Structure and Strategy in Collective Action[1]

Michael Suk-Young Chwe
University of Chicago

This article considers both structural and strategic influences on collective action. Each person in a group wants to participate only if the total number taking part is at least her threshold; people use a network to communicate their thresholds. People are strategically rational in that they are completely rational and also take into account that others are completely rational. The model shows first that network position is much more important in influencing the revolt of people with low thresholds than people with high thresholds. Second, it shows that strong links are better for revolt when thresholds are low, and weak links are better when thresholds are high. Finally, the model generalizes the threshold models of Schelling (1978) and Granovetter (1978) and shows that their findings that revolt is very sensitive to the thresholds of people "early" in the process depends heavily on the assumption that communication is never reciprocal.

INTRODUCTION

Collective action has been studied in two largely disjoint approaches, one focusing on the influence of social structure and another focusing on the incentives for individual participation. These approaches are often seen

[1] This paper has benefited from comments from participants in the Stanford Institute for Theoretical Economics, August 1997; the Social Interactions and Aggregate Economic Behavior conference at the Santa Fe Institute, August 1996; the Social Learning Workshop at State University of New York, Stony Brook, July 1996; and workshops at University of Alabama; California Institute of Technology; Johns Hopkins University; Princeton University; Rutgers University; University of California, Davis; University of California, Irvine; University of California, Los Angeles; University of Chicago; University of Illinois in Chicago; University of Illinois in Urbana-Champaign; University of Iowa; Iowa State University; University of Minnesota; University of Southern California; University of Western Ontario; and Vanderbilt University. Locally, I have benefited from suggestions received in the Social Organization of Competition Workshop, the Economic Theory Workshop, the Economic Theory Bag Lunch, and from students in my graduate game theory classes. I also appreciate the encouragement and suggestions of Ron Burt, Steve Durlauf, Jim Fearon, Eric Friedman, Douglas Gale, Lee Heavner, Morgan Kelly, Alan Kirman, Ben Klemens, Timur

0002-9602/2000/10501-0004$02.50

as competing or even opposed. This article describes a simple model that tries to bridge this methodological division, using concepts from both social network theory and game theory. Here, a group of people face a collective action problem in that an individual wants to participate only if joined by enough others; exactly how many total participants are necessary is given by the individual's "threshold." Individuals are located in a social network, and each person knows the thresholds of only her neighbors in the network; each person has "local knowledge." People are strategically rational: they are completely rational and make decisions knowing that others are completely rational.

The model makes three substantive points. First, many empirical studies of collective action make the assumption, implicit in linear regression, that a person's individual characteristics and social position enter into participation linearly and independently. My model suggests that this might not be valid. In this model, people with low thresholds, who are highly predisposed toward participation, are affected much more by social position than people with high thresholds. Intuitively, whether a low-threshold person participates or not depends greatly on whether that person happens to have some sympathetic friends, while a high-threshold person participates only if a great mass of people participate. Second, although widely scattering "weak links" seem to be better for widespread communication than more involuted "strong links" (Granovetter 1973), empirical researchers have found that strong links, not weak links, correlate positively with participation (e.g., McAdam 1986; McAdam and Paulsen 1993; Valente 1995). Our model helps resolve this puzzle by showing that when thresholds are low, strong links can be better for participation. Weak links are better at spreading information widely, but strong links are better at locally creating the common knowledge, that is, knowledge of other people's knowledge, essential for collective action. Third, the model generalizes the threshold models of Schelling (1978) and Granovetter (1978) and shows that their finding that collective action is very sensitive to the thresholds of people "early" in the process depends heavily on the assumption that communication is never reciprocal. When even a small possibility of reciprocal communication is allowed, collective action is fairly robust. These three points are made using randomly generated networks of 30 people. I conclude by discussing the model in relation to existing models, especially on the issue of network transitivity.

