The twin impact of homophily and accessibility on ideological polarization

Stephen Davies
University of Mary Washington
Fredericksburg, VA
stephen@umw.edu

ABSTRACT

We present an agent-based model to explore the causes of one aspect of ideological polarization: the extent to which members of a society have social ties only with those they agree with. Specifically, we look at two variables that affect how an artificial social network structure is built: *homophily*, or the preference of individuals to form connections with others of the same "kind"; and *accessibility*, or the ease with which agents can form connections to others distant from it, as opposed to only local agents in its immediate vicinity. Our model builds a graph according to these two parameters, and then executes the classic Binary Voter Model (BVM) process on it whereby connected nodes influence one another's opinions. We find that counter to our original hypothesis, increasing the society's accessibility decreases its polarization, especially for high levels of homophily. Also, we discover that the rate at which agents form and dissolve friendships during the simulation plays a nuanced role in the way the society evolves.

CCS CONCEPTS

• **Computing methodologies → Agent / discrete models**;

KEYWORDS

Opinion Dynamics models, political polarization

ACM Reference Format:
Stephen Davies. 2017. The twin impact of homophily and accessibility on ideological polarization. In *Proceedings of The Computational Social Science Society of the Americas (CSSSA'17).* ACM, New York, NY, USA, 8 pages. https://doi.org/10.475/123_4

1 INTRODUCTION

1.1 Two sources of friendship

Consider for a moment the friendships you have had through which meaningful mutual influence has taken place. In broad terms, these relationships can be thought of as coming from two different kinds of sources. In the first kind, you did not originally encounter the acquaintances by specifically seeking them out – they rather made an appearance in your life due to circumstances. When you were a child, there were other children down the street, and fellow students in your 1st-grade classroom. When you were older, there were dormitory hall-mates, new neighbors in new neighborhoods, and co-workers. Some of these people who by happenstance wandered into your field of view, you formed meaningful friendships with. Others you did not. But the key point that distinguishes them from the second group is that you did not discover them by seeking them out based any of their attributes. Instead, a minuscule set of people (out of all the people in the world) simply fell into your lap. And from that set, you became meaningful friends with some.

The second group consists of those you encountered *because* you had something in common, and you discovered them because you deliberately sought that something. Even before the advent of cheap, electronic communication this occurred: people attended churches and synagogues, knowing they would find others with similar worldviews. They joined rotary clubs, parent-teacher organizations, and political parties, seeking out those with similar interests. In the Internet age, this is multiplied tenfold. Anyone with any view or interest, no matter how esoteric, can find a chatroom, website, Google group, subreddit, Twitter community, or other form of online clique devoted to it. For the younger generation especially, this is an important group: recent surveys suggest that a majority of American teens form meaningful friendships online, and communicate with them daily, often never meeting the friend physically.[28]

The key point is that some of your friends were "chosen" from a very tiny subset of the people in the world who you couldn't help but run into. You might not have had anything in particular in common with them, other than geographical proximity. The others, however, you drew from a very large pool – essentially the entire online world.[1] And the tools of the Internet give you breathtaking precision with which to find and select such friends.

1.2 "Accessibility"

Admittedly, these two sources of friendship are idealized points along a continuum. Any individual relationship may have been formed due to some blend of the two. Nevertheless, it is a useful abstraction, and suggests the existence of a key parameter in modeling social networks: the relative strength of the two sources in leading to friendships. Put another way: on average, what fraction of a person's friendships develop as a result of geographical proximity and happenstance encounters, versus being due to more distant relationships where parties sought each other out based on some shared attribute?

In the language of this paper, we use the term **local** to describe friendships of the first sort, and **global** for the second. In our model

[1]Note that once a relationship has been established, the partners may never again actively consider the source of the friendship: a link has simply been made, regardless of the cause, and thereafter results in mutual influence.

CSSSA'17, October 2017, Santa Fe, New Mexico USA

ACM ISBN 123-4567-24-567/08/06…$15.00
https://doi.org/10.475/123_4

we are concerned with the relative likelihood of friendships developing from each of these two sources, and we denote as the society's **accessibility** A ($0 \leq A \leq 1$) the fraction of relationships that have *global* origination. A society with an A of 0 presents no way for its citizens to discover anyone outside the local social circle they happened to inherit. Persons in a society with $A = 1$, by contrast, have no tendency to form friendships with those "nearby" them any more than with anyone else in the world; it is as if their entire experience took place through the Internet, with equal access to all others, and the ability to search for others based on their attributes, but with all geographic information hidden.

For simplicity, we model A as a constant value throughout a society, rather than giving it a different value for each agent (as it certainly has in reality). Presumably the A of the western world in 2017 is higher than the A of thirty years ago. One question this paper tries to answer is how this might bear on ideological polarization.

1.3 The reality of homophily

One might expect that as the accessibility A of a society increases, so would the diversity of viewpoints that its citizens are exposed to. After all, a higher A means that more of a person's friendships will be drawn from the entire global field, giving them a much broader range of exposure to people with many different views and interests.

Arguing against that outcome, however, is an indisputable fact about human nature: *homophily*. One of the most reliable and long-standing observations in social psychology, homophily simply refers to the tendency for people to prefer others who are similar to them.[30] This is true across many different aspects of "similarity," whether it be race, age, religion, occupation, political affiliation, values, or common interests. Given the choice of forming ties with several individuals, people tend to choose the one(s) whom they perceive as being most like them.

