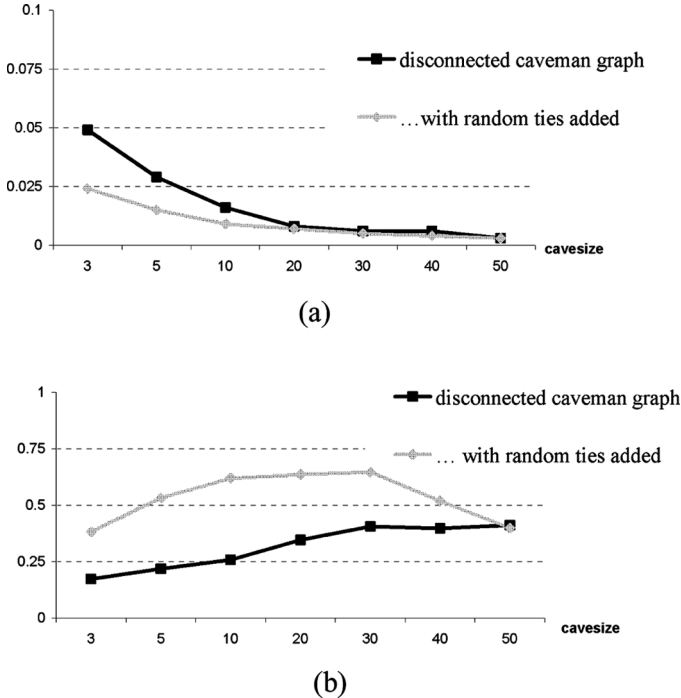



FIGURE 6 Replication of experiment 1 with varying size of caves. Polarization in endstate for cave sizes between three and 50. Averages based on 50 independent replications per condition. $K = 2$, initially 20 isolated caves with five agents per cave: (a) without negative valence of interaction, random ties added in iteration 50; (b) with negative valence of interaction, random ties added in iteration 2000.

influence and selection were included, random ties increased polarization, except for $N_c = 50$. In this condition, we also observed a nonlinear relationship between the size of caves and polarization, both with and without random ties added.


The nonlinearity reflects the interplay of two different effects of cave size. The larger the caves, the more likely a cave will initially escape consensus and polarize, due to the higher likelihood that within some dyads in the cave opinion differences are initially sufficiently large to trigger xenophobia and differentiation. We observed that 71 of 100 disconnected caves of size 5 ended in consensus before random ties were added, while consensus emerged in only 18 out of 100 caves when cave size was 50. At the same time, the larger the cave, the more likely a large majority of its local population will adopt the same extreme opinion vector, driven by differentiation from a small minority that adopts the opposite pole of the opinion space. An uneven distribution on the poles, in turn, results in low values of the polarization measure P . The emergence of an uneven distribution in large caves is caused by the more densely populated opinion space. As a consequence, most agents are initially connected to a larger number of others with similar initial opinions. We observed that in those caves in which consensus failed to develop, negative weights obtained at the outset between about 31% of all dyads when the cave size was five, while this figure dropped to about 18% when cave size was 50.

We also conducted a *ceteris paribus* replication with varying degrees of local density of caves. More precisely, we replicated experiment 1 with three conditions, removing 25%, 50%, or 75% of all local ties at the outset. In all three conditions, we found that polarization in the end state was reduced by random ties when only positive mechanisms were allowed, while random ties increased polarization with both positive and negative mechanisms.