Kuran, Michael Macy, Stephen Morris, John Padgett, Jesper Sorensen, Peyton Young, George Zanjani, and the *AJS* referees. Finally, I am grateful for financial support from the National Science Foundation (grant SBR-9712277). Direct correspondence to Michael Chwe, Department of Economics, University of Chicago, Chicago, Illinois 60637. Web site: www.chwe.net. E-mail: michael@chwe.net

THE MODEL

There is a group of n people, and each person chooses either to revolt r (participate in the collective action) or stay at home s (not participate). Each person i has an idiosyncratic threshold $\theta_i \in \{1, 2, \ldots, n + 1\}$; a person wants to revolt only if the total number of people who revolt is greater than or equal to her threshold. For example, a person with threshold 2 prefers to revolt if he is joined by at least one other; a person with threshold n prefers to revolt only if everyone else does. The social network $\rightarrow$ is a binary relation over N, where $j \rightarrow i$ means that person j talks to person i. We define $B(i) = \{j \in N: j \rightarrow i\}$ to be person i's "neighborhood," the set of people who talk to i. We assume that $\rightarrow$ is reflexive ($i \rightarrow i$) and thus $i \in B(i)$. The idea is that person i knows the thresholds of only the people in her neighborhood $B(i)$. We also assume that person i knows all network relations among the people in $B(i)$; in other words, for all j, $k \in B(i)$, he knows whether or not $j \rightarrow k$.

To model this as a game, we have to specify information, strategies, and payoffs and define equilibrium. The details of this are in the appendix; here, I try to explain the game using some simple examples. The key modeling principle here is that a person's knowledge determines his ability to distinguish between states of the world, and if a person cannot distinguish between several states of the world, he must take the same action in all of them. For example, say there are only two people: person 1 has a threshold of either 1, 2, or 3, and person 2 also has a threshold of either 1, 2, or 3. Hence, there are nine possible states of the world: 11, 12, 13, 21, 22, 23, 31, 32, 33, where 23 is the state in which person 1 has threshold 2 and person 2 has threshold 3, for example.

Say that person 1 and person 2 do not communicate; we have the "null network," as shown in figure 1 (of course $1 \rightarrow 1$ and $2 \rightarrow 2$, but throughout we leave out these "loops" for clarity). Hence, each person only knows his own threshold. We can represent person 1's knowledge by the sets {11, 12, 13}, {21, 22, 23}, {31, 32, 33}, which form a partition of the nine possible states of the world, as shown in figure 1. The idea here is that if two states of the world are in the same set, then she cannot distinguish between them; for example, person 1 cannot distinguish between states 21, 22, and 23 because she does not know person 2's threshold. Similarly, since person 2 only knows his own threshold, his partition is {11, 21, 31}, {12, 22, 32}, {13, 23, 33}, also shown in figure 1.

Person 1 chooses whether to revolt or stay at home given a state of the world. If her threshold is 1, then she is happy to revolt all by herself and thus chooses to revolt in states 11, 12, and 13, as shown in figure 1. If her threshold is 3, then she never wants to revolt under any circumstances and thus chooses to stay at home in states 31, 32, and 33. If her threshold

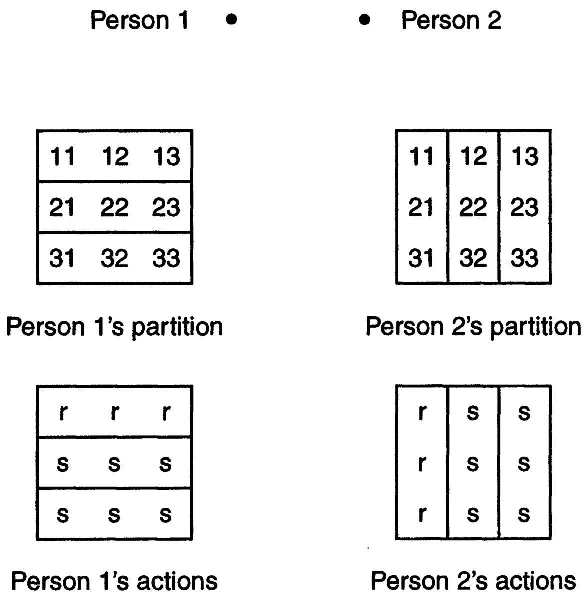

FIG. 1.—Null network (no communication)