A greater number of choices of friends, therefore, may well lead to *less* diversity within one's social circle. As with accessibility, we model the homophily H of an entire society with a single value, $0 \leq H \leq 1$, that influences friendship choices. In this case, an H of 0.5 is *neutral*: agents in the model have no preference for or against forming friendships with similar agents. At the extreme of $H = 1$, agents will always choose an agent similar to them, if possible, and at $H = 0$, they will always choose a *dis*similar agent. (This situation could be termed "heterophily.")

The interplay of these two parameters A and H and their impact on the ideological polarization of a society is the subject of this paper.

1.4 Defining a healthy society

The term "polarization" – often with the modifier "political" – abounds in recent discussion of the U.S. and other political environments[9, 18, 22, 29]. It is nearly always used with a negative connotation. Claims that the degree of polarization is increasing in western cultures have been substantiated in some ways by academics (*e.g.*, [1, 6, 34]) although with caveats ([2, 5, 20]). Numerous studies have investigated how it takes root in social networking and other online environments [3, 11, 25, 31].

Defining polarization, however, is somewhat tricky; one recent paper, in fact, spelled out nine different possible definitions[8]. Most often, the term is associated with large *differences* in views between subpopulations, especially when those views are perceived as *extreme*. If, when responding to a question on a 6-point Likert scale, half the population answered "1" and the other half answered "6," this would typically be viewed as a "polarized" population. If half answered "3" and the other half "4," on the other hand, or if equal numbers of people gave each of the six responses, that would be seen as less polarized.

In this work, however, we look at a different form of "polarization": *the extent to which adherents of one viewpoint tend to form social connections only with others of that same viewpoint.* Whether the viewpoints are themselves "extreme" on any objective scale is irrelevant, as is the percentage of individuals subscribing to each of the various viewpoints. What matters for our purposes is whether the adherents of various views form isolated pockets of communication, or "echo chambers," rather than having broad social connections with people of a variety of different opinions.

Under this interpretation, members of a society holding strong or even "extreme" views is not a negative outcome. What is important is that the members maintain fruitful dialogue with one another, and are continually exposed to views different from their own. Members sequestering themselves into ideological cliques is unhealthy. But members thoughtfully choosing to retain their opinion even in the constant and active presence of others articulating counterpoints to it is not.

We therefore use the social network's *assortativity* coefficient[32] as the key measure of polarization. If we model a social network as an undirected graph whose nodes each possess a nominal "ideology" attribute taken from a small set of possible ideologies, the assortativity gives a measure of what fraction of the edges are between likeminded nodes, compared to what we would expect if the edges were simply dispersed at random.[2]

2 RELATED WORK

Opinion Dynamics (OD) models have a robust tradition, often traced to the Binary Voter Model (BVM) of Holley and Liggett[27] and Clifford and Sudbury[10]. OD models seek to reproduce the phenomenon of individual agents forming opinions over time via mutual influence, and to draw conclusions about the overall pattern of opinions that may emerge in a society as a consequence of certain micro-behaviors.

Axelrod's work in this area[4] represented agents with multiple discrete-valued attributes occupying a cellular grid. Agents were more likely to influence one another when they had more attribute values already in common, imitating what Axelrod called "the fundamental principle of human communication": that influence occurs

[2]Formally, the assortativity coefficient of a graph is a value between -1 and 1 which is computed as follows. Let e_{ij} be the fraction of all edges in the graph which connect an agent with ideology i and an agent with ideology j, where i and j range over all pairs of possible ideologies. Let $\mathbf{e}$ be the matrix whose elements are e_{ij}, $\mathbf{x}^2$ indicate matrix multiplication, $\mathrm{Tr}\,\mathbf{x}$ be the sum of the main diagonal elements of $\mathbf{x}$, and $\|\mathbf{x}\|$ be the sum of the elements of the matrix $\mathbf{x}$. The assortativity is then $\frac{\mathrm{Tr}\,\mathbf{e}-\|\mathbf{e}^2\|}{1-\|\mathbf{e}^2\|}$. It has the value 1 when there is perfect assortative mixing (*i.e.*, *all* edges are between nodes with the same ideology), 0 when there is no assortative mixing (the ideology of the nodes has no bearing on whether they will be connected), and a negative value when nodes tend to connect to others of a *different* ideology.

more frequently between people who perceive themselves as being already fairly alike. The effect of influence in the model was to copy one of the differing attributes from one agent to the other, further increasing their similarity. Among other results, Axelrod demonstrated that as the *range* of influence increases (*i.e.*, as agents are able to interact directly with other agents 2, 3, 4, . . . squares away), the degree of overall homogeneity in the society increases. He measured homogeneity as the number of distinct, geographically isolated clusters of agents with the same attribute values.

Axelrod's result might lead us to predict that polarization would *decrease* with accessibility, rather than increase. Shibani *et al.*[35] and Grieg[24] found that subtle changes to the model, however, can produce the opposite result. Too, Flache and Macy[21] concluded that Axelrod's original result depended crucially on the opinions being discrete valued; when continuous opinions were used, and adjustments could be made to them gradually, polarization actually increased with range of influence.

