AUTOREGRESSIVE INFLUENCE AND THE DURABILITY OF POLITICAL DISAGREEMENT

Robert Huckfeldt
Indiana University

Paul E. Johnson
University of Kansas

Abstract

This paper combines results from survey research and simulation to address questions regarding the nature and impact of political communication among citizens. While influential and long-standing theories predict that political disagreement will be a rare event among people who communicate on a frequent basis, the accumulating evidence points to a different conclusion. Indeed, political disagreement is remarkably durable even within closely held networks of political communication, and while communication among interdependent citizens is frequently influential, it does not eliminate diversity. Following Axelrod's suggestion that an agent-based model can be a useful tool for "thought experiments" and the clarification of theory, we have used the Swarm simulation toolkit to investigate the formation of discussion networks and implications of theories about persuasion and information exchange.

Prepared for delivery at the Annual Meeting of the Midwest Political Science Association, April 29, 2002, Palmer House Hotel, Chicago.

AUTOREGRESSIVE INFLUENCE AND THE DURABILITY OF POLITICAL DISAGREEMENT

A central issue in the study of democratic politics is the capacity of citizens and electorates for tolerating political disagreement. The model of a free, open, and democratic society is one in which political issues are fully explored and political debates are fully aired. In such a society, citizens are open to persuasion, the social boundaries on political viewpoints are fluid and shifting, and individuals encounter the full spectrum of issue positions and political viewpoints.

How does this model correspond to contemporary analyses of citizens, communication, and disagreement in democratic politics? According to one analytic perspective, citizens employ socially supplied information as a labor saving device. By finding well informed individuals with political biases similar to their own, citizens are able to reduce information costs by relying on experts (Downs 1957). Hence, the likelihood of disagreement is reduced because individuals rely on the guidance of politically compatible experts.

Other analysts, inspired by a conformity model of social influence (Asch 1956), see a powerful social influence process within small, cohesive groups of interdependent citizens. The psychic discomfort of disagreement causes individuals to reduce dissonance through various means (Festinger 1957). In particular, individuals adopt socially prevalent viewpoints, and they avoid disagreement in the first place by censuring their patterns of social interaction to create politically homogeneous networks of political communication.

Neither of these analyses is able to accommodate the survival of disagreement in patterns of meaningful communication and deliberation among citizens. In the model of

communication as a labor saving device, disagreement is unlikely to occur due to the purposeful action of individual citizens as they seek out like-minded experts. In the conformity model, disagreement is extinguished through powerful mechanisms of social influence. Hence, the capacity of democratic electorates to consider and reconcile competing viewpoints through a meaningful process of political communication is rendered problematic.

The strategy of this paper is twofold. First, we evaluate empirical evidence regarding the survival of political disagreement among citizens within their naturally occurring patterns of social interaction. Second, we evaluate a dynamic, agent-based model of political persuasion to assess the mechanisms that might sustain disagreement among citizens. In both instances we assume that communication among citizens is politically consequential – that citizens are politically interdependent in the sense that they rely on one another for political information, expertise, and guidance. But we question both the extent to which they censure their patterns of political communication to eliminate disagreement, as well as the extent to which they are driven to conformity by the experience of political disagreement. The primary question becomes, what are the conditions under which diversity of opinion is likely to be sustained?

Conformity and the Survival of Political Disagreement

In their pioneering studies of social influence among citizens during election campaigns, the Columbia sociologists articulated a powerful model of political communication and change during election campaigns (Lazarsfeld et al. 1944; Berelson et al. 1954). They argued that political communication among citizens becomes less frequent during the period between election campaigns, and hence political preferences tend to become individually idiosyncratic. As the frequency of political communication increases in response

to the stimulus of a new election campaign, these idiosyncratic preferences become socially visible, and individuals are correspondingly brought into conformity with micro-environmental surroundings.

In the context of a high visibility election campaign, the dynamic logic of group conformity pressures is quite compelling. Before the campaign begins, people are less concerned about political affairs, and hence their conversations focus on other, nonpolitical topics: baseball, gardening, etc. As long as their political preferences are socially invisible, they are immune to conformity pressures, and hence preferences become individually idiosyncratic. As the campaign accelerates, so does the rate of political communication among associates, and individual political preferences are increasingly exposed to social scrutiny. The stage is thus set to bring individual preferences into line with the preferences that are dominant within networks of social relations.

Carried to its extreme, the logic of group conformity would seem to suggest that political disagreement should disappear within networks of social relations.[1] Pressures toward conformity might drive out disagreement in several ways (Festinger 1957; Huckfeldt and Sprague 1995). First, the discomfort of disagreement might encourage people to modify their patterns of social relations so as to exclude people with whom they disagree, or in a less extreme form, to avoid political discussion with associates who hold objectionable preferences. Second, and perhaps most importantly, individuals might bring their own preferences into correspondence with the preferences that they encounter within their networks of social relations – they may be influenced by the preferences of others.

[1] Indeed, while the production of political homogeneity within networks of political communication is a lesson that is frequently attributed to the Columbia studies, their efforts paid ample attention to the importance of disagreement and the natural limits on conformity pressures within campaign dynamics (Berelson 1954: chap. 7).

Contrary Evidence Regarding Disagreement

As compelling as the group conformity argument may be, it suffers from at least one major shortcoming – disagreement is not typically extinguished within networks of social relations, even at the end of high stimulus presidential election campaigns. The tendency for disagreement to survive within communication networks has been documented by a series of national and community-based election studies (Huckfeldt and Sprague 1995; Huckfeldt et al. 1995; Huckfeldt, Sprague, and Levine 2000). In each instance, interviews with discussants identified by the original main respondents showed less than perfect correspondence within discussion dyads. Indeed, no more than two-thirds of the discussants held a presidential candidate preference that coincided with the main respondent who named them.

These levels of disagreement become even more important when we recall that they are based on dyads rather than networks. If the probability of dyadic disagreement within a network is .7, and if the likelihood of disagreement is independent across the dyads within a network, then the probability of agreement across all the relationships within a three-discussant network drops to $.7^3$ or .34. In this scenario, ***disagreement and heterogeneous preferences become the rule rather than the exception within the micro-environments surrounding individual citizens***.

The pervasiveness of disagreement within communication networks forces a reassessment of social conformity as a mechanism of social influence, as well as a reconsideration of the dynamic implications that arise due to politically interdependent citizens. In the analysis that follows, we assess the conditions that give rise to socially sustained disagreement. Our argument is, quite simply, that complex patterns of

communication among citizens might sustain as well as extinguish patterns of disagreement among citizens.

Network Heterogeneity in the 2000 Election

Each respondent to the post-election survey of the 2000 National Election Study was asked to provide the first names of the people with whom they discussed government, elections, and politics. In a subsequent battery of questions, they were asked to make a judgment regarding the presidential candidates for whom each of these discussants voted. Seventy-four percent of the post-election respondents were able to provide at least one name, and Table 1 is based on these respondents.

In Parts A and B of Table 1, the distribution of candidate preferences within the networks is cross-classified by the respondent's reported vote choice. Both parts of the table show clear evidence of political clustering among the respondents. Nearly 63 percent of the Gore voters fail to report a single discussant who supports Bush, and 64 percent of the Bush voters fail to report a single discussant who supports Gore. This means, of course, that more than one-third of the two-party voters can identify at least one discussant who supports the candidate of the opposite party. Forty percent of Gore voters and 46 percent of Bush voters report that all their discussants share the same vote preference, and this means that more than half of these respondents perceive that they reside in politically heterogeneous networks.

