Persuasion and Political Heterogeneity within Networks of Political Communication: Agent-Based Explanations for the Survival of Disagreement

Paul E. Johnson
Dept. of Political Science
University of Kansas
Lawrence, Kansas 66045
pauljohn@ukans.edu
http://lark.cc.ukans.edu/~pauljohn

Robert Huckfeldt
Dept. of Political Science
Indiana University
Bloomington,Indiana
huckfeld@indiana.edu

Aug 22, 2001

Abstract

Several simulation models of public opinion are investigated in this paper. We begin with an implementation of Robert Axelrod's "culture model." After a discussion of the implications of that model, we then consider several variants. The central question is this: under what conditions will diversity be preserved? That is to say, when should we expect social interaction to result in the eradication of differences among individuals?

*Prepared for presentation at the Annual Meetings of the American Political Science Association, San Francisco, August 30-Sept. 2, 2001. Special thanks to Rick Riolo of the University of Michigan and Marcus Daniels of the Santa Fe Institute.

I. Introduction

In his famous essay "The Dissemination of Culture: A Model with Local Convergence and Global Polarization," Robert Axelrod presented a model of social interaction and its impact on culture. Culture is thought of as a discrete list of numerical attributes, which we might consider as "issue positions" on public policies, preferred political parties, or any other judgment. One of the most interesting issues he raises is the possibility that, over the long run, the natural processes of human interaction will homogenize the culture. All individuals will have the same cultural attributes most of the time. Over the long run, nobody will interact with anybody who is even slightly different from them. Under some conditions, the model leads to the conclusion that all people will have the same opinion about everything. Under other conditions, social opinion is not completely homogenized, but the citizens are divided into completely homogeneous and totally distinct subgroups, and none of the members of one group interact with the members of any other group.

We have some concern about this "homogenization prediction," partly for empirical reasons and partly for normative reasons. On the empirical side, recent studies have turned up evidence of sustained diversity within interpersonal networks (Huckfeldt, Johnson, and Sprague, Forthcoming). On the normative side, this model has the implication that interaction does erase all differences. People who are concerned about the survival of ways of life that are in the minority might be led to radical policy prescriptions. For example, one might argue in favor of segregation, or cultural apartheid, as the only way to preserve diversity on the aggregate social level. We find this contention to be rather discouraging, and therefore we wish to re-examine the culture model.

Our approach takes the following steps. First, we describe the original Axelrod model and our software implementation of it. Second, we consider variants of the model which alter the model’s assumptions about the way people select each other for interaction. Finally, we consider variants of the model in which the assumptions about individual

II. The Culture model

Axelrod’s Culture Model (ACM) has a design which we illustrate in Figure 1. There is a square grid on which agents are distributed, one per cell.

[Figure 1 about here]

Axelrod calls these agents "villages". The agents are like villages in the sense that they are not allowed to move about on the grid. Each village is conceived of as a unitary actor. Each village has a set of discrete-valued features, the values of which are called traits. A feature is to be thought of as an opinion, issue stance, political party allegiance, or any other construct. Suppose there are 5 features, F_i, i={0,1,2,3,4}, and on each there are 3 possible positions, {0, 1, 2}. Each village has a feature array that is assigned randomly (uniformly) at the outset. Then a village’s cultural array might be (0,1,2,1,0) or (1,2,1,0,0), for example.

We have developed a simulation model, using the Swarm Simulation Toolkit (http://www.swarm.org) which allows us to explore the impact of various changes in the model. Like Swarm itself, the code is under the GNU Greater Public License, meaning that anybody who receives an executable version of the program also has a legal right to receive the source code itself. Our software is highly modular, meaning that it is possible to "plug and play" various components to measure their effects.

From a software design point of view, the model of individual behavior has two especially important components. First, there is a "selection process" that puts agents together for a possible interaction. Second, there are "opinion adjustment rules" which determine what happens when an interaction occurs. All of the rest of the model's design is a fairly standard Swarm model, with the agents implemented as separately instantiated objects that are contained in a collection and their behavior can be monitored in a number of ways.

In the original Axelrod approach, an agent is randomly selected, and a neighbor (from the von Neumann neighborhood, which consists of neighbors on the east, west, north, and south borders) is selected, and the agents interact with probability equal to the similarity of their features. If two villages share 2 of 5 features, then they interact with probability 0.40. If they interact, the agent copies one feature on which they differ from the other. Cells that lie on the outside boundaries of the grid have fewer neighbors with which to interact. We have implemented an option to have the grid "wrap around" to form a torus, and thus eliminate these edges, but have found nothing interesting to report on that component.

Since our implementation explores some variations on the ordering of interactions among agents (a general topic known in Swarm as "scheduling"), we have designed the code with many features that can be changed when the program is created (compiled) and when it is run . Axelrod's model repeatedly selects one agent at random and conducts the interaction. Because of anomalies observed in the follow-up study by Estein and Axtell (see Axtel, et. al, 1996), it is probably more prudent to employ a Knight's tour approach, in which the agents are randomly sorted and each is given an opportunity to interact before starting through the list again. We have implemented a run-time feature to select either the one-at-a-time (Axelrod's original) scheduling or the tour through the list. In this report, we present models in which the tour approach is used.

The number of cultural issue dimensions or topics, called "features" in the culture model, and the number of positions on each topic, called "traits" in his model, can be varied. Axelrod observed that when the number of traits is small, then the system-wide homogenization tendency is the greatest, while raising the number of traits tends to lead to the formation of homogeneous subgroups. We have chosen to focus most of our attention on the difficult case, one in which there is a small number of traits.

We have also introduced some new tools for measurement of conditions within the simulation itself. Axelrod's ideas of zone and region do not extend to the general settings that we explore in extensions of the model, so new measures of diversity are required.

We have created "objective" (aggregate level) and "subjective" (agent level) indicators of diversity. Objectively, we tally the features of the agents to measure the average and variance of opinion on each feature. We also calculate a system-wide diversity measure called entropy, which is also sometimes called Shannon's information index. That is a normed measure that is equal to 0 if all objects in a set are identical and 1 if every possible type is equally represented in the set (Shannon, 1949; Balch, 2000). If there are F different features (issue dimensions) and there are T different traits (positions) for each feature, then the number of possible issue stances is T^F. If the proportion agents holding a given set of positions is p_j, then the normed total entropy is given by:

$$\text{NormedTotalEntropy} = \frac{\sum_{j=1}^{T^F} p_j * \log_2(p_j)}{\log_2(1/T^F)}$$

The normed entropy measure depends on both the number of traits and the number of features.