is 2, then ideally person 1 would like to revolt in state 21 (in which case person 2 will surely revolt) and stay at home in state 23 (in which case person 2 will surely stay at home). But she cannot distinguish between these states; she does not know person 2's threshold. If she cannot distinguish between two states, then she cannot condition her action on them. In other words, she must revolt in all three states 21, 22, and 23 or stay at home in all three states. If she chooses to revolt, then person 2 might revolt (if the state is 21) but might not (if the state is 23). In other words, if she revolts, then there is the possibility that the total number of people revolting is less than her threshold. We assume (for the sake of modeling simplicity) that a person gets a very large negative payoff or penalty if this happens and hence a person revolts only if she knows for certain that enough others will revolt. Hence, person 2 decides to stay at home in states 21, 22, and 23, as shown in figure 1. Similarly, person 2's actions are also shown in figure 1.

Now consider the case when we have the "complete graph" network in which $1 \rightarrow 2$ and $2 \rightarrow 1$, as shown in figure 2 (when an arc is symmetric, we leave out the arrows for convenience). Here, each person knows each other's threshold. Now person 1's partition (and also person 2's) is {11}, {12}, {13}, {21}, {22}, {23}, {31}, {32}, {33}, as shown in figure 2. Now each person can distinguish between all states of the world.

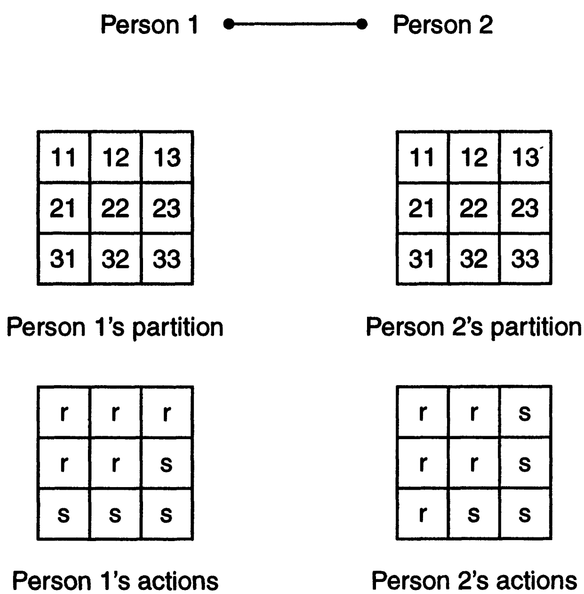

FIG. 2.—Complete network (full communication)

As before, person 1 revolts in states 11, 12, and 13 and stays at home in states 31, 32, and 33. Before, when person 1 had threshold 2, since he could not distinguish between states 21, 22, and 23, he chose to stay at home in these states. Now, person 1 can distinguish between these states and can take different actions in them. At state 21, person 1 revolts since he knows that person 2 has threshold 1 and will hence revolt. At state 23, person 1 stays at home since he knows that person 2 has threshold 3 and hence will not revolt. Similarly, person 2 revolts in state 12 and stays at home in state 32.

At state 22, when both people have threshold 2, person 1 wants to revolt if person 2 revolts, and person 2 wants to revolt if person 1 revolts. Hence, if both revolt, this is an equilibrium in the sense that each person is making the best choice given what the other person is doing (this is the standard concept of Nash equilibrium). It is also an equilibrium for both people to stay at home (if you do not revolt, I do not want to revolt, and if I do not revolt, you do not want to revolt). In our model, whenever there is this kind of indeterminacy (there is more than one equilibrium), we assume that the equilibrium that occurs is the one in which the most revolt takes place. That is, when two people with threshold 2 discover each other, we assume that they revolt.