As a final robustness test, we conducted a *ceteris paribus* replication of experiment 1 with the number of issues, K , varied across the values of $K = 1, 2, 3, 5$, and 10. Earlier work (Axelrod, 1997; Macy et al., 2003) suggested that a larger number of cultural issues fosters consensus and decreases polarization. Figure 7a shows the combined effect of random ties and cultural complexity (K) on polarization under the model with only positive influence and selection. Results confirm that more issues increase consensus and reduce polarization with and without random ties. With initial opinions uniformly distributed and drawn independently from each other, the average initial distance is less likely to exceed 1 as K increases, and, therefore, the initial weight is more likely to be positive. Without random ties added, the consequence is that the distinct opinion vectors on which the disconnected



(a)



(b)

FIGURE 7 Replication of experiment 1 with different numbers of salient cultural issues, K . Effect of K on polarization in endstate, with and without random ties added. $N = 100$, $K = 2$, initially 20 disconnected caves with five agents per cave: (a) without negative valence of interaction, random ties added in iteration 50; (b) with negative valence of interaction, random ties added in iteration 2000.

caves locally converge become more similar. With random ties, an additional effect is that the initial similarity amplifies positive influence. Correspondingly, Figure 7a exhibits a negative effect of random ties on polarization across all numbers of issues that we inspected.

Figure 7b shows that the model with both positive and negative valence of interaction generates the same qualitative result for the effect of K . More cultural dimensions make it less likely that polarization develops in local caves in the first place. Polarization can only develop in those caves where at the outset overall opinion distances in a sufficient number of dyads are sufficiently large to trigger negative influences that can split the cave into opposing factions. The likelihood that from the outset two agents largely disagree across all

salient cultural issues declines in the number of these issues, given that opinions are initially distributed randomly and independently from each other. At the same time, we find that random ties increase polarization across almost all levels of K , which is consistent with the results of experiments 1. The exception is $K = 10$, for which polarization virtually disappears, such that there is no discernible difference between the two structural conditions.

In sum, experiment 1 confirms that the effect of greater connectedness on social integration and polarization may decisively depend on the micromechanisms underlying cultural interaction. With only positive influence and selection, long-range ties promote greater cultural integration and assimilation. When both positive and negative valences are assumed, the effect is reversed. Long-range ties become conduits for the spread of locally developed polarization to regions of the network that would have otherwise remained nonpolarized.

3.2. Experiment 2

The previous experiment showed how random ties that bridge network components, or caves, can decisively affect polarization dynamics. This effect could be due to one of the following two possibilities: (1) random ties create connections between previously disconnected caves, or (2) random ties increase network density. Experiment 2 was designed to disentangle these possibilities, using a connected caveman graph and adding short-range ties that increase density without reducing the average geodesic versus long-range ties that increase density while reducing the mean geodesic. Figure 8 shows these graph structures.

In the connected caveman graph (Fig. 8a), there is only one component, because each cave is connected to the “next” cave in a “ring” of connected caves. A cave is linked to each of the two caves that are adjacent to it in the spatial arrangement, with one bridge tie per adjacent cave. We added at random 20 additional undirected “short-range ties” to obtain structures like the one shown in Figure 8b. The short-range ties that were added linked a random node in each cave to a random node of the cave on the right side to it in the circular spatial arrangement shown in Figure 8a. The example shows that the additional short-range ties reduce both the average path length L and the level of clustering only moderately, from $L = 10.7$ to $L = 8.58$ and from $C = 0.84$ to $C = 0.67$, respectively. By contrast, when the same number of ties between random pairs of nodes drawn from the entire graph were added (shown in Fig. 8c), the average path length decreased about three times as much, from $L = 10.7$ to




FIGURE 8 Graph structures for comparison of effects of long-range and short-range ties in experiment 2. $N = 100$, five agents per cave. (a) Connected caveman graph. $D = 0.044$, $C = 0.84$, $L = 10.7$. (b) Connected caveman graph with 20 random short-range ties added. Black ties = short-range ties added. $D = 0.048$, $C = 0.67$, $L = 8.58$. (c) Connected caveman graph with 20 random ties added. $D = 0.048$, $C = 0.728$, $L = 4.73$.

$L = 4.73$, while clustering remained at a high level ($C = 0.73$). In other words, these random ties are “long-range ties,” in the sense that on average they have a larger range than the short-range ties shown in Figure 8b.