Our subjective measures are built up from the experiences of the individual agents. Each agent keeps a "running tally" of its experiences. The code for this book-keeping makes the model take somewhat longer to run, but the substantive payoff justifies the effort. We present three individual-level measures. First, for each other agent that is encountered, the agent checks to see if the two agree about a randomly chosen feature. The proportion of encounters on which there is a shared feature is kept as a moving average that we call "acquaintance." Think of it as the individual's belief that it will have something in common with a randomly chosen other. In other words, they could become acquaintances. When an interaction occurs, the agents "compare notes" and find out how much they have in common. We use that information to construct the second and third measures. The level of "harmony" is the proportion of opinions that the two agents share. Each agent also keeps track of the proportion of others with whom it is identical. The "harmony" and "identical" indicators reflect information about agents only with whom an interaction occurs, while the "acquaintance" measure is collected for all agents contacted. Please note that these are not accurate summaries of the actual state of the system, but rather the experience-based beliefs of the individual agents. Generally, we calculate averages and standard deviations as summary measures of these measures.

The simulation model can be run interactively (with graphical displays of the grid and various conditions in the model) or it can be run an a batch mode. Interactive runs allow parameters to be adjusted "on the fly" and various on-screen displays can be used. When we run the model in batch mode, say for 100 or more runs, we use prespecified random number seeds. A batch of 100 runs under one set of conditions is thus directly comparable to a second batch of 100 runs because the random components that go into the creation of the start-up conditions for the models are the same. The model can be run for various parameter settings, but we are assured that the random initial conditions (features) of the i'th run of the model under a given set of conditions is equivalent to the i'th run of the model under another set of conditions. As a result, any difference between two runs can be attributed to changes which occur according to the logic of the model itself.

After a good deal of experimentation with this model, we conclude, as Axelrod did that, over the long run, there is not likely to be very much cultural diversity. While the tendency toward homogeneity is greater for some parameter settings than others, it is powerful in all cases. As Axelrod observed, the number of traits is a vital element. If there is only a small number of traits, then interaction is likely to occur (and change of individual opinion is certain). Repetition inevitably wipes out diversity. On the other hand, if the number of traits is huge, then two agents are unlikely to interact because they have nothing in common. Pockets of culture develop which are completely isolated from one another. Diversity is preserved in the aggregate sense, but none of the individual agents interact with others who are different from them, and their perception tends to be that all agents are identical and like themselves.

The summary of 100 runs of a model in which there are 5 features and 3 traits per feature is presented in Table 1. The numerical results are in-line with the original Axelrod results. The simulation continues until 10 passes are made through the list of agents without any changes of opinion. In every single run, the level of entropy was 0, meaning all citizens were identical in the end.

[Table 1 about here]

[Figure 2 about here]

Perhaps the summary statistics do not do justice to the results. The model generates output data in the form of time-plots of some variables and it also creates a set of color-coded "rasters" to display the state of the grid for each feature. A view of the distribution of opinion in the grid can be seen in the top part of Figure 2. The value of a feature at each spot in the grid is represented by a colored square. Values of the feature are color coded, with the dark representing the value of 0, and lighter shades used to indicate a progression of values. (In a model with more traits, we use all shades between the extremes.) As the model runs, it is readily apparent through time that the features are homogenized one-at-a-time. As soon as a single feature is homogenized, we know for certain that all features will be homogenized. This is so because any 2 agents in the society will have at least one feature in common, so they will interact, and the homogenization will continue. The process of homogenization accelerates as each piece of the puzzle falls into place. In Figure 2a, we show the starting conditions of each of the five features in its own grid. The typical outcome of the simulation is the second part of the figure, showing that each feature has been homogenized because the grid is filled with only one color.

In the bottom panel of Figure 2, we present the time plots of some of the diagnostic variables. Note that entropy begins at 0.8, near the maximum, and then it falls to a value of 0. Individual perceptions are roughly in line with that aggregate indicator. The perceptions of potential for acquaintance, harmony, and identicality begin with values near zero and gradually climb near 1.0. They don't quite reach 1.0 because the agents use a 20 period moving average to keep their records, and when the simulation stops after 10 periods without any culture changes, some agents can still remember past periods when disagreement was experienced. If we let the simulation continue ten more time steps, then the individual perception indicators would converge to 1.0.

In his article, Axelrod points out that it is possible to find parameter settings such that the grid is not homogenized, but rather subdivided into groups that do not interact with one another. That never happens when there are 5 features and 3 traits per feature, but it can happen if the number of traits is increased.

III. Selection Variations

The tendency of the Axelrod model to extinguish diversity is the source of trouble that is inspiring our effort. We wonder if it is possible to redesign the model so that interaction among differing types of people can occur without the homogenization effect taking place. We seek to understand the conditions under which, over the long term, entropy remains at some nonzero level. We hope also that agents who are not identical with each other interact with greater-than-zero frequency and, in addition, the agents believe that the levels of harmony and identicality are below 1. That is to say, it should not be the case that all agents think that everybody with whom they might interact is exactly like them.

As we mentioned above, the individual behavior model of the Axelrod model has two components, a selection component and an adjustment component. We believed at the outset that it would be possible to adjust the selection component of the model in order to preserve diversity. We have explored a number of variations, all of which are rather discouraging. In fact, we feel safe in saying that, short of erecting behavioral "firewalls" that completely block interaction among people who differ, there is not much that can be done to preserve diversity by tinkering with the selection part of the model.

III.A. Multi-agent villages

Consider what might happen if each village had not just one agent, but many agents, all of whom agree with each other at the outset, and they interact with one another frequently. Perhaps the interaction process might lead to the formation of "opinion clusters." In order to implement this version of the model, we have developed a general purpose multi-agent grid called MultiGrid2d, which allows us to create containers in the grid into which agents can be inserted and removed. We assign features at random to each village, but then insist that all agents within the village are identical.

The interaction component of the model requires us to introduce a new concept, parochialism. Parochialism is the likelihood that an agent will look within its own village when it seeks a partner for interaction. If the level of parochialism is sufficiently high, meaning agents interact with their "own kind" frequently enough, then maybe the homogenizing impact of contact with outsiders can be ameliorated.

We have run batches of simulations with multi-agent villages, and we find that, while the homogenization of opinion takes longer, it generally still occurs. Contrary to our hopes, this model does not lead to the development of self-maintaining clusters. Table 2 the summary statistics for one batch of run can be found. These results are for an extreme case in which parochialism is equal to 0.95, meaning that agents interact with people in their cell almost all of the time. When the simulation stops because no agent has changed its culture features for 10 successive time steps, the agents believe that they are in a highly harmonious environment.