One popular approach to modeling *continuous* opinion dynamics is the Bounded Confidence (BC) model (originally in [16, 26]). This assumes that agents will be influenced only by the opinions of others that their own opinion is already sufficiently "close" to (*i.e.*, within some threshold ϵ); other opinions are viewed as too extreme from one's own, and therefore untrustworthy. (This is similar in spirit to Axelrod's "fundamental principle," but in the context of a single attribute, not an array of them.) The result of such influence, when it does occur, is an averaging operation that pulls each agent's opinion closer to the other. The BC mechanism is one way of preventing a graph of continuous opinions from converging to absolute homogeneity, as will happen if the averaging operation happens unconditionally.

All of these studies inspired by Axelrod place agents on a rectangular grid. Work has been done, too, on agents connected via a more general network/graph structure, whether a complete graph, possibly with edge weights[15, 17], or a general graph[13]. Several of these studies have explored the interplay between homophily and (various definitions of) polarization in the context of continuous-valued opinions. Dandekar *et al.*[13] in particular use a variant of the assortativity coefficient called the network degree index (for continuous attributes). However, their work was still based on continuous attributes, and they invoke a more complex opinion formation process than we do, incorporating confirmation bias.

Recently, Gargiulo and Gandica[23] explored the connection between homophily and polarization in the context of continuous-valued opinions under a BC dynamic. To build the initial graph for their simulation, they extend the well-known preferential attachment mechanism [7] to incorporate homophily: when a new node chooses which existing node to connect to in the graph, it incorporates information not only about the degree of existing nodes, but also about the similarity of their opinions to its own. In this way, not only are nodes with more neighbors more likely to be chosen for attachment (as in [7], [16]), but nodes with similar opinions to the new node are also more likely.

Gargiulo and Gandica discovered that under these assumptions, as homophily is increased (*i.e.*, as similarity is weighed more heavily than degree when attaching new nodes), polarization – measured as the number of distinct opinion clusters at equilibrium – *decreases*. This counterintuitive result can be explained as follows: when homophily is low, a node does not form as many initial connections to nodes with similar opinions (within the threshold ϵ of its own). Therefore, it is more likely that that node will get "stuck" during the bounded confidence process: since it can't find many neighbors whose opinions seem plausible, it stubbornly sticks to its own. Increasing homophily equips each node with more neighbors whose opinions are close to its own, such that it can gradually be nudged towards the emerging consensus.

Our work differs from these other studies most significantly in that we measure "polarization" as the graph's assortativity, rather than as the number of distinct opinion clusters that emerge, or the extremity of views. Dandekar *et al.*'s is the only work we are aware of that focuses on something akin to the assortativity coefficient, and their model is quite different: it uses continuous-valued attributes, weighted edges, confirmation bias, a starting graph based on a stochastic block model, and a BC-like process of repeated opinion averaging. Ours instead models a BVM process on discrete attributes, with the starting graph produced through a local/global morphogenesis process as previously described.

3 THE MODEL

We present the model using an abbreviated version of the ODD protocol [33].

3.1 Purpose

This abstract agent-based model simulates an evolving social network of agents that each possess an "ideology" attribute representing some opinion. As it evolves, nodes are randomly chosen to influence their neighbors by propagating their ideologies to them. The "polarization" of the network, measured as the tendency of nodes to be connected to others with the same ideology, will change as this process takes place.

The purpose of the model is to investigate the effects on polarization of three parameters: two which affect how the initial graph is constructed, and one which affects the graph structure as it evolves. The first, "**accessibility**," controls the set of possible neighbors that a node chooses from each time it adds a friendship during graph construction. When accessibility is low, neighbors will more often be chosen from a small (random) pool that is initially available to a node. When it is high, neighbors will more often be chosen from the entire network at large. Accessibility thus models the extent to which a society's citizens form **local** friendships (*e.g.*, geographical neighbors) versus **global** ones (*e.g.*, over the Internet).

The second parameter, "**homophily**," controls the strength of each node's preference to attach to other nodes with the same ideology. When it is high, nodes will almost always form connections to nodes that have the same ideology, if possible. This is true regardless of whether they are formed locally or globally. When low, nodes prefer to attach to nodes with *different* ideologies, and when medium, they are indifferent.

Finally, as the simulation progresses, nodes might make new friendships and dissolve old ones. This is to reflect the fact that not only do real-world persons' ideologies fluctuate over time, but so do their social connections. If this model variant is enabled, the "**rebalancing rate**" parameter controls the frequency with which a node re-evaluates its friendships in light of its tendency to homophily.

With this model, we hope to gain general insight into how societies that provide different means of communication and discovery, and societies that encourage different levels of tolerance for opposing opinions, may differ in the prevalence of "echo chambers" whereby dissenting views are rarely heard.

3.2 Entities, State Variables and Scales

The entities in the model are Agents, and have the following attributes:

ID A unique ID number.

Ideology One of a discrete set of possible ideologies, represented as integers $0, 1, \ldots I$. The ideology of an Agent will change during the simulation, possibly many times, as it is influenced by its neighbors.

Friendships A set of references to other Agent entities with whom this Agent has a social connection. The entire set of Agents and their Friendships form an *undirected* graph: if Agent X is friends with Agent Y, Agent Y is also friends with Agent X. Once the initial graph is built, each Agent's set of Friendships is fixed over the lifetime of the model (unless the **DynamicRebalancing** Policy is enabled; see below.)