To make a prediction for more general networks with n people, it is

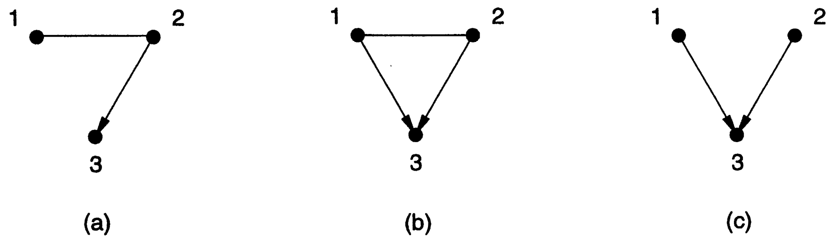

FIG. 3.—Three networks: everyone has threshold 2

simply a matter of writing down all states of the world, each person's information partition, and proceeding in a similar way. This becomes difficult, however, for large n since there are $(n + 1)^n$ states of the world. But it is possible to solve the model without writing down all states of the world (of course, all but the simplest cases still require a computer; a program is available from the author). To take an example, consider the three networks in figure 3 and say everyone has threshold 2; thus, the true state of the world can be written 222.

First, consider network (*a*). Here, person 1 knows his own and person 2's threshold; thus, he knows that the state of the world is either 221, 222, 223, or 224. Person 2 knows her own and person 1's threshold and thus similarly knows that the state of the world is either 221, 222, 223, or 224. Person 3 knows his own and person 2's threshold and thus knows that the state of the world is either 122, 222, 322, or 422. Who revolts? If person 1 revolts at states 221, 222, 223, and 224, and person 2 revolts in states 221, 222, 223, and 224, this is an equilibrium since in all four states each person can count on the other to revolt. Will person 3 revolt? As far as person 3 can tell, the state of the world could be 422. In state 422, person 1 has threshold 4 and surely stays at home. In state 422, person 2 also stays at home because he cannot count on having a partner: person 1 stays at home, and person 2 does not know anything about person 3's threshold. So at state 422, person 3 will surely not have a partner. Since person 3 cannot distinguish between the states 122, 222, 322, or 422, he will not take the risk and hence will stay at home. So only persons 1 and 2 will revolt.

Note that here person 3 does not revolt even though he has threshold 2 and has a neighbor, person 2, who revolts. Person 3 knows that person 2 will revolt if person 1 is a willing companion, but person 3 does not know anything about person 1; as far as person 3 knows, person 1 could have threshold 4, in which case person 2 will surely stay at home.

Now, consider network (*b*). Again, persons 1 and 2 know that the state of the world is either 221, 222, 223, or 224. Now, person 3 knows every-

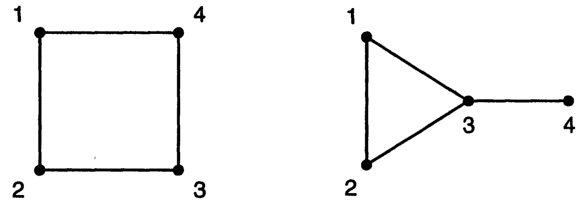

FIG. 4.—Square and kite

one's threshold and thus knows that the state of the world is exactly 222. Who revolts now? As before, person 1 and person 2 can count on each other and revolt in all four states 221, 222, 223, or 224. Person 3 knows that the state of the world is 222 and thus persons 1 and 2 will both revolt; hence, he revolts also. So now everyone revolts.

Finally, consider network (*c*). As before, person 3 knows everyone's threshold, but now person 1 only knows his own threshold and person 2 only knows her own threshold. As far as person 1 can tell, the state of the world could be any state in the set {211, 212, 213, 214, 221, . . . , 244}, and thus person 1 does not revolt. Similarly, person 2 does not revolt. Person 3 knows that the state is 222, but since persons 1 and 2 do not revolt at state 222, person 3 stays at home.

Note that person 3 revolts in network (*b*) but not in network (*c*), even though his knowledge of others' thresholds is the same in both cases. The reason is simply that in network (*b*) person 3 knows that persons 1 and 2 can communicate reciprocally and hence revolt, while in network (*c*), person 3 knows that they cannot. In other words, it is crucial here that person 3 knows whether persons 1 and 2 know each other; our model relies on the assumption that each person knows the relationships among her neighbors.