The level of entropy drops to a very low, but nonzero value. Why not all the way to zero? The agents mostly interact within their own cells, so the differences are preserved most of the time. Still, in most runs of the model, the homogenization pressure through the 5 percent of interactions which occur outside the cell is strong enough to overcome the parochialism effect. We have found a few runs in which one cell on the edge of the grid remains distinct on one or more features. This outcome is the result of a "knife edge" balancing act. On the edge, the agents from that cell can interact with only a restricted number of neighbors, so their exposure is limited. Furthermore, by the "luck of the draw," when they consider interact with outsiders, they find nothing in common with them. This kind of "diversity" requires that all agents within a cell agree with each other exactly, and they avoid interaction with strangers assiduously.

[Table 2 about here]

There is, however, one very interesting property of these multi-agent grid simulations with high parochialism: the individual agents believe they are in a state of harmony thoughout the simulation, but aggregate diversity goes from the maximum to zero during the run of the model. Consider the extreme case in which parochialism is set at 0.95. In each cell we create 5 identical agents, and then we set them interacting with each other in random order, once per time step. In Figure 3, we show one of the runs of this model. Note that, when the simulation begins, the level of entropy is in fact high, because opinions are assigned randomly to the cells. However, agent perceptions of homogeneity are quite high because they frequently interact with other people within their cell, and at the start they are all identical. The measures of acquaintance, harmony, and identicality remain on the high side throughout the simulation. But, as time goes by, entropy goes down. Although the agents do not perceive it, the homogenization process is occurring. If the parochialism coefficient is set at a lower level, the homogenization occurs more quickly, of course, but the agents are more aware of their diversity as well.

III.B. Less selective agents.

One is struck by the fact that these simulations tend to homogenize opinion because people who agree with each other are more likely to interact. The "self-selection" process snowballs to create a situation in which the only people who interact are totally identical, and it appears to be just a matter of randomness whether or not there will be is one homogeneous society or a few homogeneous subgroups. Suppose that people were open to a broader array of interpersonal contacts? Wouldn't it be possible to forestall the homogenization by raising the interaction rate among people who would not ordinarily interact? It turns out the answer is a decisive "no".

Building on some ideas in the early work of Coleman (1964: chap. 16), we explored a variant of the model that causes interaction to occur more often among people who disagree. As in the original Axelrod model, at each time period, an individual encounters another individual, and an interaction occurs with probability proportional to similarity. If the encounter does produce an interaction, the search for an interaction partner is complete in that period. Alternatively, if the encounter fails to produce an interaction, there is some probability that the interaction will take place anyway. We have called the probability of interaction under these circumstances the Coleman coefficient. If the interaction does not take place, the individual continues to search up to 10 times until an interaction partner is located.

[Table 3 about here]

To people who expect the increased interaction among people who differ to preserve diversity, he results of this "Coleman model" are surprising. And disappointing. A summary of 100 runs for three values of the Coleman coefficient is presented in Table 3. Compared to the original Axelrod model with one agent per cell (which, implicitly, as the Coleman coefficient equal to 0), the number of iterations that occur before the culture is homogenized is reduced. Furthermore, as the Coleman coefficient is raised from 0.2 to 0.8, the number of iterations required to homogenize the society gets smaller and smaller. The conclusion, of course, is that the impact of encouraging interactions among dissimilar agents is the accelerated destruction of diversity. Or, viewed more positively, the ACM's component of self-selection serves to ***sustain*** diversity! People who ***avoid*** interaction with politically disagreeable encounters are acting to sustain their own beliefs (and, by extension, the beliefs of their neighbors). The attenuation of self-selection does not change the fact that, over the long haul, disagreement disappears. But the preservation of these small clusters tends to delay the process of political homogenization.

III.c. Neighborhoods and Workplaces

The work presented thus far has stayed within the fundamental constraints imposed by the Axelrod model. The agents, or villages, do not move around, they are never exposed to new context or new information. We wonder if redesigning the model so that agents move among contexts will help. We strive to create conditions under which agents can experience political interaction with a more-or-less unpredictable set of others in a variety of contexts. We pursue this strategy with the caveat that we don't want the agents to wander pointlessly seeking political interaction, but rather they wander with a purpose. Agents have "home" neighborhoods, where they might interact with family or neighbors, and they also interact with people in social occasions, such as the workplace, church, a labor union meeting, and so forth.

[Figure 4 about here]

The description of this next step in our modeling exercise requires an explanation of a very significant departure in the design of the simulation model. In Figure 4, we present a sketch of the geographical arrangement of the agents and their movement/interaction opportunities. There can be one or more "home grids," which are standard square arrays of agents. There can also be workplaces where agents might go and seek out interactions. In the home grid, there is some regularity to interaction, because the identity of the agent with whom one is interacting (if interaction occurs) is relatively more predictable. On the other hand, interactions in the work grids are less predictable. This is so because agents are unevenly "crowded" into the work grids. Agents from any home grid can be assigned to any of the cells in any of the work grids. Some cells in a work grid can have several agents assigned to them, while others might have none. In order to implement this in a meaningful way, it is necessary to redesign the scheduling substructure of the model to allow for movement among the milieux as well as the ability to search for available discussants and interact with them.

We have taken the following approach to scheduling agent actions. This is a discrete event simulation, but up to this point we have thought of a time step as a single pass through the list of agents. Now we take a more explicit approach of thinking of time as the passage of days, each of which is made up of a number of small time steps (say, hours). In the simulations described here, we (arbitrarily) set the number of steps per day to 10 and then we allow the agents to autonomously decide when to initiate interactions with neighbors who happen to be in their vicinity at a given instant. Each agent begins the day at home, and stays there, on average, for 5 time steps. Some agents spend almost no time at home, while others spend all their time at home. When the time comes to move to the work grid, then the agent removes itself from the home neighborhood and inserts itself into its designated work grid. During the time the agent is removed from one grid and "on the way to another," it cannot initiate any interactions and none of the other agents find it available for interaction. At the start of each day, each agent randomly selects a time step during which to seek an interaction. If the chosen time is before movement into the work grid, then the agent seeks interaction at home. The agents are moving about, but they still search for discussants in the von Neumann neighborhood.