We wanted to know how short-range versus long-range ties affect opinion dynamics when they appear at a moment that opinion distributions have settled into stable patterns within local caves. Accordingly, we entered the additional ties once the model reached equilibrium (at iteration 50 for the model with only positive valence of interaction and at iteration 2000 for the model with both positive and negative valence of interaction.)




FIGURE 9 Dynamics of polarization for connected caveman graph without random ties, with 20 short-range ties added and with 20 random ties added, both in iteration 50. Model *without* negative valence of interaction. $N = 100$, $K = 2$, initially 20 connected caves with five agents per cave.

Figure 9 shows the dynamics of polarization that we obtained with the model with positive valence of interaction only, averaged across 50 independent realizations per condition. Figure 9 illustrates that there is no discernible difference between the conditions. As Abelson (1964) and Harary (1959) have shown, with only positive influence consensus is inevitable when the network is connected. Comparison with the results of experiment 1 shows that the predominant effect of adding ties occurs when the ties bridge disconnected caves. The range of ties makes no further difference in a connected graph.

Figure 10 charts the results when interaction can have both positive and negative valence of interaction. The results reported in




FIGURE 10 Dynamics of polarization for connected caveman graph without random ties, with 20 short-range ties added and with 20 random ties added, both in iteration 2000. Model *with* negative valence of interaction. $N = 100$, $K = 2$, initially 20 connected caves with five agents per cave.

Figure 10 are strikingly different from those obtained with only positive valence of interaction. The range of additional ties affects opinion dynamics. While additional short-range ties do not discernibly change the dynamics of polarization, polarization increases sharply from about $P=0.4$ to about $P=0.55$ when additional long-range ties are added. In a connected but highly clustered network, both assimilation and differentiation occur locally among agents in the same local region of the network. Local region refers here to a set of caves that are neighboring each other in the spatial arrangement. The resulting spatial distribution of opinions is characterized by the mix of local regions with large consensus on moderate opinions and some local regions in which local polarization has developed such that neighbors differ sharply from each other. Short-range ties tend to fall by definition within the same local region such that they either link two agents who both have moderate opinions or two agents who have extreme opinions either on the same or the opposite corner of the opinion space than the other. In both cases, influence via the added tie can not increase opinion differences. By contrast, long-range ties are more likely to link agents in polarized local regions to agents with moderate opinions in other local regions that are initially nonpolarized. When this happens, the influences from extremists destabilize the local moderate consensus, while the extremists are “anchored” in their extreme views by the need to differentiate from their local extremist neighbors.

Experiment 2 shows how the effect of the range of ties on polarization interacts with the micro-level mechanisms of influence and selection. With only positive mechanisms, we obtained no discernible effect of the range of ties. With both positive and negative valence of interaction, however, long-range ties increased polarization much more than did short-range ties. Robustness tests showed that this result generalizes to a wider range of conditions. Figure 11 shows effects of the same cave size manipulation that we used for experiment 1, this time as *ceteris paribus* replication of experiment 2. Across all cave sizes inspected, we found that the range of random ties had virtually no effect on polarization with only positive mechanisms, but long-range ties increased polarization when both positive and negative valences of interaction were included. This pattern also obtained when we varied the local density of caves across three conditions (25%, 50%, and 75% of local ties removed).

Finally, as Figure 12 illustrates, the qualitative result was also robust across different numbers of salient cultural issues and variation of the number of issues has the same effects that we also observed in experiment 1.