PLURALISTIC IGNORANCE AND COMMON KNOWLEDGE

Now, say that we have four people, each with threshold 3, and consider two networks, the "square" and "kite," as shown in figure 4. These two networks nicely illustrate that a person's revolt depends not only on him knowing others' thresholds but also on him knowing what other people know.

Consider first the square. Since everyone has threshold 3, the actual state of the world can be written 3333. In the square, person 1 knows that person 2 and person 4 have threshold 3 but knows nothing about person 3's threshold. Hence, all person 1 knows is that the state of the world is in the set {3313, 3323, 3333, 3343, 3353}, where state 3353, for example, is the state in which person 3 has threshold 5 and everyone else has threshold 3. Person 1 will revolt only if he knows that two other people will revolt in each of these states. Consider state 3353, which is perfectly possible as far as person 1 can tell. At state 3353, what, for example, will person 2 do? Person 2 does not know person 4's threshold and thus cannot distinguish between the states in the set {3351, 3352, 3353, 3354, 3355}. If person 2 revolts in these states, it is possible that she will have only one partner (the state could be 3355, in which case persons 3 and 4 never want to revolt). Hence, person 2 will stay at home at state 3353. So, if person 1 revolts in state 3353, at best only person 4 will join him. Hence, person 1 does not revolt at the states {3313, 3323, 3333, 3343, 3353} and thus does not revolt at state 3333. We can show similarly that persons 2, 3, and 4 do not revolt at state 3333 either.

Consider now the kite. Here, person 3 knows everyone's threshold and thus knows that the state of the world is 3333. Persons 1 and 2 know everything but person 4's threshold; they thus know the state of the world is in {3331, 3332, 3333, 3334, 3335}. Person 4 only knows his own and person 3's threshold; he thus knows that the state of the world is in the set {1133, 1233, . . . , 5533}. Person 4 clearly does not revolt; as far as he can tell, the state of the world is 5533, in which case he has at most one willing partner. But, if persons 1, 2, and 3 revolt in all states in {3331, 3332, 3333, 3334, 3335}, this is an equilibrium since in all three states at least three people revolt and each person has threshold 3.

So, according to our model, no one revolts in the square, but three people revolt in the kite. Note that this difference cannot be accounted for by macroscopic characteristics such as the total number of links (four symmetric links in both cases), or even by finer measures such as the number of neighbors each person has (in the kite, two revolters have only two neighbors, as in the square). The difference between the square and kite is truly a structural difference.

What is the difference between the square and the kite? In the square, each person knows that there are three people with threshold 3: himself and his two neighbors. That is, each person knows that conditions are such that revolt is possible. But, say I am considering whether to revolt. What do I know about, say, my neighbor to the right? I know that he has threshold 3. I also know that he knows that there is at least one other person with threshold 3: me. But I do not know anything about his other neighbor ("across" from me). Hence, I cannot count on him revolting, and so I do not revolt. In this case, even though everyone knows that revolt is possible, no one in fact revolts. In the kite, however, each person in the "triangle" knows not only that two other people have threshold 3, but they also know that these two other people know each other's threshold.

In other words, in the square, each person knows that there is sufficient discontent (there are three people with threshold 3), but each person does not know that anyone else knows. It is not enough for everyone to simply know that there is sufficient discontent; what is required is "common knowledge" (Lewis 1969; Aumann 1976): everyone has to know that there is sufficient discontent, everyone has to know that everyone knows, everyone has to know that everyone knows that everyone knows, and so on. In the kite, among members of the triangle, the fact that there are three people with threshold 3 is common knowledge: each person knows it, knows that every other person knows it, and so on.