LocalAssociates A set of references to other Agent entities with whom this Agent *may* form a friendship when it chooses locally (see below). Conceptually, these represent the (small, relative to the whole population) group of other people to whom an agent is geographically proximate. Some of these may actually become the agent's friends; perhaps many or all of them if the society has low accessibility. Importantly, once the initial graph is constructed, the LocalAssociates attribute is discarded and no longer used in the simulation.

3.3 Process Overview and Scheduling

Once the initial graph is built (see "Initialization," below), we carry out the standard Binary Voter Model (BVM) process on it, using the **selection with replacement** and **neighbor influences node** variants (see [14].) At each of T iterations:

(1) An agent X is chosen at random from the entire graph. (This is the **selection with replacement** variant; see [14].)
(2) One of its friends Y is chosen at random from its set of Friendships. (If the agent has no Friendships, skip this iteration.)
(3) If Y's current ideology is different than X's, copy Y's to X's. (The **neighbor influences node** variant; see [14].)
(4) If the **DynamicRebalancing** Policy is enabled, and its rate R is such that it should take place now, carry out the **DynamicallyRebalance** on agent X (see below).
(5) Compute and store the graph's assortativity coefficient with respect to Ideology (using the `assortativity_nominal` function from the `igraph` Python package[12]).

3.4 Initialization (morphogenesis)

When the simulation begins, generate an undirected Erdos-Renyi random graph[19] with N nodes and probability of edge connection p. Call this the **LocalAssociatesGraph**. Each node represents an Agent, and the edges of this graph constitute its **LocalAssociates**. With uniform probability, assign each node an initial Ideology from the set of I ideologies.

Then, generate the **Friendships** graph as follows:

(1) Create an empty **Friendships** graph with N nodes (agents).
(2) For each node X (in ID order), generate a_X friendship connections to other nodes, where a_X is the degree of node X in the **LocalAssociatesGraph**. Choose the node for each connection as follows:
   - With probability A (the accessibility parameter), select a node (without replacement) from X's **LocalAssociates**, weighted by H (the homophily parameter, see below).
   - With probability $1-A$, select a node (without replacement) from *all* nodes, weighted by H.

   In either case, if there are no nodes available with whom X is not already friends, skip the step.

"Weighted by H" means that candidate nodes whose ideologies are the same as node X's are assigned an (unnormalized) probability of H to be selected, and candidate nodes of different ideologies are assigned $1 - H$. All probabilities are then normalized to sum to 1, and a node is chosen.[3]

3.5 Submodels

DynamicRebalancing. If this optional Policy is enabled, a "rebalance rate" parameter R ($0 \leq R \leq 1$) controls how frequently the **DynamicallyRebalance** submodel is executed. R=1 means it will execute *every* iteration of the main simulation loop (immediately after copying the neighbor's ideology.) R=0 means it will never execute, and any value in between means it will execute some fraction of the iterations.

When it executes for a node X, the **DynamicallyRebalance** process is as follows:

(1) Perform *one* of the following:
   - With probability .5, connect a new node to X, chosen (without replacement) from all other nodes, weighted by the homophily H (exactly as above).[4]
   - With probability .5, *dis*connect a node from X (breaking the friendship), chosen from its current Friends, weighted by $1-H$. This means that for $H > .5$, nodes with the same ideology as X are *less* likely to be chosen than nodes with different ideologies.

All of the numerical parameters are configurable. The simulation's default values for them are T=500, I=3, N=50, p=.06, A=.5, H=.7, R=0, and **DynamicRebalancing** disabled. The code is written in Python, is open source, and available at https://github.com/hzontine/polarbear/tree/master/wide.

4 HYPOTHESES

We form the following hypotheses about the above model's behavior.

[3] Therefore, if there are c_X candidate nodes for X to connect to, s_X of whose ideologies match X, the probability that each same-ideology node will be chosen is αH and the probability for each different-ideology node is $\alpha(1-H)$, where α is the normalizing constant $\frac{1}{s_X H + (c_X - s_X)(1-H)}$.

[4] In the unlikely event that X is already connected to all other nodes, skip this step.

Hypothesis 1a (H_{1a}): For a given level of accessibility A, the polarization P *of the initial graph* will increase with homophily H.

Hypothesis 1b (H_{1b}): For a given A, P will *continue* to increase with H as the BVM process takes place on the graph.

Hypothesis 2a (H_{2a}): For a given H>.5, the P *of the initial graph* will increase with A.

Hypothesis 2b (H_{2b}): For a given H>.5, P will *continue* to increase with A as the BVM process takes place on the graph.

Hypothesis 3 (H_3): If the **DynamicRebalancing** policy is enabled, and H>.5, P will increase more rapidly during the BVM process, and more rapidly with increasing R. This will be unaffected by A.

We expect H_{1a} to hold simply because as H increases, the preference for nodes to attach to others of the same ideology increases, which should increase assortativity. Less obviously, H_{2a} should hold because when A increases nodes have greater freedom of choice in who they connect to. Low values of A mean that as the graph is built, nodes will often be "forced" to choose a friend from the limited number present in their LocalAssociates set, and will thus more often have no choice but to make a friend of a different ideology. Thus we expect more cross-ideology friendships when A is low. (Obviously this only holds when the homophily is greater than .5, indicating a preference for same-ideology nodes.)