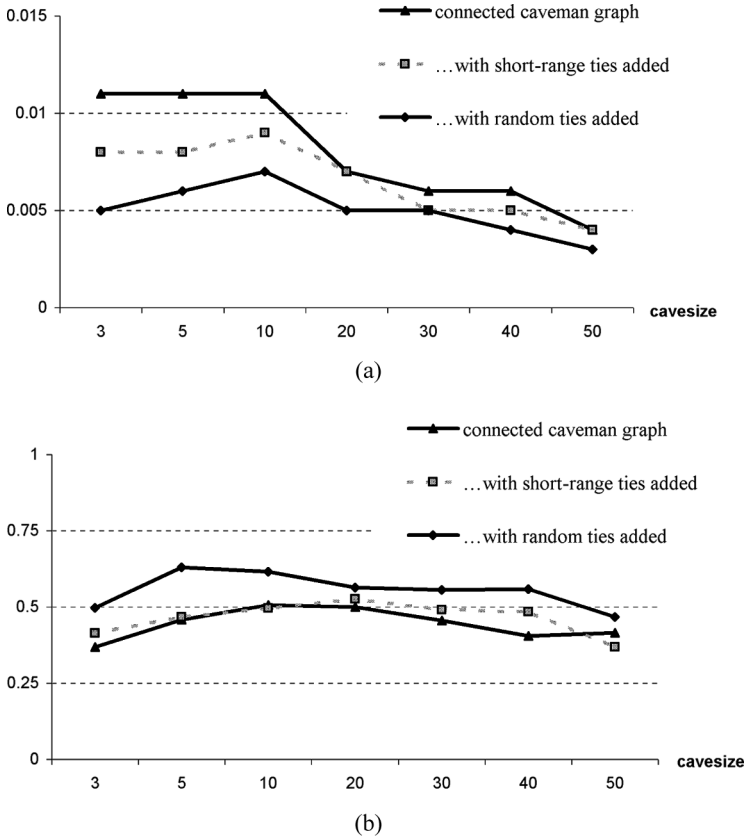


FIGURE 11 Replication of experiment 2 with varying size of caves. Polarization in endstate for cave sizes between three and 50. Averages based on 50 independent replications per condition. $K=2$, initially 20 connected caves with five agents per cave; (a) without negative valence of interaction, ties added in iteration 50; (b) with negative valence of interaction, ties added in iteration 2000.

In sum, experiment 2 shows that model implications for the effect of the range of ties in connected networks depend decisively on the underlying behavioral assumptions. When only positive influence and selection are assumed, consensus is the inevitable outcome in connected networks. Neither adding short-range nor long-range ties has any effect on the level of consensus that develops because the graph is connected without additional ties. When both positive and negative influence and selection are allowed, long-range ties increase

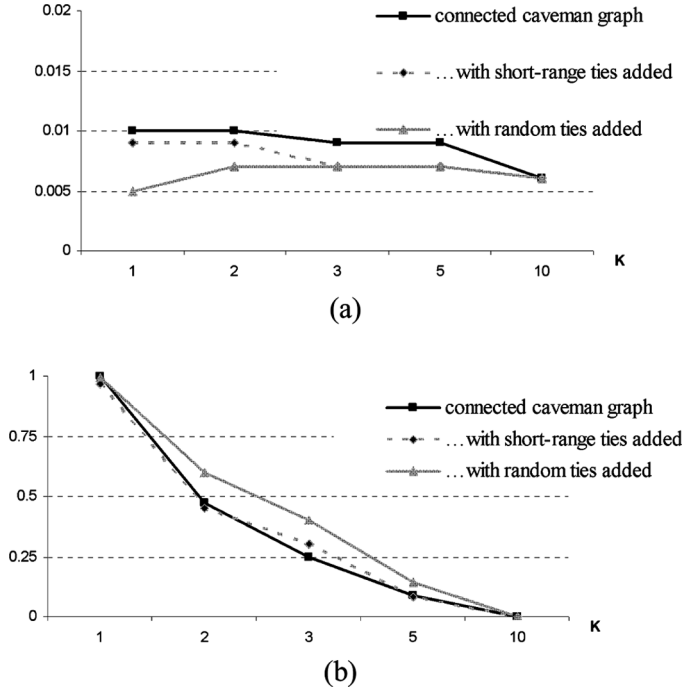


FIGURE 12 Replication of experiment 2 with different numbers of salient cultural issues, K . Effect of K on polarization in endstate. Connected caveman graph without ties added, with 20 short-range ties. and with 20 long-range ties added. $N = 100$, initially 20 connected caves with five agents per cave: (a) model *without* negative valence of interaction, ties added in iteration 50; (b) model with negative valence of interaction, ties added in iteration 2000.