In presenting a tale of glasses that are **both** half empty **and** half full, Table 1 reinforces the findings of earlier studies. While the vote choices of respondents are clearly interdependent with the distributions of vote choices within their networks, these patterns of interdependence are far from complete. Indeed the evidence of disagreement is as noteworthy

as the evidence of agreement.[2]

Autoregressive Patterns of Agreement and Disagreement

How do individuals respond to political heterogeneity within their communication networks? Other analyses suggest that the views of individuals are discounted if they run counter to the dominant view within the network. In this way, the message conveyed by any particular discussant is autoregressively weighted by the distribution of preferences in the remainder of the network (Huckfeldt, Johnson, and Sprague 2002). Hence, we would expect agreement within dyads to depend not only on the political distance between the discussant's candidate preference and the respondent's political orientation, but also on the distance between the discussant's preference and the distribution of preferences in the remainder of the network.

Using the logit models in Table 2, the presence or absence of agreement within each of the identified dyads is considered in terms this argument. In each of the dyads analyzed in the first model, the respondent perceives that the discussant voted for Bush, and hence agreement is defined as an instance in which the respondent also reported voting for Bush. Similarly in the second model, the respondent perceives that each of the discussants supported Gore, and agreement exists if the respondent also reports voting Gore. Finally, in the third model, respondents perceive these discussants voted neither for Bush nor Gore, and agreement exists

[2] Table 1 inevitably includes modest biases that underestimate levels of diversity. While the 2000 National Election Study did not include interviews with the discussants of the main respondents, other studies have incorporated snow ball designs that pursue interviews with members of the networks named by the respondents (Huckfeldt and Sprague 1987; Huckfeldt et al. 1998; Huckfeldt, Sprague, and Levine 2000). These studies use the discussant's self-reported vote as the criterion against which to judge the presence of systematic biases in the respondent's perception of political preferences in the network. While respondents are reasonably accurate in their perceptions of discussant preferences, respondents are less likely to recognize the preferences of discussants with whom they disagree.

if the respondent reports voting for neither candidate. (Based on these measurement procedures, the participants in 60 percent of all dyads in the sample hold the same political preferences.)

Several explanatory variables are included in the models: the respondent's partisanship, as well as the numbers of discussants in the remaining network whom the respondent perceived as voting for Bush, Gore, and neither. In the models for discussants who are perceived to support Gore and Bush, respondent partisanship is measured on the traditional 7-point scale, where -3 is strong Democrat and 3 is strong Republican. In the model for discussants who support neither candidate, respondent partisanship is measured as strength absent direction – as the absolute value of the party identification measure. Recall that a maximum of four discussants is recorded for each respondent, and hence the maximum number of discussants in the residual network is three.

These models allow us to address the conditions that enhance and diminish the probability of agreement within particular dyads. The models consistently produce statistically discernible coefficients for respondent partisanship. Strong Demcrats are more likely to agree with Bush voters and less likely to agree with Gore voters. Correspondingly, strong Republicans are more likely to agree with Bush voters and less likely to agree with Gore voters. Strong partisans of either variety are less likely to agree with voters who support neither of the major party candidates.

The preference distributions within the residual networks produce more complex (and perhaps more interesting) results. Support for Bush in the residual network enhances the probability of agreeing with a discussant who supports Bush and attenuates the probability of agreeing with a discussant who supports neither of the major party candidates. Similarly,

support for Gore enhances the probability of agreement with a discussant who supports Gore and diminishes the probability of agreement with a discussant who supports neither candidate. In addition, increased Gore support attenuates the probability of agreeing with a discussant who voted for Bush. Finally, the number of discussants who support neither candidate fails to produce a discernible coefficient in any of the three models.

The magnitudes and implications of these patterns of relationships are seen most readily in Figure 1, which displays the estimated probabilities of agreement within dyads for respondents who identify as political independents. The independent respondents are more likely to agree with discussants who voted for Bush to the extent that support for Bush occurs in the remainder of their networks. Similarly, they are more likely to agree with Gore-voting discussants to the extent that Gore voting is more common in the remainder of their networks. In both instances, the agreement probability is only modestly attenuated by the distribution of other preferences or non-preferences. Finally, and in contrast, the probability of agreement with a discussant who supports neither candidate is not enhanced by the presence of other non-supporters in the remainder of the network, but is diminished by the presence of either Bush or Gore voters.

Figures 2 and 3 replicate Figure 1, but for respondents who identify as strong Republicans and strong Democrats respectively. These figures demonstrate the strong bias toward (1) agreement with discussants whose vote preferences correspond with the respondents' partisan loyalties and (2) disagreement with discussants whose vote preferences run counter to the respondents' partisan loyalties. The pattern of agreement and disagreement that depends on preference distributions in the residual network is also present in these figures, but its magnitude is greatly reduced.

In summary, this analysis yields several results.

1. A strong partisan is highly likely to **agree** with a discussant who supports her own party's candidate – and to **disagree** with a discussant who supports the opposite party's candidate – **regardless** of the partisan division in the remainder of the network.

2. As the strength of partisanship decreases, the probability that individuals will agree or disagree with a discussant who supports either candidate increasingly becomes contingent on the distribution of candidate preferences in the remainder of the communication network. In particular, the probability of agreement is enhanced by the presence of other discussants who hold the same political preference.

3. The probability of agreeing or disagreeing with a discussant who supports **neither** of the major party candidates is also contingent on preference distributions in the remainder of the network. While the probability is not enhanced by the presence of other discussants who support neither candidate, it is dramatically diminished by the presence of discussants who support either of the candidates.

What do these results suggest? Agreement within dyads is typically sustained by larger networks of communication that simultaneously support the preferences of <u>both</u> individuals within the dyad. Hence, disagreement is also socially sustained, but by politically <u>divergent</u> networks that serve to pull the two members of the dyad in politically <u>opposite</u> directions (Huckfeldt, Johnson, and Sprague 2002). In summary, the survival of disagreement within dyads is profitably seen within larger patterns of association and communication that occur at the intersection between the networks that surround individual citizens.

How are we to evaluate the implications of these conclusions? We certainly do not possess the data that would be needed for a full evaluation of political homogeneity and diversity among and between the networks of communication within which citizens are imbedded. Indeed, such a body of information is, as a practical matter, quite nearly inconceivable. And hence our efforts point in a different analytic direction.

Agent Based Models of Social Influence

In the remainder of this paper, we pursue an agent-based modeling strategy. Two of the leading simulation models that explore these questions are Axelrod's (1997) Culture Model (ACM) and the social impact model (SIM), which is based on a theory developed by Bibb Latané (Nowak, Szamrej, and Latan é, 1990; Latané, Nowak, and Liu, 1994; Nowak and Lewenstein, 1996). In these models, the opinion-holders are thought of as cells arranged evenly on a square grid and opinions are adjusted in response to the status of neighboring cells. Axelrod's model is based on dyadic interactions, while the SIM updates cells according to a sum of influences exerted by all cells (simultaneously).