The design which allows both movement and flexible interaction patterns diverges conceptually from the standard Axelrod model, of course. The redesign causes a several impacts on the selection of discussants. First, there is the simple act of movement. Agents who move take themselves out of the context of other agents and give themselves opportunities to find new contacts. The workplace offers a relatively unpredictable set of interactions because there can be many agents in each cell. Second, the home neighborhoods are discrete from one another. Since the residents of different home grids never interact directly with each other, each home grid is, at least on the surface, insulated from the homogenizing pressures observed in the other models. Perhaps we will generate neighborhoods in which there is local homogeneity (all people within grid i agree on an issue) but also global diversity (the agents in grid j disagree with the agents in grid i).

The results, which we now present, are somewhat disappointing. Using the standard Axelrod rules for the selection of discussion partners (modified as above for multiple-occupancy grids) and adjustment of opinion, we find that, over the long run, diversity is eradicated. This is true for any geographical setup which follows the home/work grid dichotomy we have described.

[Table 4 about here]

In Table 4, we give the summary statistics for a model in which there are five 10x10 home grids and three 5x5 work grids. One agent is assigned to each position in each home grid (meaning there are 500 agents in all), and agents are assigned positions in the work grids in a random way, meaning that some cells in the work grid are empty while others have several occupants. The rules for multi-agent grids described above were used. The parochialism coefficient has been set to 0.5, meaning that an agent who is in a cell with other agents will pick among those agents with equal likelihood one half of the time, and the rest of the time a discussant in one of the four neighbors will be chosen.

IIId. Identity-based Selection.

In the models we have considered thus far, we have found several likely culprits in our search for the explanation of the model's homogenizing tendency. First, agents are likely to have a basis upon which to interact--one or more features in common. Second, the agents interact if they have anything at all in common. If agent A has one out of five features in common with agent B, there is a 20 percent chance they will interact, and A will copy a feature from B. That can happen even if, on the other side of A, there is an agent that agrees with A about everything. Perhaps the shortcoming of this model is that it assumes agents initiate interactions with a randomly drawn individual even though there are more agreeable agents in their immediate vicinity.

We now reconceptualize the selection process in the following way. Rather than giving an agent the opportunity to interact with a randomly chosen neighbor, what if we allow agents to search their neighborhoods for discussants that they expect to be the most agreeable with them. Perhaps doing so will create enough selection bias to create and maintain diversity.

The model is implemented in the following way. Each agent keeps a record about each other that it has contacted. When the simulation proceeds, the agent searches its neighborhood and draws one possible contact from each neighboring cell. (If the agent's own cell contains other agents, one of them is selected at random). This creates a list of possible discussants. From this list, the agent then selects the one that is expected to have the most in common with it. If several in the list are equally "appealing," then one is selected at random with equal likelihood.

The only substantive complication is that, at the outset, the agents have no accumulation of records upon which to make their decisions. At time 0, tere are no acquaintances, and all agents are "strangers." To deal with this, we allow the agents to learn from experience. At the outset, the assumed harmony value for strangers is 0.5. Every time the agent interacts with a stranger, that value is adjusted to reflect experience. The upshot of all of this is that the agents are able to remember the others they have interacted with.

These selective agents can be placed into any of the interaction models, either the original Axelrod grid, the multi-agent grid, or the neighborhood/workplace models. To our surprise, in none of these implementations did we find that homogenity was preserved. In Table 5, we present two sets of summary statistics for the state of the model when the simulation ends. In 5a, the results for the standard ACM are shown after the identity-based selection of discussants is introduced. It is not particularly surprising that the model reaches a steady state in a smaller number of time steps after identity based selection is introduced. It is not a surprise that the agents believe that, when they interact, they are likely to find others who agree with them about most things. And, furthermore, it is not particularly surprising that agents believe that the probability that they have something in common with a stranger is around 0.5. It is, however, a surprise that the level of entropy is negligible and that the variance of opinion on each feature is 0.

[Table 5 about here]

The overall implication of these experiments is that the homogenization of public opinion is a pervasive (irresistible?) tendency when agents interact according to the logic of Axelrod's original culture model. Tinkering with the structure of the neighborhoods or attempting to condition interaction processes does not generate self-reinforcing social subgroups or diversity of opinion. As a result, we have not yet been able to rebut the nasty implication of this model, which is that in order to preserve diversity on a social level, one must eliminate diversity on an individual level by blocking interactions among people who differ.

IV. Opinion Adjustment Variations

Since we have investigated the selection process, and found no relief in several variants, we now turn our attention to reconsider the model of opinion adjustment.

The model we have been working with shares certain features with models of population genetics (Gillespie, 1998). An agent's culture array, e.g., (1,1,2,1,0), is scarcely different from a genetic string for an individual, and while interaction in the culture model is not exactly like genetic reproduction, there are certain similarities. In population genetics one of the central questions focuses on the question of whether each generation is more or less diverse than its predecessor. One theoretically compelling result is known as genetic drift. Genetic drift--the complete takeover of a particular part of the genetic code by one particular trait--happens over the long run if there is random interaction among individuals and reproduction follows the standard genetic laws. The homogenization predicted by the population geneticist corresponds closely to the pattern of homogenization we discussed in the culture model, although the logic of individual interaction and adjustment in the two models is not exactly the same.

If genetic drift (or cultural drift?) is to be expected when interactions occur among paired individuals, then perhaps a more significant conceptual departure is needed if diversity is to be preserved. Rather than tinkering with the selection process for individual interactions, we instead should attempt to reconceptualize outcome of the interaction. Perhaps the original culture model is too individualistic, in the sense that individual agents interact in an isolated dyad. What if, instead, social context plays a role?

To spare the reader some suspense, we note at the outset that this final avenue is promising in both respects, preserving diversity in the local network of people who interact as well as the aggregate society. We do not know of a theory exactly like this in the extant literature, but surely a similar one must exist. It might be called a "network-embedded resistance" model, for lack of a better term. Simply put, agents do not copy features from each other in a "willy nilly" fashion. Rather, they change only when there is "good reason" to do so. Good reason, is found when a majority of existing "friendly contacts" also support the proposed new opinion.

The essence of this approach is as follows. The agents come in contact with one another, and they keep records on identity and the opinions of the other agents. They are able to "make notes" to indicate if the other agent is diametrically opposed to them (disagrees on all issues) or is in complete agreement, or somewhere in between. An agent says another agent is agreeable, or is a "friend," if they agree on half or more of the features. When two agents interact, then a feature on which they differ is randomly selected and the agent who initiated the contact considers adopting the other's point of view. When the agent is presented with a contradictory opinion, it can conduct a "small poll" of the other agents it has met in the past to find out if the new opinion is good one. (We have two variants of this model. The agent can base its decision on its recollections of their friend's attitudes when they last met, or the agent can actively conduct a new survey of people who are listed as friends.) If more than one-half of the contacts who are recorded as agreeable support the new point of view, then it will be adopted. As far as we can tell, this new adjustment rule works to maintain diversity with any interaction model as long as the agent is exposed to a relatively large number of other agents .