polarization sharply, but short-range ties do not. Short-range ties tend to link pairs of agents whose opinion distances vis-à-vis each other have already been shaped by the same local influences that they exert upon each other via the short-range tie. Long-range ties can link agents from neighborhoods with different local equilibria, pushing in particular agents in nonpolarized clusters to assimilate to or differentiate from the extreme opinions taken by agents in polarized clusters. This results in a higher level of polarization in the overall population.

4. DISCUSSION AND CONCLUSION

Building on Granovetter’s work, research on “small-world” networks suggests that a small proportion of long-range ties that

bridge otherwise distant or disconnected clusters can promote cultural integration. We show that this macrosocial effect depends decisively on the micro-level assumption that interaction is limited to positive influence and selection. Using a computational model similar to those in earlier studies of cultural convergence and polarization, we found that our model generates an integrating effect of long-range ties when only positive influence and selection are allowed. However, long-range ties turned out to have the opposite effect when homophily and assimilation are combined with their negative moments, xenophobia and differentiation. In that case, long-range ties fostered cultural polarization rather than integration.

An important scope condition for our argument is that actors who are connected with a long-range tie are more likely to differ sharply from each other than are actors connected with a short-range tie. Under that condition, interactions in long-range ties are more likely to be driven by the negative mechanisms of social interaction, xenophobia and differentiation, than are interactions within short-range ties. In our model, this condition arises because from a random start, there is on average more assimilation between agents than there is differentiation. This implies that agents within close proximity in the network become more similar in their opinions than agents within large distance from each other. In the online appendix of this article, we have explored alternative assumptions about the initial distribution of weights that render positive ties initially equally likely than negative ties. The results show that then additional long-range ties in the network no longer have a stronger effect on polarization than additional short-range ties.

Our models are obviously highly stylized. But their main elements, small-world network structures and social influence and selection processes, capture important features of intercultural contact in the social world. People mostly interact with others in geographical proximity, and within similar cultural and social segments of the population. Yet increasing cultural diversity of Western societies, spatial mobility, and more powerful communication technologies also create long-range interpersonal contacts that bridge network clusters. Social influence and social selection are well-documented fundamental social mechanisms in interpersonal interaction. Accordingly, we have focused on alternative specifications of these mechanisms in our theoretical exploration of the implications of long-range ties on cultural diversity.

Admittedly, our stylized models do not address other important elements of intercultural contact. For example, much of exposure to

other cultures occurs through media, not via interpersonal contacts. Moreover, national governments, political parties, and organizations can also influence people's perception of intercultural differences in order to foster intercultural integration or to exploit intercultural conflict for their own purposes. Future research should carefully explore how such complications affect the mechanisms that we have addressed.

Beyond that, this study relies on a number of further simplifications that need to be carefully examined. Analytical studies of Klemm and his co-authors (2003a, 2003b) have shown that stable diversity may critically hinge upon the absence of random cultural perturbation. We expect that noise is less critical for our results because the polarized outcomes we are interested in are highly stable against a small level of random changes in the opinion states. In a polarized outcome, any agents who randomly experiment with different cultural states are either immediately drawn back to the same profile of their neighbors or they change sides and move to the other pole if the random exploration was large enough.

Another possible limitation is that we used only highly stylized network structures that do not resemble empirical networks. However, we found that our results were robust across variation of cluster size and local density. Tentatively, we interpret this as evidence that our results may generalize to a broader range of clustered "small-world" networks in which the addition of long-range ties increases connectivity without reducing clustering.