The simulation model that we present begins as an extension of Axelrod's culture model, a model which describes evolution of culture due to small-scale (localized) social interaction. In Axelrod's approach, the agents, or cells in a grid, are described as villages, each village has a culture, represented by a vector of features. These features are integer-valued, e.g., 0, 1, 2. The values of the features are referred to as traits, and are randomly assigned at the outset. In the ACM, a village is randomly selected and can look "up", "down", "left" or "right" to find another village (that is, it is a truncated von Neumann neighborhood; in contrast, see Epstein and Axtell, 1996). After a random neighbor is selected, an interaction occurs with probability equal to the similarity in the traits of the two agents. If the interaction occurs, then an issue on which the two disagree is selected at random and the agent's opinion on the issue is changed to match that of the selected neighbor. Hence, influence automatically follows whenever interaction occurs.

Axelrod made a number of observations on the basis of his model, the most striking being that, over the long run, cultural diversity tends to be eliminated. While the tendency

toward homogeneity is greater for some parameter settings than others, it is powerful in all cases. When diversity survives in the Axelrod model, it is a diversity of the most extreme sort. Different cultural clumps are completely homogeneous and totally isolated from one another. If a village interacts, it interacts with villages that are identical to it. As Axelrod shows, separate groups do not form in some conditions, but they are more likely to form if the number of traits per feature is high. Under those conditions, two agents are less like to have anything in common and so they never interact. He shows that the number of clusters decreases as the number of features increases, and the number of clusters increases as the number of traits increases.

Axelrod's conclusion poses a challenge for the current effort. A useful model of political communication within small networks of citizens should accommodate homogeneity as well as heterogeneity across interactions because that is what repeated empirical analyses have demonstrated. One solution is to create agents who are individually resistant to environmental influences, similar to the strong partisans in Figures 2 and 3. That is not the route explored here, for two reasons. First, as Figure 1 suggests, not all individuals ***are*** resistant to influence. Second, creating influence-resistant agents constitutes a ***deus ex machina*** that avoids the more theoretically demanding task – the development of a more intricate understanding of the formation of networks and the formulation of public opinion.[3]

We have constructed a general purpose agent-based model of social interaction and influence. The model is implemented in Objective-C using the Swarm Simulation Toolkit (Minar, et. Al, 1996; our current version uses features offered in Swarm-2.2), which is

[3] We do not mean to suggest that the presence of influence-resistant agents do not have emergent consequences for patterns of agreement and disagreement. And it is a problem that we are currently addressing in other work.

currently supported by the Swarm Development Group, a nonprofit membership organization (http://www.swarm.org). The original Axelrod model can be produced as a special case. In our model, agents are able to move about among multiple grids that have various properties, exercise varying levels of self-selection in the formation of interaction networks, and, perhaps most importantly, they can react to each other in a number of different ways (Huckfeldt, Johnson, Sprague and Craw 2000; Johnson and Huckfeldt, 2001). The model has a range of variables for analysis, including such basic issues as the size of the grid, the number of features and traits, the scheduling of agent actions, and so forth. Substantively important additions concern the processes through which others are sought out for discussion and opinions are adjusted. In addition to introducing a number of system and individual level parameters, we also have introduced summary measures for the diversity of opinion (entropy) as well as measures for the individual experience of diversity. These are discussed below (see also, Johnson 1999 and Johnson and Huckfeldt, 2001).

In terms of the formal implementation, the most significant departure is the incorporation of code that allows us to have several agents inhabit the same cell at the same time. This "multi-agent grid" structure allows a considerable amount of freedom in model design (in contrast, the standard approach restricts agent movement by requiring agents to search for open cells). Each agent has a "home grid" and a "work grid". Agent movement is controlled by a variant of "dynamic scheduling" (see Johnson and Lancaster, 2000: Chapter 9.6). Each agent plans its activities over the course of a “day”, which is a predetermined (in this case, 10) number of time steps. At the beginning of each day, the agents are randomly sorted and each is told to schedule its movements throughout the day and to select (at random)

a time during the day at which to initiate an interaction. We can reproduce the Axelrod model by creating only one home grid and forcing agents to spend all of their time there.

In this paper we are primarily concerned with the dynamic consequences of dyads that are autoregressively dependent on preference distributions in the larger network. The empirical evidence—in this paper and elsewhere (Huckfeldt, Johnson, and Sprague 2002)—indicates that influence within dyads depends on preference distributions in the larger network. In this spirit, our model departs from the Axelrod formulation of making influence automatic once interaction occurs. In the current effort, we implement the previously demonstrated empirical result that dyadic influence is in fact contingent—autogressively contingent—on the presence of various opinions and preferences in the remainder of the network.

The autoregressive influence model differs from the ACM in that each agent takes into account a broader array of opinions, but it does not go so far as the social impact model (SIM). In the SIM, all cells are "simultaneously" updated against a "snapshot" of the whole society. Each cell in the grid is "acted upon" by every other cell according to a distance-based law of influence (Nowak, Szamrej, and Latane, 1990; Latane, Nowak, and Liu, 1994; Nowak and Lewenstein, 1996). As such, the SIM invokes a sort of "social telepathy" (Erbring and Young, 1979) that we seek to avoid. In our model, agents know about the state of the world only through direct two-way interactions, and they build up their networks through experimental adjustment. Unlike the SIM, which predicts a world in which certain opinions are held in tightly clustered subgroups of homogeneous cells, we strive to understand ways in which diversity can be preserved. While it is possible to design a replication of the SIM as a special case of our general framework, such a model is quite slow and inefficient. Because of

that fact, we have implemented a separate, generalized version of the SIM. While we do not explore that model in detail here, the code is available from the authors upon request.

A Baseline Model

The baseline model is intended to show that our generalized model can generate the same results as the original ACM. In the full computer model, each agent in the model is conceived of as a separate "citizen" object with the ability to move about, initiate interactions, and adjust opinions. The baseline model is a restricted version, since the agents are distributed evenly over a 10 by 10 grid and they are fixed in positions. In the models described in this paper, we have set the number of features at five and the number of traits at three.

In the baseline model, agents are not allowed to move—an agent looks for a discussion candidate in the way that Axelrod described, choosing a discussion candidate at random from the neighborhood with an interaction probability equal to the similarity of the two agents.[4] As in the ACM, the result of interaction is that the agent copies one feature on which the agent and the discussant differ.

As the simulation proceeds, the agents keep records about the others they have encountered. They note, first, what fraction of the discussion candidates they encounter agree with them about a randomly chosen feature (when they find such a common feature, we call them "acquaintances" because an interaction will follow). Among the people selected for interaction, the agent makes note of the proportion of features on which it agrees with the discussant (the degree of "harmony"), and it also notes if the agent's features are identical to

[4] If one sets the day to length 1, and selects only one agent for an interaction per day, then this model is identical to the original Axelrod model.

its own. Each agent uses a 20 period moving average to tally these observations. We can aggregate these individual experiences by calculating various summary statistics.

The baseline model reproduces dynamics of the original ACM. A graph depicting three summary measures calculated from one run of the model is presented in Figure 4. The most obvious feature of Figure 4 is that all three measures converge to unity. First, the "acquainted" line indicates the average proportion of random encounters that produce interaction. A higher acquaintance rate reflects a higher level of shared preferences among individuals within a neighborhood. As time goes by, more and more neighbors find themselves open to interaction with a randomly chosen neighbor. Second, the "harmonious" line indicates the level of agreement between people who interact. It reflects experience ***within networks of interaction***, indicating that the chances of disagreeing about any particular issue are diminished over time. Finally, the "identical" line indicates the average of individual experience regarding the extent to which agents are identical to the discussants with whom they interact. Here again, the focus is on those agents engaged in interaction, and particularly the proportion of interacting agents who hold identical positions on all five issues. Not only are people open to more interaction, but those interactions are increasingly likely to result in total agreement between the agents.