[Table 6 about here]

In Table 6, we present summary statistics for this "context-embedded resistance" model in a couple of contexts. In 6a, the results for the impact of this change on the original Axelrod culture model. Recall that the original model has a 10x10 grid, discussants are selected at random from the four-sided neighborhood, and interaction occurs with probability equal to the similarity of the agents. The only difference in this model is that agents adopt new points of view when there is support for that view among a majority of their friends. As in the previous cases, we stop the simulation when ten passes are made through the list of agents without a single agent changing a single feature. In each of the 100 runs of the model, the level of entropy is in the middle ranges when the simulation stops. The variance of the features is also far from zero.

Furthermore, the perceptions of the agents indicate that they correctly perceive that they are in a mixed environment.

In Table 6b, the summary results for the model with 5 home grids and 3 work grids are presented. These summary results are not substantially different from the one neighborhood model represented by 6a. We have selected (at random) one of the runs of this model for presentation in Figure 5. Note how entropy starts at a relatively high level but settles down into a steady state in the middle range, while perceptions of homogeneity rise. Agents think they agree about two-thirds of the time with the people with whom they interact and they expect less than 1/3 of their discussants to be identical with them.

[Figure 5 about here]

We do not hold out this simulation exercise as a model of reality. Rather, as Axelrod (1997b) argues, it is a "thought experiment" through which we investigate the implications of our ideas. However, we hasten to add that we think we are in the right empirical "ball park." First, our survey-based studies of discussion partners generally unearth diversity within the networks of discussants. Furthermore, there is evidence that individuals who are presented with views that contradict their own are more likely to change their minds when there is a high level of support for the new view within their network of friends (see Huckfeldt, Johnson, and Sprague, forthcoming).

V. Conclusion

The tendency of repeated social interaction to squeeze out cultural diversity appears to be persuasive when opinion adjustment occurs according to the model originally specified by Axelrod. A number of efforts to adjust the logic through which agents are "put together" did not change that basic tendency. If we dub this phenomenon "cultural drift," it indeed appears pervasive in models of dyadic interaction.

A change in the way that agents decide whether or not they should adopt new culture features can change the model's fundamental tendencies, however. If agents view new information within the context of their political experience, taking into account information from other discussion partners, then the tendency toward homogenization is abated. While the interaction process does result in homogenization to a degree, it appears that, across a relatively wide range of initial conditions, and within a number of interaction environments, the use of the network--embedded adjustment will preserve a measurable amount of diversity.

On the face of it, there does not seem to be any necessary reason why the network--embedded model of decision should preserve diversity. After all, if a society were completely homogeneous, then interaction would sustain that pattern, and sprinkling in minority points of view here and there would not change that. Furthermore, we acknowledge that it is possible to design starting positions for the simulation that would cause diversity to be eradicated. A carefully crafted chain of networks could be constructed, no doubt, that would fold under the pressure of repeated interaction. However, when preferences are assigned in a random way, we have found that network-embedded individuals act in such a way as to preserve diversity both within their immediate spheres as well as on an aggregate level.

The ironic aspect of this outcome is, of course, that the agents themselves are not seeking diversity in any sense of the word. Quite the opposite, they are seeking to cooperate with the majority of the agents that are most like them. As a consequence, we are inclined to call the resulting diversity an emergent property of the complex system.

Finally, in conclusion, we would like to mention that we have considered many minor adjustments in this model that we have not reported here. We considered a model in which agents were spread unevenly on the grid. We considered, for example, a change in the interaction process so that either agent might copy a feature from the other, or that both might do so. We have considered variations on the way agents decide how appealing strangers might be. None of these changes yielded results which differ from those described here.

Upon publication, we will eagerly make available the source code for this model and we will promise to keep it up-to-date with changes in the Swarm toolkit itself.

References

Axelrod, Robert, 1997a . "The Dissemination of Culture: A Model with Local Convergence and Global Polarization," *Journal of Conflict Resolution*, vol. 41, pp. 203-26.

Axelrod, Robert. 1997b. The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton, NJ: Princeton University Press.

Axtell, Robert, Robert Axelrod, Joshua M. Epstein, and Michael D. Cohen. 1996. Aligning Simulation Models: A Case Study and Results. *Computational and Mathematical Organization Theory* 1: 123-41.

Balch, Tucker. 2000. Hierarchical Social Entropy: An Information Theoretic Measure of Robot Team Diversity. *A*utonomous Robots 8(3):.

Gillespie, John H. 1998. *Population Genetics*. Baltimore, MD: Johns Hopkins University Press.

Huckfeldt, Robert, Paul E. Johnson, John Sprague. Forthcoming. Political Environments, Political Dynamics, and the Survival of Disagreement. *Journal of Politics*.

Huckfeldt, Robert, and John Sprague. 1995. Citizens, Politics, and Social Communication. New York: Cambridge University Press.

Shannon. C.E. 1949. The Mathematical Theory of Communication. Champaign-Urbana, University of Illinois Press.

Table 1

Axelrod Culture Model		
5 Features, 3 Traits		
Summary of 100 Simulations		
	Mean	Standard Deviation
Iterations	441.89	146.49
Variance of Opinion		
feature 0	0	0
feature 1	0	0
feature 2	0	0
feature 3	0	0
feature 4	0	0
totalEntropy	0	0
Average of Perceptions		
Acquaintance	0.99	0
Harmony	0.99	0
Identical	0.97	0.01

Table 2
Multi-Agent Grid

5 Agents per cell (10x10 grid)		
5 Features, 3 Traits, Parochialism=0.95		
Summary of 100 Simulations		
	Mean	Standard Deviation
Iterations	8220.61	2511.32
Variance of Opinion		
feature 0	0.00168	0.0011
feature 1	0.00038	0.0038
feature 2	0	0
feature 3	0.00008	0.0008
feature 4	0	0
totalEntropy	0.000408	0.002008
Average of Perceptions		
Acquaintance	0.999	0.00042
Harmony	0.999	0.00029
Identical	0.997	0.00144