The rationale behind H_{1b} and H_{2b} is as follows. When the initial graph is constructed so as to be more polarized, the BVM will have more raw material to work with in order to make it further polarized. We expect increasing returns as the structure of the network, already built to put like-minded nodes largely together, causes the remaining vestiges of local disagreement to be snuffed out.

Finally, we predict H_3 because dynamically rebalancing should have a strictly positive impact on assortativity (when $H > .5$). New edges will be added to the graph, and old edges dropped, with preference to forming/maintaining friendships with likeminded nodes. This should accelerate the degree to which the graph becomes polarized, and the rate of acceleration should increase when we add/remove friendships more often. We have no *a priori* reason to suspect that the A used to form the initial graph will come into play.

5 RESULTS

5.1 H_{1a} and H_{2a}

To verify H_{1a} and H_{2a}, we need only consider the starting graph. We generate a large number of initial graphs according to the process described in section 3.4, using a range of values of H and A, and measure their assortativity. Figure 1 shows a box plot of the results of generating 1000 such graphs for each combination of six accessibility values and six homophily values. (For these and all other results in this paper, N was set to 50 agents and I to 2 ideologies.) As expected, the polarization of these initial graphs clearly increases with both H and A, establishing H_{1a} and H_{2a}.

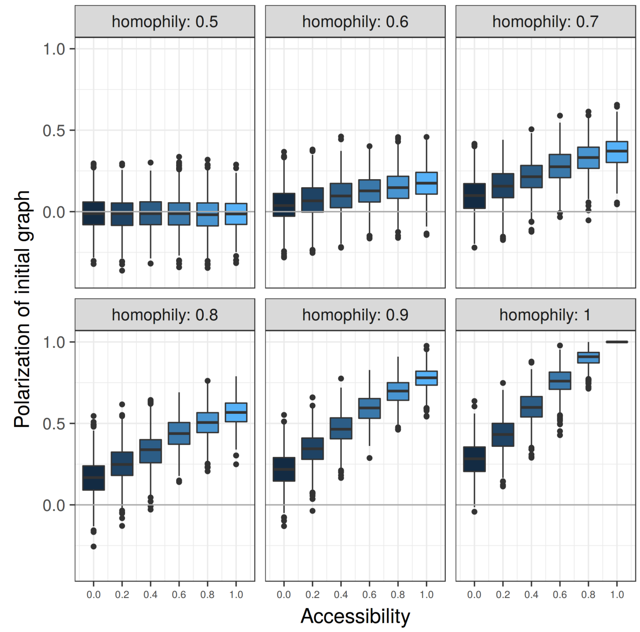

Figure 1: Polarizations of initial graphs, defined as nominal assortativity on the ideology attribute, for varying levels of homophily H and accessibility A. The box for each pair of values represents 1000 randomly generated starting graphs.

5.2 H_{1b} and H_{2b}

Hypotheses H_{1b} and H_{2b} are a bit trickier to evaluate, since we seek to discover whether the polarization of "the graph" increases as a result of the BVM. But of course the BVM produces an entire time series of graphs, not a single one. Several approaches are possible: we could take a snapshot of the graph at some fixed number of iterations, and measure the assortativity at that point; we could take the maximum (or minimum) assortativity anywhere in the sequence; we could compute the mean assortativity of all graphs produced during the process; *etc.* We choose the last of these approaches here, but do not begin computing the mean until the 50^{th} of 200 iterations. This admittedly arbitrary boundary point is an attempt to measure the assortativity only after the process has had a chance to emerge from a cold start. To summarize, then: we choose our measure of "the polarization that the BVM process induces" to be the mean assortativity of the graph at iterations 50 through 200 of the BVM process.

The result, quite surprising to us, is in the top half of Figure 2. For high levels of homophily ($H \geq .8$), additional accessibility does result in higher polarization. But this effect seems to be less than it was on the initial graph, and in fact for moderate levels of homophily ($.5 \leq H \leq .7$), higher accessibility actually *lowers* the polarization.

To get a clearer view of the BVM's effect, we "normalize" these iterations-50-through-200 polarization values by subtracting each *initial* graph's polarization from them. In this manner, we isolate the effect of the BVM process (section 3.3) from the effect of the initial graph-construction process (section 3.4). The result is the bottom plot in Figure 2. Clearly, for most values of homophily, the accessibility has a *moderating* effect, rather than an amplifying effect, on the graph's polarization.