This particular run is not significantly different from the others we conducted with these parameter values. We set the model so that it would terminate the simulation if no opinion change was observed for 10 consecutive days, or 100 time steps. The average number of steps to termination is 8972.9, and in each of the 100 runs, all diversity was eliminated. Entropy, an index of diversity across the population of opinion, drops to 0 in all cases.

Quite clearly, this model leads to the same outcome as the Axelrod model. Equally clearly, the model does not correspond to our empirical observations. Uniformity and a lack of disagreement are not standard features of the political landscape, and our objective has been to consider several changes in the specification of the model that would yield a more believable world. Variations in levels of self-selection, as well as the assignment of agents to multiple grids, have not altered the model's outcome – homogeneity continues to be the stable equilibrium outcome of the process. In the analysis that follows, we consider the consequences of autoregressive patterns of influence within and among discussion dyads.

Separating Persuasion from Interaction

The Axelrod model conflates interaction with persuasion. Whenever an agent interacts with a discussant who differs, one feature is automatically copied from the discussant to the agent. In this way, agents are wholly indiscriminate in their adoption of opposing points of view. For many purposes, this is perhaps a wholly adequate model. If you need information regarding web sites for vacation alternatives, you might indeed seek out information from people with travel interests similar to your own and take whatever information they provide.

In contrast, the value of political information taken through social interaction is problematic. Even if you acquire information from a generally trustworthy individual suggesting that George W. Bush is just another rich fraternity kid who would make a terrible president, you might want to evaluate the worth of that information. The important point is that communicated information does not necessarily translate into influence, and in this sense the influence of even effectively communicated information is quite problematic.

How do people evaluate the worth and credibility of political information? What makes for political information on the part of a communicated opinion or preference? Indeed, a range of factors could be considered: the clarity with which individuals communicate, the imputed expertise of political discussants, and more. In this analysis we focus on the incidence of opinions within networks of political communication (Huckfeldt, Johnson, and Sprague, 2002).

If you think that George W. Bush is high quality presidential material, and one of your friends tells you that George Bush is just another rich fraternity kid, how might you respond? According to the baseline model you would simply change your opinion, but an alternative strategic response is to contextualize the information provided by the discussant relative to information provided by other discussants. Hence if you like Bush, but your friend Joe dislikes him, you might take account of other opinions about his capabilities. If all your other information sources suggest that he is a good guy, you might downgrade the credibility that you place on Joe's opinion. On the other hand, if all your other information sources agree with Joe, you might reconsider your own opinion on the matter (see McPhee 1963).

As an alternative, we suggest that people respond by comparing "new" opinions against the opinions of other people in their personal networks. If one is a Gore supporter, and an acquaintance suggests that Bush is an excellent presidential candidate, then a switch to favor Bush should only be forthcoming if many other acquaintances also support Bush. People tend to maintain an opinion as long as there is support for it in their networks of personal contacts. Any single piece of information is seen within the context of all the information that is available. The social influence of any single interaction ceases to be

determinate, and the agent becomes an evaluator of information received through a successive autoregressive process of social interaction.

Autoregressive Influence and the Durability of Disagreement

In this final section we consider the consequences of an autoregressive influence process for the durability of disagreement in the agent-based model. The discussants are selected in the same manner as previous model, but agents keep records on the contacts they have experienced and use those records when formulating their response to new points of view. The implementation of the autoregressive logic is somewhat primitive, but it captures the essential logic we are trying to portray. Each time an agent interacts, it counts the number of features it holds in common with the other. When an interaction occurs, and the other offers a different opinion on a randomly chosen feature, then the agent polls the people that it agrees with on more than one-half of the issues, and if more than one-half of those "friends" agree with the new point of view, it is adopted. Thus, new ideas or novel preferences should take longer to catch on, and individual agents should be less susceptible to persuasion.

As Figure 5 shows, when the influence of an opinion is proportional to its incidence within an individual's network of contacts, diversity is maintained both within the larger population and within networks of political communication. First, the level of acquaintanceship is lower than in the previous models, reflecting the fact that the opinions of the agents are not so homogeneous. People are regularly put in contact with others with whom they disagree on a randomly chosen issue. Second, only a relatively small proportion of networks are composed of dyads with identical preferences. Finally, the average proportional agreement with any discussion partner (harmony) is only slightly above .6. That

value, which is consistent with the empirical results, indicates that among discussants there is a considerable level of agreement, but by no means complete homogeneity.

We hasten to add that this figure is very much representative of the 100 runs we performed with these settings. The number of steps to convergence averaged 732.1 with a standard deviation of 148. However, the averages (and standard deviations) of the acquainted, harmony, and identical variables were .44 (0.03), .36 (0.05), and .6 (0.04). Entropy is not zero at the end of any of the runs.

What do these results suggest? First and foremost, these results point to the importance of separating the ***communication*** of information from the ***persuasiveness*** of information. Even effectively communicated messages may lack influence, and this analysis points to the importance of interdependent citizens as discriminating consumers of political information. Second, these results suggest that political influence is imbedded within an autoregressive process of social influence. People judge new information the context of old information, and to the extent that new information does not correspond with information that they have already collected, it is less likely to be persuasive. Finally, the autoregressive structure that underlies political influence is responsible, perhaps ironically, for sustaining political heterogeneity and diversity within the larger population.

In order to demonstrate the stabilizing role of social networks, we introduce one piece of analysis that derives from the larger model. This final model exposes agents to a greater variety of inputs and, as a result, generates larger, more complicated, personal networks. In this larger model, there are five 10x10 "home" neighborhoods, and each day the agents spend part of their time at home and some also go into three 5x5 "work" grids. Agents initiate interactions at random, choosing randomly among others that they happen to find and, when

they find someone, they interact with probability equal to their similarity. When presented with a difference, an agent will adopt the other's opinion if more than one-half of its "friends" (people with whom it has agreed in the past on more than one-half of the issues) support that new opinion.

At the outset, agents do not have any "friends" and they are just wandering about, forming acquaintances and making notes. After a few iterations, however, they begin to influence each other. Out of 500 agents, the number who are persuaded to change in each day is typically less than 10, and as the networks stabilize, that number declines. After 1000 days (time = 10,000 timesteps), as is shown in Figure 6a, the overall experiences have stabilized: agents report neither complete homogeneity nor complete strife. After 2000 days (time = 20,000 timesteps) we stopped the simulation and archived its state. Then we subjected the 500 agents in the system to a series of "random shocks" or pulses. The impact of these exogenous impacts is demonstrated in Figure 6b. Recall that the opinions can be valued 0, 1, or 2, and that the original conditions in the model are randomly assigned, so the average of opinion at the outset is always close to 1.0.

The results indicate, generally, that pulses which go against the grain of the local networks are soon forgotten, while large pulses which go along with predominant tendencies are longer lasting. For example, at 20,000 timesteps, we selected (at random) 50 agents for whom feature 4 was 0 and changed it to 2. (The pulse effect is labeled "A" in Figure 6b.) After a while, most of the impact of the opinion pulse has been erased by the regular pattern of re-exposure to views in the local network.

On the other hand, at 25,000 timesteps, we took feature 1, on which the average was 1.2, and changed 50 agent opinions from 0 to 2. That pulse, labeled B, has a more permanent effect because there was more support for that view in the residual networks.