Table 3
Restricted Self-Selection: The "Coleman" Model

Coleman Model

5 Features, 3 Traits
Summary of 100 Simulations

Coleman Parameter	0.2		0.5		0.8	
	Mean	St. Dev.	Mean	Std.Dev	Mean	Std.Dev.
Iterations	240.85	79.1	262.92	77.74	272.26	94.58
Variance of Opinion						
feature 0	0	0	0	0	0	0
feature 1	0	0	0	0	0	0
feature 2	0	0	0	0	0	0
feature 3	0	0	0	0	0	0
feature 4	0	0	0	0	0	0
totalEntropy	0	0	0	0	0	0
Average of Perceptions						
Acquaintance	0.99	0	0.99	0	0.99	0
Harmony	0.99	0	0.99	0	0.99	0
Identical	0.96	0.02	0.96	0.01	0.96	0.02

Table 4
Five Neighborhoods (10x10) and Three Workplaces (5x5)

5 Features, 3 Traits, Parochialism=0.5
Summary of 100 Simulations

	Mean	Standard Deviation
Iterations	20286	29.01
Variance of Opinion		
feature 0	0	0
feature 1	0	0
feature 2	0	0
feature 3	0	0
feature 4	0	0
totalEntropy	0	0
Average of Perceptions		
Acquaintance	0.999	0.002
Harmony	0.998	0.00155
Identical	0.988	0.0077

Table 5
Identity-based Selection of Discussants

a) Identity-based selection in the standard Axelrod model (10x10 grid)

5 Features, 3 Traits
Summary of 100 Simulations

	Mean	Standard Deviation
Iterations	282.03	90.88
Variance of Opinion		
feature 0	0	0
feature 1	0	0
feature 2	0	0
feature 3	0	0
feature 4	0	0
totalEntropy	0	0
Average of Perceptions		
Acquaintance	0.599.	0.010
Harmony	0.99	0.0034
Identical	0.958	0.0166

b) Identity-based selection in the 5 neighborhood, 3 workplace model

5 Features, 3 Traits
Summary of 100 Simulations

	Mean	Standard Deviation
Iterations	10220.6	3588.51
Variance of Opinion		
feature 0	0	0
feature 1	0	0
feature 2	0	0
feature 3	0	0
feature 4	0	0
totalEntropy	0	0
Average of Perceptions		
Acquaintance	0.599.	0.005
Harmony	0.999	0.00181
Identical	0.983	0.009

Table 6
Network-embedded Resistance

a) Resistance incorporated in the standard Axelrod culture model (10x10 grid)

5 Features, 3 Traits
Summary of 100 Simulations

	Mean	Standard Deviation
Iterations	71.5	13.82
Variance of Opinion		
feature 0	0.64	0.11
feature 1	0.64	0.056
feature 2	0.64	0.061
feature 3	0.64	0.062
feature 4	0.63	0.062
totalEntropy	0.71	0.016
Average of Perceptions		
Acquaintance	0.44	0.031
Harmony	0.60	0.030
Identical	0.36	0.046

b) Resistance incorporated in the 5 neighborhood, 3 workplace model

5 Features, 3 Traits
Summary of 100 Simulations

	Mean	Standard Deviation
Iterations	7871.3	29.01
Variance of Opinion		
feature 0	0.61	0.12
feature 1	0.64	0.13
feature 2	0.64	0.12
feature 3	0.61	0.11
feature 4	0.61	0.12
totalEntropy	0.65	0.054
Average of Perceptions		
Acquaintance	0.49	0.025
Harmony	0.68	0.023
Identical	0.37	0.043

Figure 1
Axelrod Culture Model

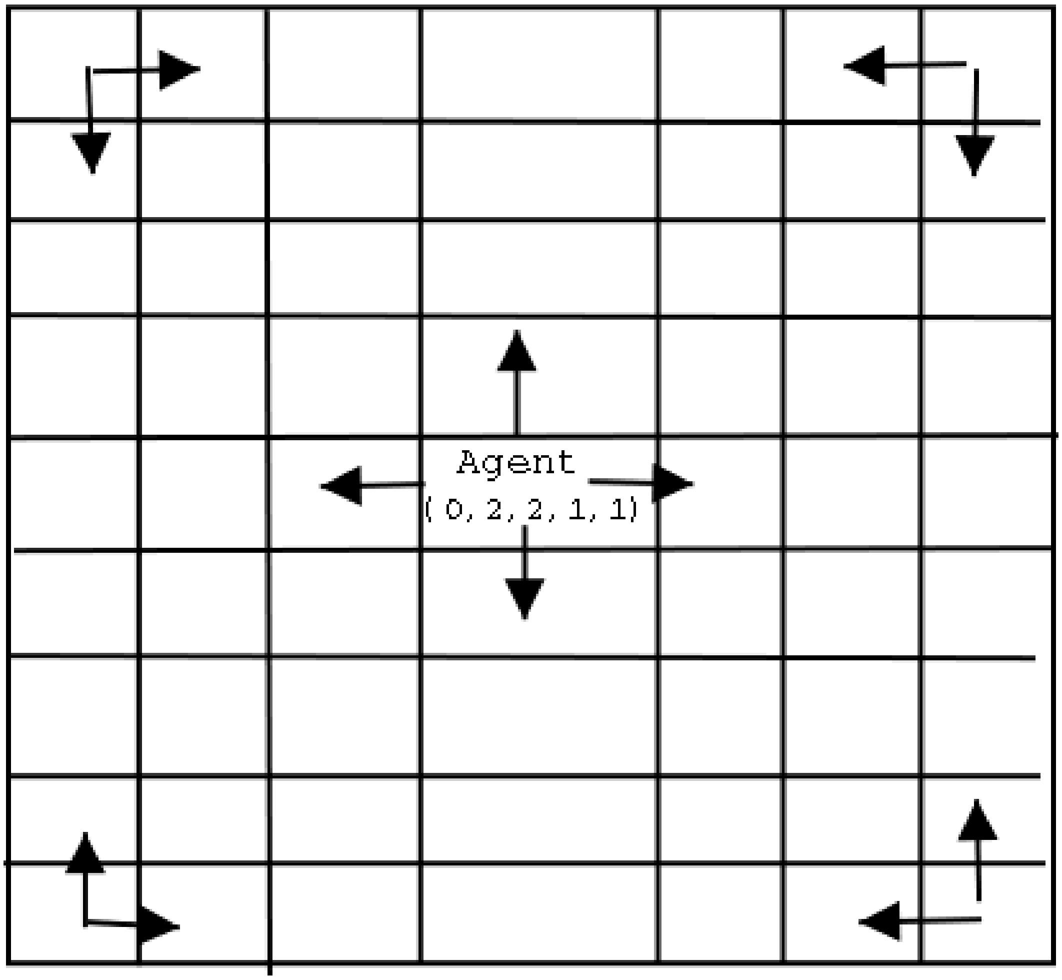

Basic model

- Grid of agents (villages)
- Feature array for each agent, e.g., (0,2,2,1,1) for a 5 feature model
- Interaction within truncated vonNeumann neighborhoods