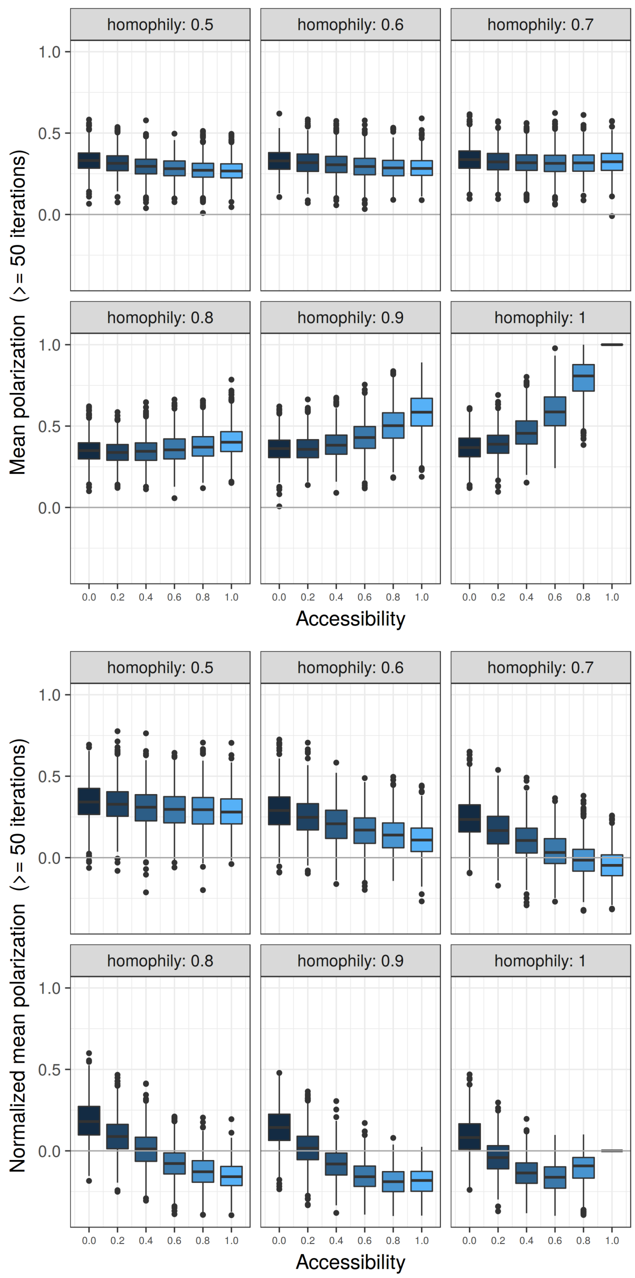

Figure 2: Top: mean polarization of the graph at iterations 50 through 200 of the BVM process. Bottom: the "normalized" polarization, defined as the difference between the mean polarization (top) and the initial polarization (Figure 1).

We thus not only fail to verify H_{1b} and H_{2b}, but we radically refute them. The opposite is true.

5.3 H_3

To evaluate H_3, we compare the normalized polarizations of simulations that have **DynamicRebalancing** enabled with those that don't. The results are presented in Figure 3 (for clarity, we omit outliers from the plots and show DynamicRebalancing at only a single rate, R=1.) Interestingly, hypothesis H_3 is confirmed only for *some* homphily values. When $H \leq .6$, dynamic rebalancing acts as a moderating effect on polarization. Only for values of .7 and above does it increase the polarization relative to the initial graph, as predicted.

This effect can possibly be explained by the small size of the graph (only 50 nodes). There are a fixed number of other nodes in the graph with the same ideology as a given node X. Therefore, the more likeminded friends a node has, the fewer additional nodes with that Ideology remain in the pool. If forced to dissolve friendships often and replace them with new friends (which happens for large values of R), there will be proportionately fewer similarly likeminded nodes to replace that broken friendship with. Only if the homophily is so high that X aggressively cherry-picks likeminded neighbors with strong preference will it be able to overcome this tendency prevent the polarization from drifting lower. (Admittedly, this is only a conjecture about the reasons for this puzzling effect.)

6 DISCUSSION AND CONCLUSION

The fact that H_{1b} and H_{2b} were refuted so soundly frankly astonished us. But perhaps considering the literature in this area we should not be so surprised. Different studies based on rather subtle extensions to the original Axelrod model have come to very different conclusions as to whether more accessibility will increase or decrease polarization. These differences appear to depend on exactly how accessibility, polarization, homophily, and the attribute in question are defined.

Clearly there is a delicate interplay here, at the heart of which remains an unresolved question about what the end result of enhanced communication will be for a society. Will the expanded freedom of selection lead to people simply forming more homogeneous factions? Or will the greater exposure to more "remote" parts of the society result in greater diversity of one's social groups? Our model suggests the latter is more true, but the answer does not appear to be simple. A careful study of the models presented in section 2 is called for in order to tease out which specific differences are responsible for which effects. Then, the social psychology question can be applied: how to synthesize aspects of these models to arrive at a composite, more complex description of how human beings actually act?

As for H_3, the interaction between variables is interesting and non-trivial. When agents exhibit a strong preference to form ties with their own "type" and dissolve those with others, then changing friends often will accelerate polarization, as expected. But if this homophily preference is milder, not only will frequently changing friends fail to exacerbate polarization, in fact it will dampen it.

Perhaps one interpretation is as follows. If a society is to be tolerant, its citizens must each *either* be open to friendships with people of opposing views, *or* be reluctant to dissolve ties and form new ones. This somewhat tangled statement becomes intelligible if we consider that the freedom to change friends freely is a dangerous tool in the hands of someone with strong homophily. He is likely to use that tool to aggressively seek only duplicates of himself. A person with low homophily, on the other hand, is "safe" to entrust