Finally, to demonstrate one of the interesting and complicated aspects of these models, consider the small pulses labeled C. Each small spike in feature 2 represents the fact that we have found 25 agents and changed their opinion from 1 to 2 on that feature. It appears that the autoregressive influence in the networks quickly erases each of the 5 shocks that we applied on feature 2 at 1000 timestep intervals. If anything, the small positive shocks are followed by a decline in the average for that feature. However, a side-effect of those "tweaks" is a shift on features 0 and 3, which presumably results from the temporary ripples in opinion on feature 2 which create new "friendships" and break old ones. The simulation run shown in Figure 6b is broadly representative of our experience with the model.

The results of this larger model should serve to address the concern that the stabilizing impact of the autoregressive influence process is an artifact of the small (maximum: 4) networks that are allowed in the original Axelrod design. The average number of others with whom each agent has come in contact is 27 and the number of friends is 14. While these values are considerably greater than we find in cross-sectional survey data, they do serve to drive home one important point: diversity is not being preserved by isolating agents from opinions with which they disagree. Rather, diversity is preserved within the small personal networks by giving the agents the ability to respond to the diversity of their experience in a sensible way.

Summary and Conclusion

A substantial body of evidence has accumulated regarding the distribution of preferences within citizens' networks of political communication. Contrary to a great deal of conventional wisdom, these networks are ***not*** safe havens from political disagreement. Indeed, it would appear that disagreement is the modal condition among citizens – most citizens experience disagreement and divergent political preferences within these networks. Moreover, this conclusion is based on the closely held, self-reported relationships of the citizens themselves, and on perhaps the most visible of contemporary political choices – support for a particular presidential candidate. Hence the question becomes, what is the nature of the dynamic process that sustains disagreement among citizens?

As long as persuasion is the inevitable consequence of interaction within discrete dyads, the elimination of political diversity and disagreement may be a foregone conclusion, at least over the long haul. In contrast, a far different outcome emerges when we treat persuasion within dyads as a problematic and less than automatic consequence of interaction across an individual's entire network of contacts. Based on earlier empirical results, we conceive the probability of persuasion as a function of an opinion's incidence within an individual's network of relationships. That is, individuals are less likely to be persuaded by opinions that win only limited support among the participants within their communication networks. Indeed, this model of persuasion serves to maintain diversity and disagreement both in the short run and over the long haul.

In many ways this is a surprising outcome. The model of influence we are describing rewards majority opinion at the same time that it punishes the political minority, but it produces an aggregate outcome in which the minority does not disappear. The potential of

this mechanism for maintaining political disagreement is that the influence of majorities and minorities are defined according to the distribution of opinion within closely held micro-environments of political communication. Hence, people are able to resist divergent viewpoints within the network because every opinion is filtered through every other opinion.

Finally, the power of the mechanism we are describing is wholly dependent on the low levels of network density that are built into the model (Granovetter 1973; Burt 1992). If the network densities were high – if networks were wholly self-contained so that all members shared the same interaction partners – then disagreement would disappear even though diverse preferences would be sustained in the larger environment. That is, no one would ever encounter diverse preferences because every particular network is wholly self-contained and entirely homogeneous. In contrast, low network densities, combined with influence that is predicated on the incidence of particular opinions within networks, serve to sustain political diversity in the larger environment as well as the experience of disagreement within citizens' closely held networks of political communication.

References

Asch, S.E. 1956. Studies on Independence and Conformity: A Minority of One Against a Unanimous Majority. Psychological Monographs 70: 416.

Axelrod, Robert. 1997. The Complexity of Cooperation. Princeton: Princeton University Press.

Berelson, Bernard R., Paul F. Lazarsfeld, and William N. McPhee. 1954. Voting: A Study of Opinion Formation in a Presidential Election. Chicago: University of Chicago Press.

Burt, Ronald S. 1992. Structural Holes. Cambridge, MA: Harvard University Press.

Downs, Anthony. 1957. An Economic Theory of Democracy. New York: Harper and Row.

Fabrigar, Leandre R. and Jon A. Krosnick. 1995. ``Attitude Importance and the False Consensus Effect." Personality and Social Psychology Bulletin 21:468-479.

Festinger, Leon. 1957. A Theory of Cognitive Dissonance. Palo Alto, California: Stanford University Press.

Granovetter, Mark. 1973. "The Strength of Weak Ties," American Journal of Sociology 78: 1360-80.

Huckfeldt, Robert, Paul E. Johnson, and John Sprague. 2002. "Political Environments, Political Dynamics, and the Survival of Disagreement," Journal of Politics: in press.

Huckfeldt, Robert, Paul E. Johnson, John Sprague, and Michael Craw. 2000. Influence, Communication, and the Survival of Political Disagreement among Citizens." Prepared for delivery at the annual meeting of the Midwest Political Science Association, September, 2000.

Huckfeldt, Robert, Paul A. Beck, R. Dalton, and Jeffrey Levine. 1995. "Political Environments, Cohesive Social Groups, and the Communication of Public Opinion," American Journal of Political Science 39: 1025-1054.

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

Huckfeldt, Robert, John Sprague, and Jeffrey Levine. 2000. "The Dynamics of Collective Deliberation in the 1996 Election: Campaign Effects on Accessibility, Certainty, and Accuracy," American Political Science Review 94: 641-651.

Johnson, Paul E. and Robert Huckfeldt. 2001. Persuasion and Political Heterogeneity within Networks of Political Communication: AgentBased Explanations for the Survival of Disagreement. Prepared for delivery at the annual meeting of the Midwest Political Science Association, September, 2001. (updated version available: http://lark.cc.ukans.edu/~pauljohn/ResearchPapers/APSA01)

Johnson, Paul E. 1999. Protests, Elections, and Other Forms of Political Contagion. Paper presented at the Annual Meeting of the American Political Science Association, Atlanta, GA.

Johnson, Paul E. and Alex Lancaster. 2000. Swarm User Guide. URL: http://www.santafe.edu/projects/swarm/swarmdocs/userbook/userbook.html.

Lazarsfeld, Paul, Bernard Berelson, and Hazel Gaudet. 1948. The People's Choice. New York: Columbia University Press.

Latane, Bibb, Andrzej Nowak, and James H. Liu. 1994. Measuring emergent social phenomena: dynamism, polarization, and clustering as order parameters of dynamic social systems. Behavioral Science, 39:1-24.

McPhee, William N. , with Robert B. Smith and Jack Ferguson. 1963. "A Theory of Informal Social Influence." In William N. McPhee, Formal Theories of Mass Behavior. New York: Free.

Miller, Warren. 1956. "One Party Politics and the Voter," American Political Science Review 50: 707-725.

Minar, Nelson, Roger Burkhart, Christopher Langton, and Manor Askenazi. 1996. The Swarm Simulation System: A Toolkit for Building Multi-Agent Simulations. Technical Report 96-04-2, Santa Fe Institute, Santa Fe, New Mexico.

Nowak, Andrzej, and Maciej Lewenstein, "Modeling Social Change with Cellular Automata," in R. Hegselmann et al., eds. Modeling and Simulation in the Social Sciences from a Philosophy of Science Point of View. Amsterdam: Kluwer, pp. 249-285.

Nowak, Andrzej, Jacek Szamrej, and Bibb Latané. 1990. From Private Attitude to Public Opinion: A Dynamic Theory of Social Impact." Psychological Review 97(3): 362-376.