Figure 2
Axelrod Culture Model

2a) Initial conditions

Feature: 0 1 2 3 4

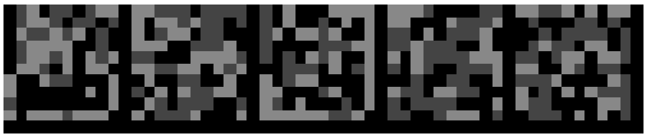

2b) Final conditions

2c) Time paths of summary measures

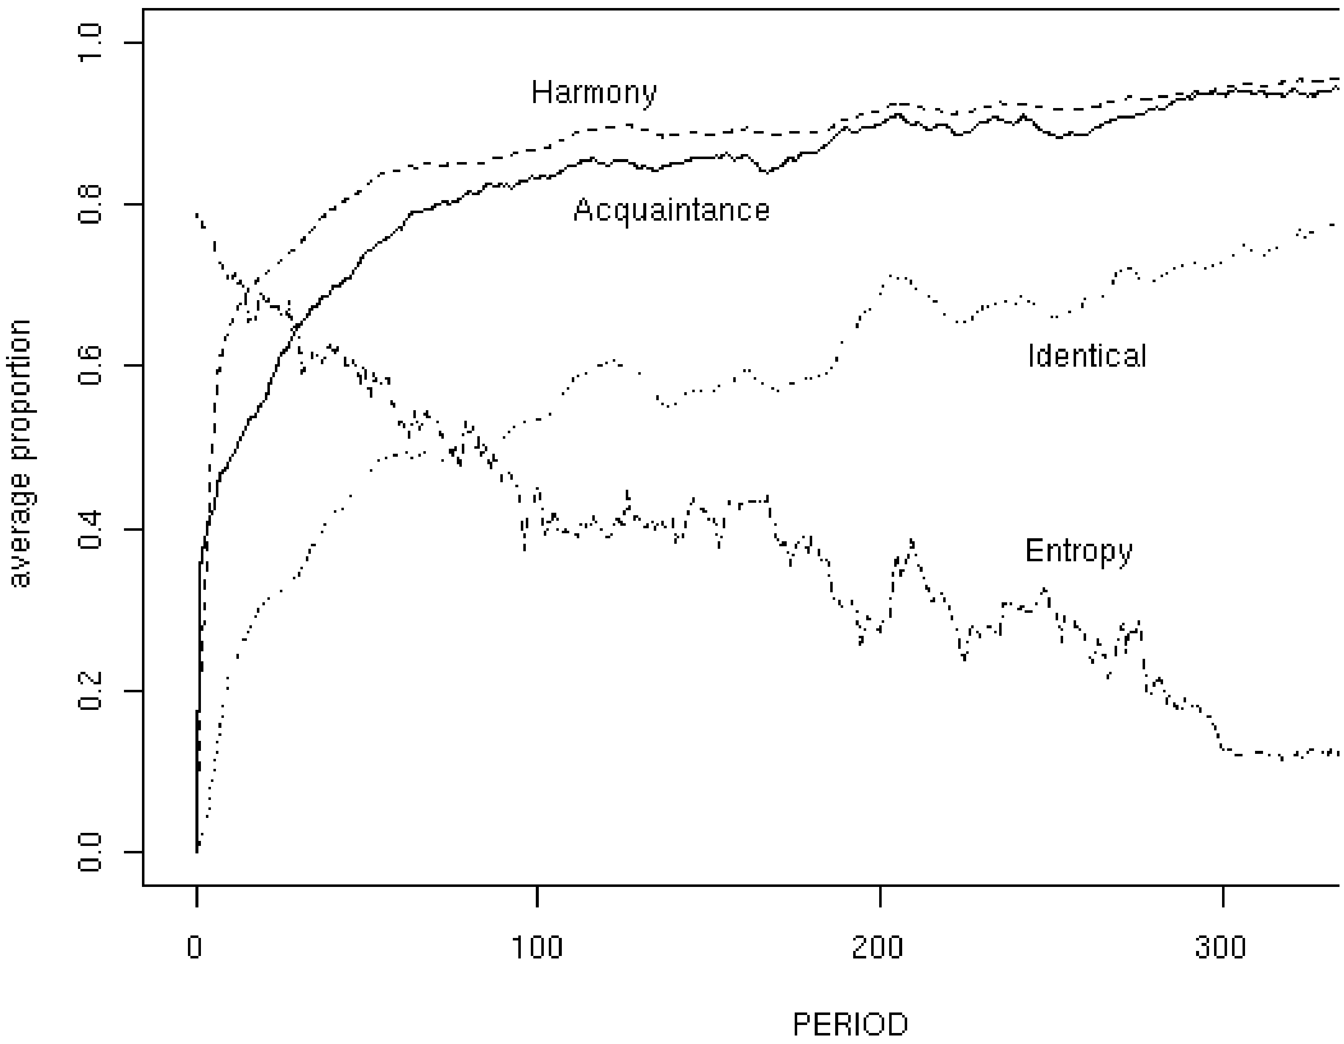