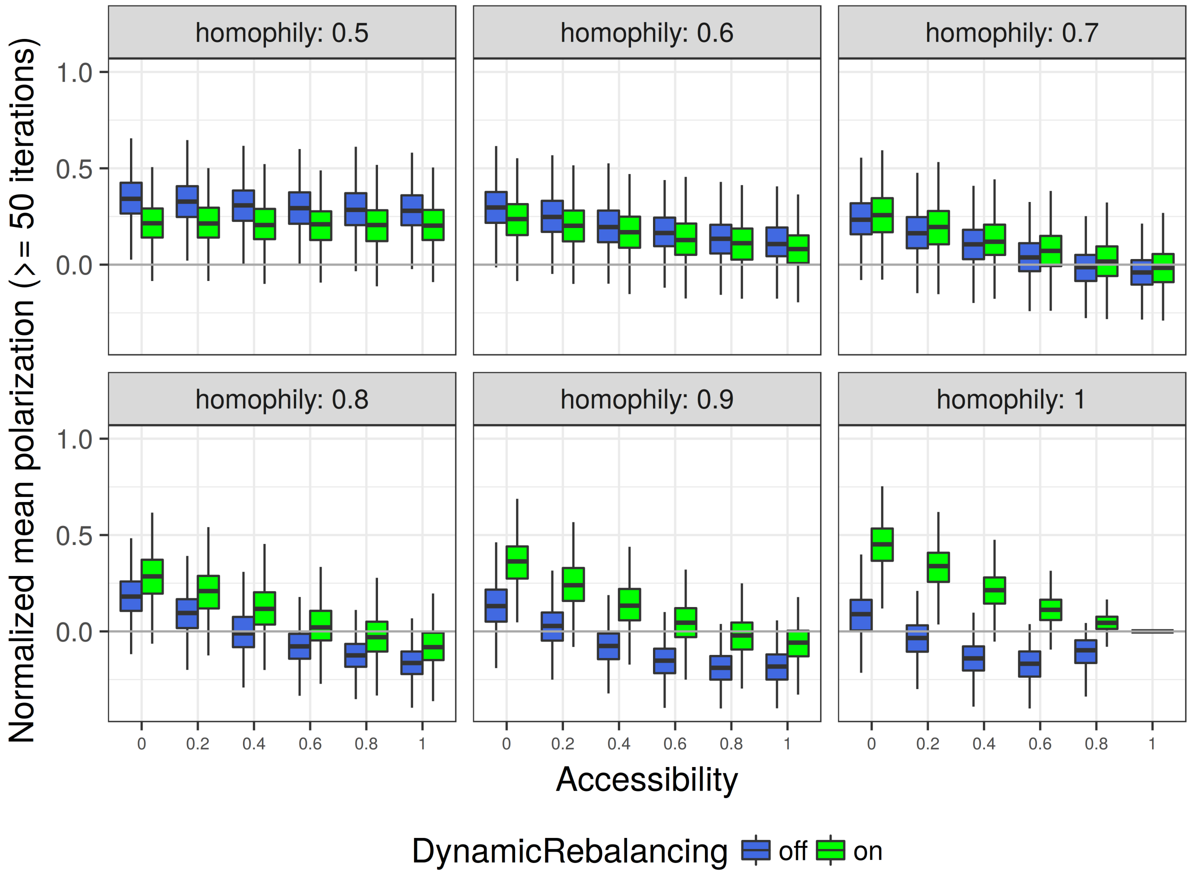

Figure 3: Effect of the DynamicRebalancing policy on (normalized) mean polarization. The blue boxplots are identical to those in the bottom half of Figure 2. The green boxplots use the same starting graphs, but with dynamic rebalancing enabled at a rate of R=1.

that freedom to; she may change friends often, but that will be a good thing, since she will form diverse ties.

REFERENCES

[1] Alan I. Abramowitz. 2015. The New American Electorate. *American Gridlock: The Sources, Character, and Impact of Political Polarization* (2015), 19.

[2] Samuel J. Abrams and Morris P. Fiorina. 2015. Party sorting: The foundation of polarized politics. *American gridlock: the sources, character, and impact of polarization* (2015), 113–29.

[3] Lada A. Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*. ACM, 36–43. http://dl.acm.org/citation.cfm?id=1134277

[4] Robert Axelrod. 1997. The dissemination of culture: A model with local convergence and global polarization. *Journal of conflict resolution* 41, 2 (1997), 203–226. http://journals.sagepub.com/doi/abs/10.1177/0022002797041002001

[5] Delia Baldassarri and Peter Bearman. 2007. Dynamics of political polarization. *American sociological review* 72, 5 (2007), 784–811. http://journals.sagepub.com/doi/abs/10.1177/000312240707200507

[6] Delia Baldassarri and Andrew Gelman. 2008. Partisans without constraint: Political polarization and trends in American public opinion. *Amer. J. Sociology* 114, 2 (2008), 408–446. http://www.journals.uchicago.edu/doi/abs/10.1086/590649

[7] Albert-Laszlo Barabasi and Reka Albert. 1999. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509–512. http://science.sciencemag.org/content/286/5439/509.short

[8] Aaron Bramson, Patrick Grim, Daniel J. Singer, Steven Fisher, William Berger, Graham Sack, and Carissa Flocken. 2016. Disambiguation of social polarization concepts and measures. *Journal of Mathematical Sociology* 40, 2 (April 2016), 80–111. https://doi.org/10.1080/0022250X.2016.1147443

[9] James Campbell. 2016. The source of America's political polarization? It's us. *Los Angeles Times* (June 2016). http://www.latimes.com/opinion/op-ed/la-oe-campbell-political-polarization-20160627-snap-story.html

[10] Peter Clifford and Aidan Sudbury. 1973. A Model for Spatial Conflict. *Biometrika* 60, 3 (1973), 581–588. https://doi.org/10.2307/2335008