Table 1. Level of diversity within political communication networks during the 2000 election.

A. Percent of network voting for Gore by respondent's vote.

	Gore	neither	Bush
None (0%)	15.9	56.6	64.0
Some	44.2	28.7	28.8
All (100%)	40.0	14.8	7.2
	473	244	430

B. Percent of network voting for Bush by respondent's vote.

	Gore	neither	Bush
None (0%)	62.8	49.6	12.8
Some	31.7	32.8	41.2
All (100%)	5.5	17.6	46.0
	473	244	430

Source: 2000 National Election Study.

Table 2. Respondent agreement with discussants who support Bush, Gore, and neither candidate by partisanship of the respondent and distribution of preferences in the residual network. (Logit models.)

	respondent agreement with a discussant who supports:		
	Bush	Gore	neither
constant	-.41	-.28	1.46
	(1.73)	(1.20)	(4.17)
party identification	.83	-.81	
	(11.19)	(10.00)	
partisan strength			-.73
			(5.06)
residual network support for:			
Bush	.71	-.16	-.64
	(4.30)	(1.30)	(3.99)
Gore	-.39	.57	-.94
	(2.57)	(3.21)	(4.98)
neither	-.24	-.07	-.06
	(1.48)	(.40)	(.34)
N (clusters)	1183(665)	1071(649)	545(395)
Chi-square/df/p	160/4/.00	121/4/.00	52/4/.00

respondent agreement within dyad: 1=if respondent perceives that the discussant supports their own candidate choice, 0=other
party identification: seven point scale= -3 (strong Democrat) to 3 (strong Republican
partisan strength: four point scale= 0 (independent) to 3 (strong partisan)
residual network support for Bush: number of discussants in remainder of network that are perceived to support Bush
residual network support for Gore: number of discussants in remainder of network that are perceived to support Gore
residual network support for neither: number of discussants in remainder of network that are perceived to support neither candidate

Figure 1. Predicted probability of agreement within network dyads, by the candidate preference of the discussant in the dyad and the levels of candidate support in the remainder of the network. All respondents are independents.

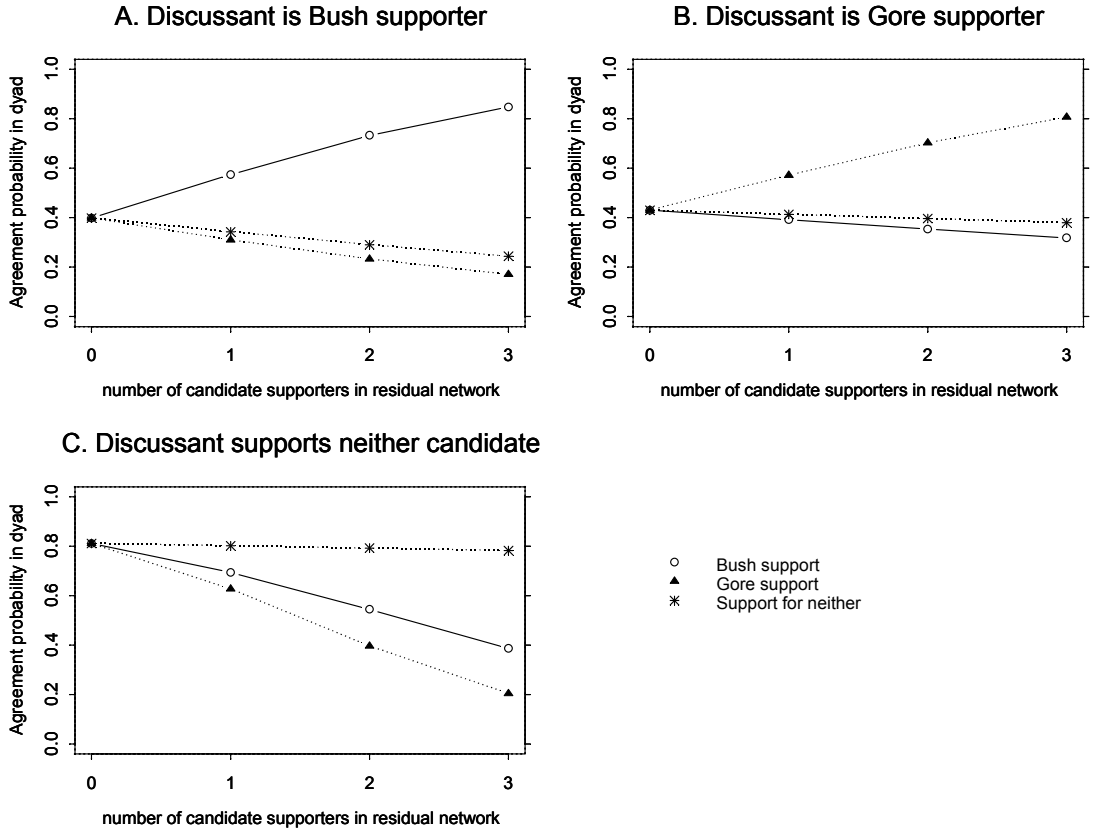

Note: When the level of support for a particular candidate preference in the residual network is held constant at 0, 1, 2, or 3 discussants, support for the other preferences is held constant at 0.

Figure 2. Predicted probability of agreement within network dyads, by the candidate preference of the discussant in the dyad and the levels of candidate support in the remainder of the network. All respondents are strong Republicans.