Figure 3
Parochialism

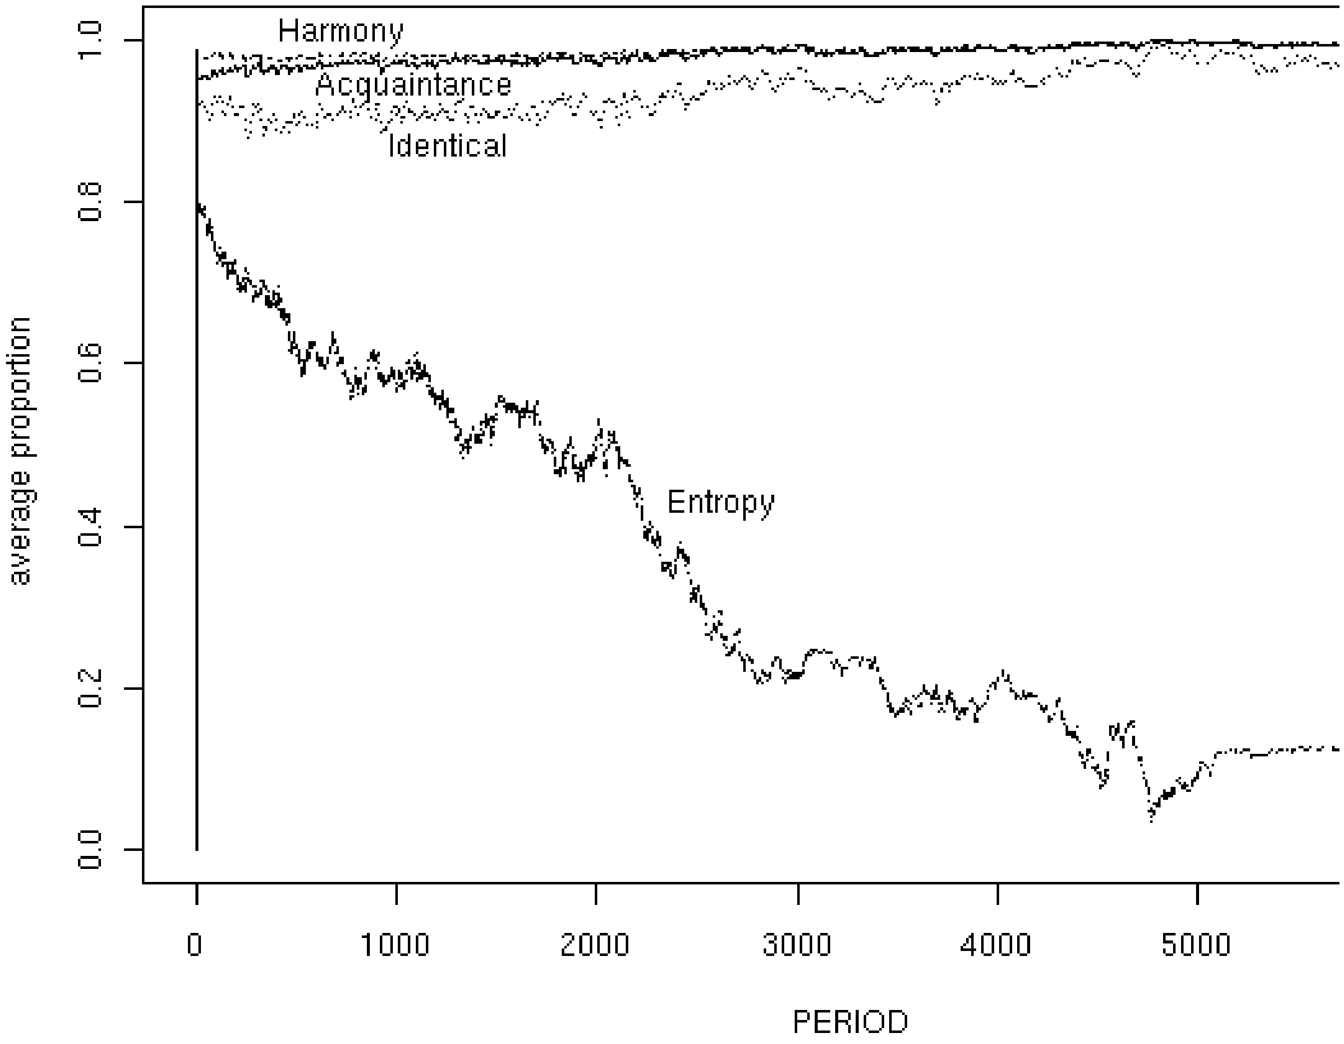

Figure 4
Multiple Home Grids and Work Grids

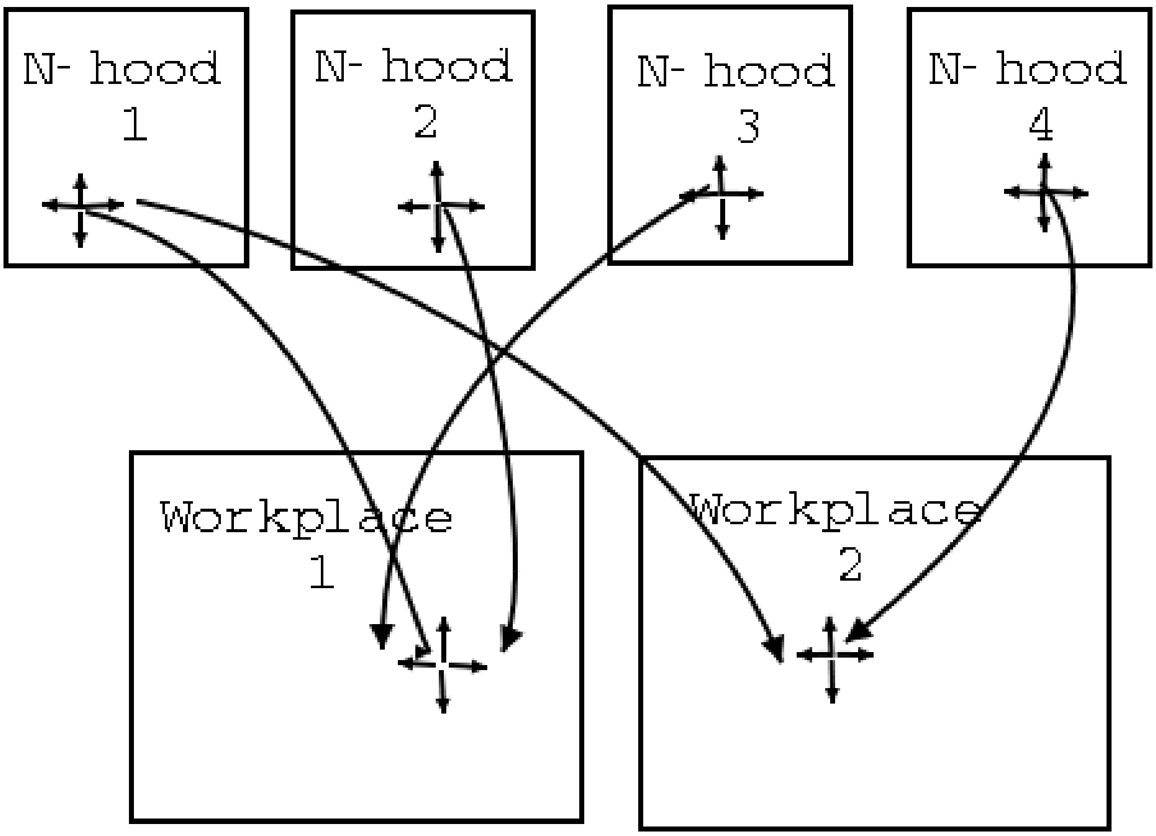

Figure 5
Network-embedded Resistance to Homogenization of Culture