[11] Michael Conover, Jacob Ratkiewicz, Matthew R. Francisco, Bruno Goncalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. *ICWSM* 133 (2011), 89–96. http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/download/2847/3275

[12] Gabor Csardi and Tamas Nepusz. 2006. The igraph software package for complex network research. *InterJournal* Complex Systems (2006), 1695. http://igraph.org

[13] Pranav Dandekar, Ashish Goel, and David T. Lee. 2013. Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences* 110, 15 (2013), 5791–5796. http://www.pnas.org/content/110/15/5791.short

[14] Stephen Davies and Hannah Zontine. 2016. The Computational Social Science Society of the Americas CSSSA 2016. Santa Fe, New Mexico. http://cs.umw.edu/~stephen/writings/DaviesZontine2016.pdf

[15] Guillaume Deffuant, Frederic Amblard, Gerard Weisbuch, and Thierry Faure. 2002. How can extremism prevail? A study based on the relative agreement interaction model. *Journal of artificial societies and social simulation* 5, 4 (2002). http://jasss.soc.surrey.ac.uk/5/4/1.html

[16] Guillaume Deffuant, David Neau, Frederic Amblard, and Gerard Weisbuch. 2000. Mixing beliefs among interacting agents. *Advances in Complex Systems* 3, 01n04 (2000), 87–98. http://www.worldscientific.com/doi/abs/10.1142/

S0219525900000078
[17] Morris H. DeGroot. 1974. Reaching a consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121. http://www.tandfonline.com/doi/abs/10.1080/01621459.1974.10480137
[18] Michael Dimock, Carroll Doherty, Jocelyn Kiley, and Russ Oates. 2014. *Political Polarization in the American Public.* Technical Report. Pew Research Center. http://www.people-press.org/2014/06/12/political-polarization-in-the-american-public/
[19] Paul Erdos and Alfred Renyi. 1960. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci* 5, 1 (1960), 17–60. https://pdfs.semanticscholar.org/4201/73e087bca0bdb31985e28ff69c60a129c8ef.pdf
[20] Morris P. Fiorina and Samuel J. Abrams. 2008. Political polarization in the American public. *Annu. Rev. Polit. Sci.* 11 (2008), 563–588. http://www.annualreviews.org/doi/abs/10.1146/annurev.polisci.11.053106.153836
[21] Andreas Flache and Michael W. Macy. 2006. Why more contact may increase cultural polarization. *arXiv preprint physics/0604196* (2006). https://arxiv.org/abs/physics/0604196
[22] Davied French. 2017. We're Not in a Civil War, but We Are Drifting Toward Divorce. (June 2017). http://www.nationalreview.com/article/448385/liberal-conservative-divide-americans-self-segregate-culturally
[23] Floriana Gargiulo and Yerali Gandica. 2017. The Role of Homophily in the Emergence of Opinion Controversies. *Journal of Artificial Societies and Social Simulation* 20, 3 (2017), 8.
[24] J. Michael Greig. 2002. The end of geography? Globalization, communications, and culture in the international system. *Journal of Conflict Resolution* 46, 2 (2002), 225–243. http://journals.sagepub.com/doi/abs/10.1177/0022002702046002003
[25] Eszter Hargittai, Jason Gallo, and Matthew Kane. 2008. Cross-ideological discussions among conservative and liberal bloggers. *Public Choice* 134, 1 (2008), 67–86. http://www.springerlink.com/index/p7m41t21344130t7.pdf
[26] Rainer Hegselmann and Ulrich Krause. 2002. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation* 5, 3 (2002). http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.454.3597&rep=rep1&type=pdf
[27] Holley and Liggett. 1975. Ergodic theorems for weakly interacting systems and the voter model. (1975).
[28] Amanda Lenhart. 2015. Teens, Technology and Friendships. (Aug. 2015). http://www.pewinternet.org/2015/08/06/teens-technology-and-friendships/
[29] Nolan McCarty, Keith T. Poole, and Howard Rosenthal. 2016. *Polarized America: The dance of ideology and unequal riches.* mit Press.
[30] M. McPherson, L. Smith-Lovin, and J. M. Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* (2001), 415–444. http://www.jstor.org/stable/10.2307/2678628
[31] Seyedreza Mousavi and Bin Gu. 2014. The role of online social networks in political polarization. (2014). http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1559&context=amcis2014
[32] Mark EJ Newman. 2003. Mixing patterns in networks. *Physical Review E* 67, 2 (2003), 026126. https://journals.aps.org/pre/abstract/10.1103/PhysRevE.67.026126
[33] J. Gary Polhill, Dawn Parker, Daniel Brown, and Volker Grimm. 2008. Using the ODD protocol for describing three agent-based social simulation models of land-use change. *Journal of Artificial Societies and Social Simulation* 11, 2 (2008), 3. http://jasss.soc.surrey.ac.uk/11/2/3.html
[34] Markus Prior. 2013. Media and political polarization. *Annual Review of Political Science* 16 (2013), 101–127. http://www.annualreviews.org/doi/abs/10.1146/annurev-polisci-100711-135242
[35] Yasufumi Shibanai, Satoko Yasuno, and Itaru Ishiguro. 2001. Effects of global information feedback on diversity: extensions to Axelrod's adaptive culture model. *Journal of Conflict Resolution* 45, 1 (2001), 80–96. http://journals.sagepub.com/doi/abs/10.1177/0022002701045001004