The importance of common knowledge for collective action can be demonstrated in a wide variety of contexts (starting with Lewis [1969]; see also, e.g., Rubinstein 1989; Chwe 1998). Several related ideas exist in the literature. One is the concept of "pluralistic ignorance" from social psychology (Allport 1924; Katz and Allport 1931; O'Gorman 1986), which refers to a situation in which people hold very incorrect beliefs about the beliefs of others. For example, in a 1972 survey, 15% of white Americans favored racial segregation, but 72% believed that a majority of the whites in their area favored segregation (O'Gorman 1979; see also Shamir 1993). Pluralistic ignorance is not usually understood as related to collective action; more often it is understood as a matter of individual psychology (O'Gorman 1986; Mullen and Hu 1988): a person reduces dissonance by thinking that her own view is the majority view, for example. Recently, however, it has been applied to the former Soviet Union and Eastern European states, the idea being that dissatisfaction was widespread but that few people knew how widespread it was. These accounts focus on limited communication due to criminal penalties for self-expression, a government-controlled press, and in the spirit of our model, a lack of social ties (Coser 1990; Kuran 1991). "The reduction of pluralistic ignorance," due to modern communication technology and increased foreign contacts, led to "a political wave of tremendous power" (Coser 1990, p. 182).

Another related idea is James Scott's (1990) distinction between the "public transcript," what subordinates say when talking publicly to their superiors, and the "hidden transcript," what subordinates say among themselves. For example, "if the sharecropping tenants of a large landowner are restive over higher rents, he would rather see them individually and perhaps make concessions than to have a public confrontation" (Scott 1990, p. 56). This is because public declarations enable subordinates to find out about each other: *"It is only when this hidden transcript is openly declared that subordinates can fully recognize the full extent to which their claims, their dreams, their anger is shared by other subordinates"* (Scott 1990, p. 223; emphasis in the original). The knowledge of others, which makes revolt possible, is formed not only by public declarations but also by social structure: "the informal networks of community . . . through kinship, labor exchange, neighborhood, ritual practices, or daily occupational links" (Scott 1990, p. 151).

Even more explicitly, Roger Gould (1995, pp. 18–20) argues that "potential recruits to a social movement will only participate if they see themselves as part of a collectivity that is sufficiently large and solidary to assure some chance of success through mobilization. A significant source of the information they need to make this judgement is . . . social relations [which are] the mechanism for mutual recognition of shared interests (and of recognition of this recognition, and so on)." Translated into the language of our model, a person will participate only if she is assured that enough others will (the number of people revolting will be at least her threshold), and crucial in that decision is what she learns about the interests (thresholds) of others in her social network. Also, it is not just a matter of recognizing other people's interests (knowing other people's thresholds) but recognizing other people's recognition (knowing what others know), and so forth.

This idea, that social networks influence a group's knowledge of itself and thereby influence its ability to collectively act, is plausible and relevant enough to have arisen independently in multiple contexts. By giving it explicit mathematical form as in our model, it can be used not just to broadly motivate an argument but also to make specific predictions.

DYNAMICS

So far, our model is static. A simple way to add some dynamics is to assume that each person's neighborhood expands in time. We assumed earlier that a person's neighborhood is those people who are immediately adjacent, in other words, within one link away; we can think of this as the first "period" in time. We can make the simple assumption that in the second period each person's neighborhood is those people who are within two links away; in the third period, each person's neighborhood is those people within three links; and so on. To say this precisely, define the distance $d(j, i)$ from person j to person i to be the length of the shortest path from j to i. Define $B(i, q) = \{j \in N: d(j, i) \leq q\}$ to be the "neighborhood of radius q" centered at i. The assumption is that at time t, person i's neighborhood is $B(i, t)$: he knows the thresholds of everyone in $B(i, t)$ and the network relations among the people in $B(i, t)$. The idea here is that information flows through the network over time: first people know about their neighbors, second they learn about their neighbors' neighbors too, and so on. It is easy to show that as time progresses, and people learn more about each other, revolt never decreases: revolt either increases or remains constant.