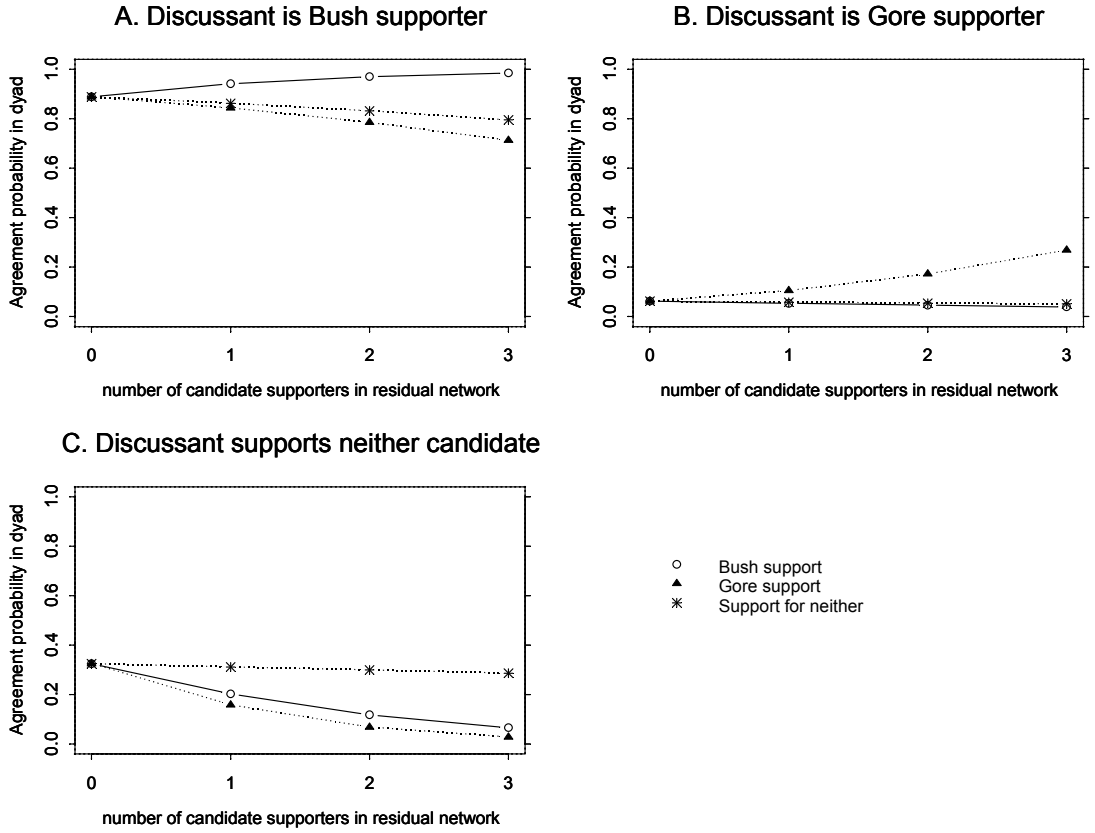

Note: When the level of support for a particular candidate preference in the residual network is held constant at 0, 1, 2, or 3 discussants, support for the other preferences is held constant at 0.

Figure 3. Predicted probability of agreement within network dyads, by the candidate preference of the discussant in the dyad and the levels of candidate support in the remainder of the network. All respondents are strong Democrats.

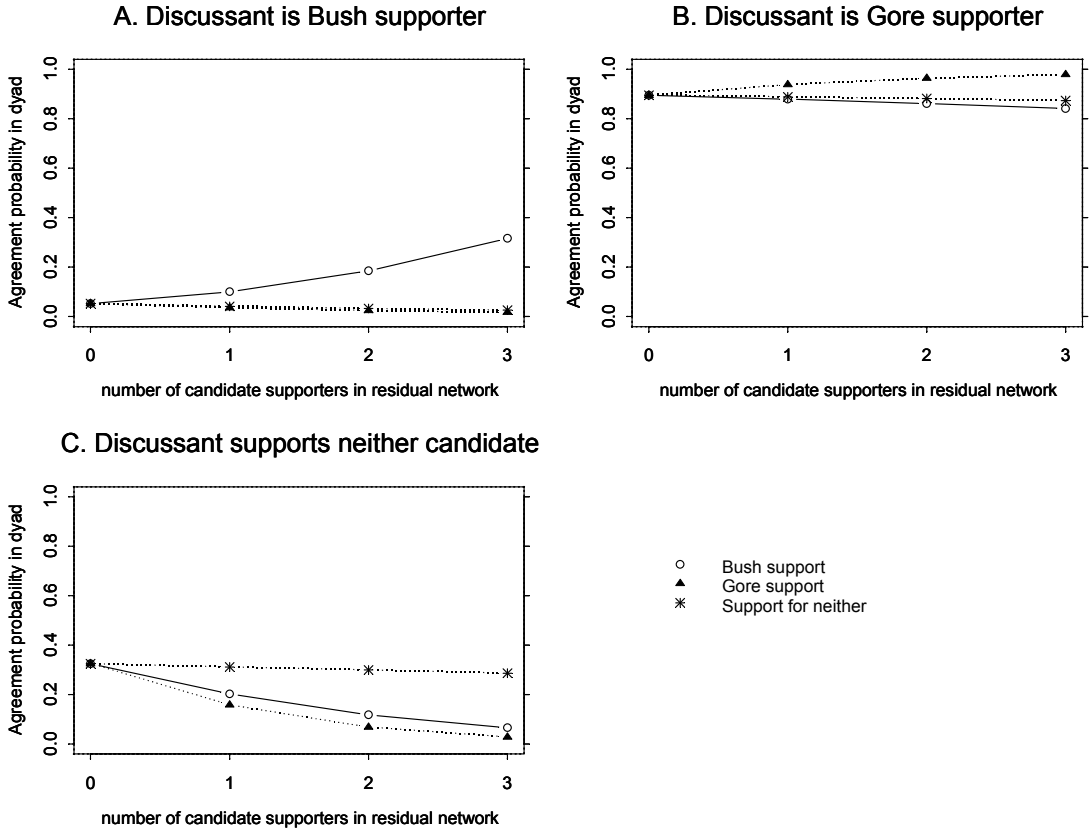

Note: When the level of support for a particular candidate preference in the residual network is held constant at 0, 1, 2, or 3 discussants, support for the other preferences is held constant at 0.

Figure 4. Baseline model of communication and influence.

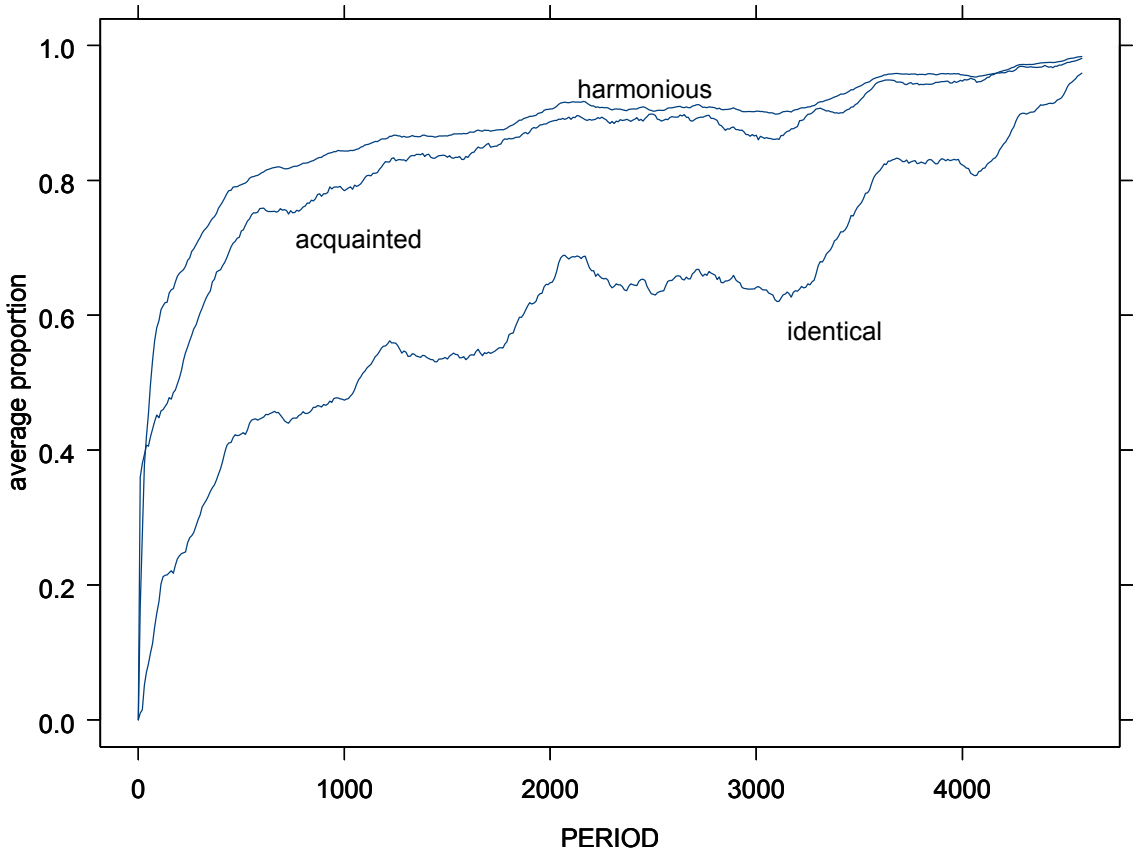

Figure 5. Autoregressive influence model.