As a final simplification, we have neglected that agents may differ not only on "flexible" opinions but also on "fixed" demographic characteristics (see also Macy et al., 2003; Flache and Mäs, 2008a, 2008b). Empirically, demographic differences tend to correlate with geographical location. Given this correlation, we expect that including fixed agent attributes increases rather than decreases the potential for xenophobia and differentiation. Hence, this complication may not limit but actually strengthen the effect of long-range ties on polarization.

Our results challenge the prevailing view that greater opportunities for cultural contact in an increasingly small world promote cultural convergence and integration. This implication of "small-world" theory depends on the micro-level behavioral assumption that interaction is exclusively positive in valence. This result should caution modelers of cultural dynamics against overestimating the integrative effects of greater cultural contact.

REFERENCES

- Abelson, R. P. (1964). Mathematical models of the distribution of attitudes under controversy. In N. Frederiksen & H. Gulliken (Eds.), *Contributions to Mathematical Psychology* (pp. 142–160). New York, NY: Holt, Rinehart & Winston.
- Adams, J. & Roscigno, V. J. (2007). White supremacists, oppositional culture and the World Wide Web. *Social Forces*, *84*, 759–777.
- Axelrod, R. (1997). The dissemination of culture: a model with local convergence and global polarization. *Journal of Conflict Resolution*, *41*, 203–226.
- Baldassarri, D. & Bearman, P. (2007). Dynamics of political polarization. *American Sociological Review*, *72*, 784–811.
- Buskens, V., Corten, R., & Weesie, J. (2008). Consent or conflict: coevolution of coordination and networks. *Journal of Peace Research*, *45*, 205–222.
- Centola, D. & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, *113*, 702–734.
- Early, P. C. & Mosakowski, E. (2000). Creating hybrid team cultures: an empirical test of transnational team functioning. *Academy of Management Journal*, *43*, 26–49.
- Festinger, L., Schachter, S., & Back, K. (1950). *Social Pressures in Informal Groups*. New York, NY: Harper.
- Flache, A. & Mäs, M. (2008a). How to get the timing right: a computational model of the effects of the timing of contacts on team cohesion in demographically diverse teams. *Computational and Mathematical Organization Theory*, *14*(1), 23–51.
- Flache, A. & Mäs, M. (2008b). Why do faultlines matter? A computational model of how strong demographic faultlines undermine team cohesion. *Simulation Modelling Practice and Theory*, *16*, 175–191.
- French, J. R. P. (1956). A formal theory of social power. *Psychological Review*, *63*, 181–194.
- Friedkin, N. E. (2001). Norm formation in social influence networks. *Social Networks*, *23*(3), 167–189.
- Friedkin, N. E. & Johnsen, E. C. (1990). Social influence and opinions. *The Journal of Mathematical Sociology*, *15*, 193–206.
- Friedkin, N. E. & Johnsen, E. C. (1999). Social influence networks and opinion change. *Advances in Group Processes*, *16*, 1–29.
- Friedman, T. (2005). *The World Is Flat: A Brief History of the Twenty-First Century*. New York, NY: Farrar, Straus and Giroux.
- Glaeser, E. L. & Ward, B. A. (2006). Myths and realities of American political geography. *Journal of Economic Perspectives*, *20*(2), 119–144.
- Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology*, *78*, 1360–1380.
- Greig, J. M. (2002). The end of geography? Globalization, communications, and culture in the international system. *Journal of Conflict Resolution*, *46*(2), 225–243.
- Harary, F. (1959). A criterion for unanimity in French's theory of social power. In D. Cartwright (Ed.), *Studies in Social Power* (pp. 168–182). Ann Arbor, MI: University of Michigan, Institute of Social Research.
- Hegselmann, R. & Krause, U. (2002). Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, *5*(3). Retrieved from <http://jasss.soc.surrey.ac.uk/5/3/2.html>

- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79, 2554–2558.
- Huberman, B. A. & Glance, N. S. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Science*, 90, 7716–7718.
- Jager, W. & Amblard, F. (2005). Uniformity, bipolarization and pluriformity captured as generic stylized behavior with an agent-based simulation model of attitude change. *Computational and Mathematical Organization Theory*, 10, 295–303.
- Kirkman, B. L., Rosen, B., Gibson, C. B., Tesluk, P. E., & McPherson, S. O. (2002). Five challenges to virtual team success: lessons from Sabre, Inc. *Academy of Management Executive*, 16, 67–79.
- Kitts, J. (2006). Social influence and the emergence of norms amid ties of amity and enmity. *Simulation Modelling Practice and Theory*, 14, 407–422.
- Kitts, J., Macy, M. W., & Flache, A. (1999). Structural learning: conformity and social control in task-oriented groups. *Computational and Mathematical Organization Theory*, 5, 129–145.
- Klemm, K., Eguiluz, V. M., Toral, R., & San Miguel, M. (2003a). Global culture: a noise induced transition in finite systems. *Physical Review E*, 67, 045101(R).
- Klemm, K., Eguiluz, V. M., Toral, R., & San Miguel, M. (2003b). Non-equilibrium transitions in complex networks: a model of social interaction. *Physical Review E*, 67, 026120(R).
- Latané, B. & Nowak, A. (1997). Self-organizing social systems: necessary and sufficient conditions for the emergence of clustering, consolidation and continuing diversity. In G. Barnet & F. Boster (Eds.), *Progress in Communication Science: Persuasion* (pp. 43–74). Norwood, NJ: Ablex.
- Lazarsfeld, P. F. & Merton, R. K. (1954). Friendship as a social process: a substantive and methodological analysis. In M. Berger (Ed.), *Freedom and Control in Modern Society* (pp. 18–66). New York, NY: Van Nostrand.
- Macy, M. W. & Flache, A. (2009). Agent-based modeling: social order from the bottom up. In P. Hedström & P. Bearman (Eds.), *The Oxford Handbook of Analytical Sociology* (pp. 245–268). New York, NY: Oxford University Press.
- Macy, M. W., Kitts, J., Flache, A., & Benard, S. (2003). Polarization in dynamic networks: a Hopfield model of emergent structure. In R. Breiger, K. Carley, & P. Pattison (Eds.), *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers* (pp. 162–173). Washington, DC: The National Academies Press.
- Mark, N. (1998). Beyond individual differences: social differentiation from first principles. *American Sociological Review*, 63, 309–330.
- Mark, N. (2003). Culture and competition: homophily and distancing explanations for cultural niches. *American Sociological Review*, 68, 319–345.
- Massey, D. S. & Denton, N. A. (1998). Hypersegregation in U.S. metropolitan areas: Black and Hispanic segregation along five dimensions. *Demography*, 26, 373–391.
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Homophily in social networks. *Annual Review of Sociology*, 27, 415–444.
- Nowak, A. & Vallacher, R. (1997). Toward computational social psychology: cellular automata and neural network models of interpersonal dynamics. In S. Read & L. Miller (Eds.), *Connectionist Models of Social Reasoning and Social Behavior* (pp. 277–311). Mahwah, NJ: Lawrence Erlbaum.

- Parisi, D., Cecconi, F., & Natale, F. (2003). Cultural change in spatial environments. The role of cultural assimilation and internal changes in cultures. *Journal of Conflict Resolution*, *47*, 163–179.
- Shibanai, Y., Yasuno, S., & Ishiguro, I. (2001). Effects of global information feedback on diversity: extensions to Axelrod's adaptive culture model. *Journal of Conflict Resolution*, *45*, 80–96.
- Watts, D. J. & Strogatz, S. (1998). Collective dynamics of “small-world” networks. *Nature*, *393*, 440–442.
- Watts, D. J. (1999). Network dynamics and the small-world phenomenon. *American Journal of Sociology*, *105*, 493–527.