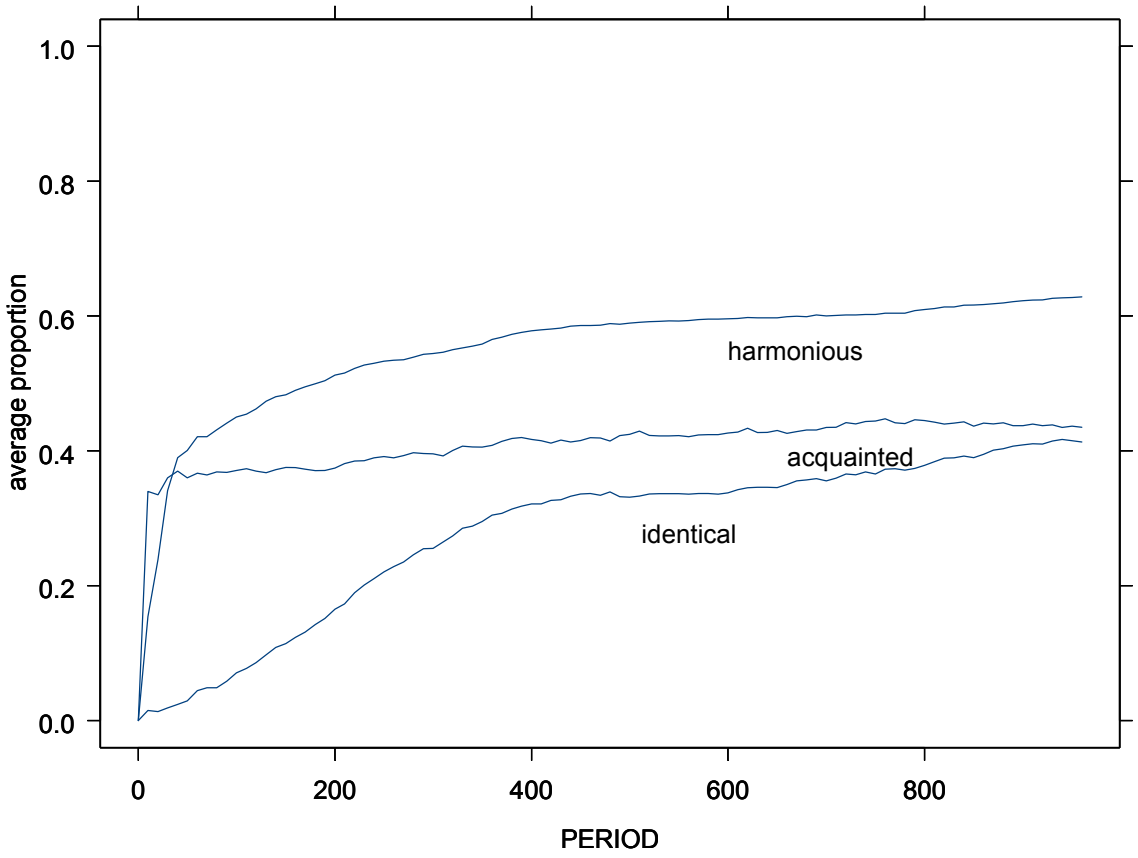

Figure 6a

Diversity and stability in a model with 5 home neighborhoods and 3 workplaces

A. Experience of Diversity

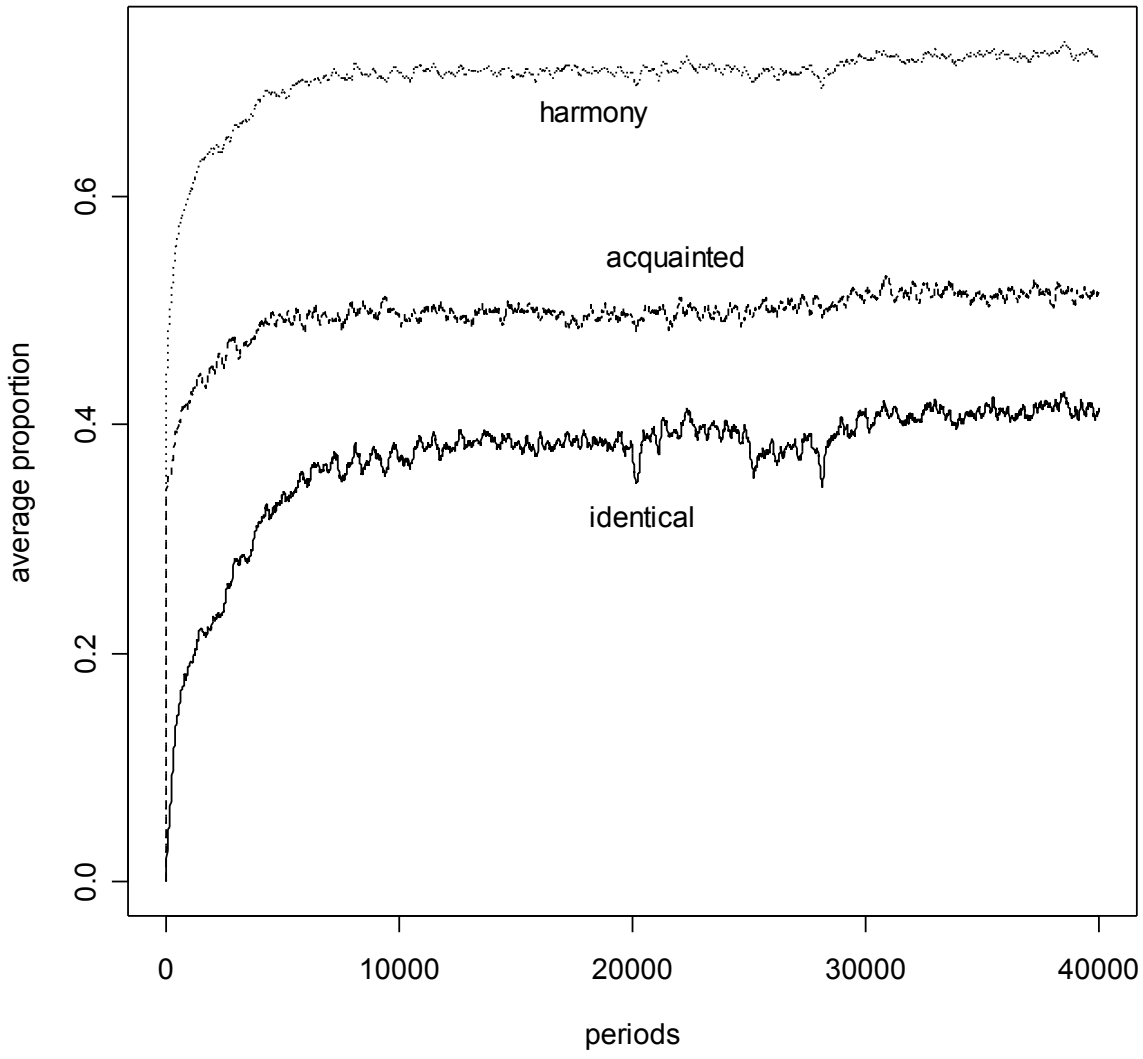

B. The Impact of Exogenous Shocks