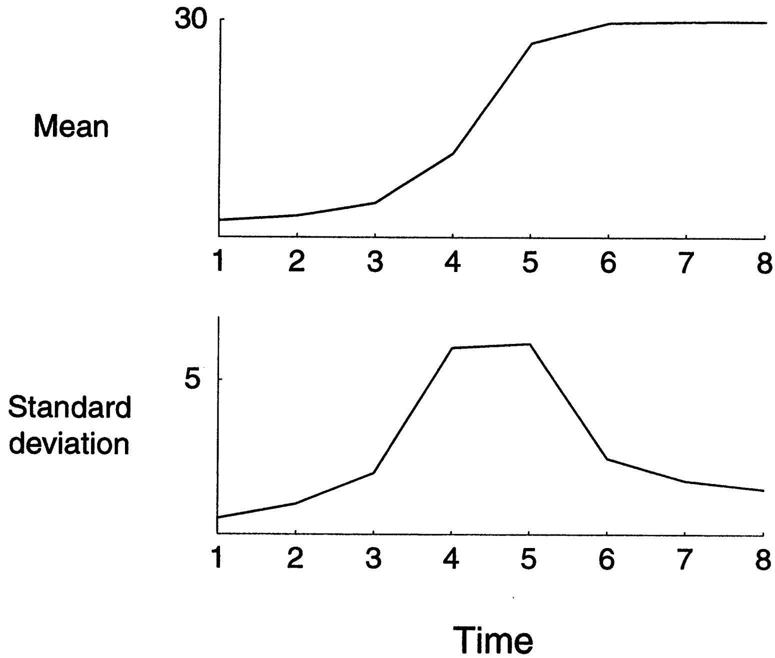

FIG. 5.—Mean and standard deviation of number of people revolting over time

THRESHOLDS AND NETWORK POSITION

Since our model incorporates both network and threshold heterogeneity, it is natural to look at their combined effects. Consider networks of 30 people, in which each person has two neighbors selected at random; these are "unbiased" networks since any person is equally likely to be a given person's neighbor. Say that there are two people with threshold 1, two with threshold 2, and so on up to 15. We randomly generate 1,000 of these networks and find the average and standard deviation of the number of people revolting at a given time, as shown in figure 5.

The mean number of people revolting grows slowly at first, increases quickly, and finally slows down again as in the classic "logistic" curve. (Of course, in general in our model, depending on the thresholds and the network, the growth of revolt can be almost any curve, as shown in later examples; if growth is not logistic, one cannot conclude, as does, for example, Burt [1987], that the network has no effect.) By $t = 6$, in 958 of our 1,000 networks, all 30 people have revolted; by $t = 8$, in 994 of our 1,000 networks, all 30 have revolted. The standard deviation is highest in intermediate periods and can be quite substantial: for example, at $t = 4$, the standard deviation is 6.0, more than half of the mean 11.6. This variance comes entirely from the randomness in network positions, since the distribution of thresholds is always the same; in other words, network variation alone can have a large effect on aggregate behavior.

One might expect that the people who revolt earliest are the ones with low thresholds. Figure 6 shows the threshold distribution of people who first revolt at time $t = 1$, $t = 2$, and so on; recall that there are two people for each threshold from 1 to 15. At $t = 1$, the two threshold 1 people both revolt, and occasionally a threshold 2 person revolts. At $t = 2$, the people who revolt are mainly threshold 2 and threshold 3 people; similarly, at time $t = 3$, the only people who revolt are those with low thresholds. At $t = 4$, some high-threshold people revolt, but the threshold distribution is still slightly skewed toward lower thresholds. At $t = 5$ and $t = 6$, the threshold distribution is skewed toward higher thresholds but is still fairly "flat." So it is true that people who revolt early tend to have low thresholds, but it is also true that people who revolt late do not necessarily have high thresholds. For example, when $t = 5$ and most of the people revolting have high thresholds, there are still many low-threshold people who are just starting to revolt.

Another way to see this is to plot the mean and standard deviation of the time at which a person of a given threshold starts to revolt, as in figure 7 (the rare cases in which a person does not revolt by $t = 8$ are not included in this calculation). People with higher thresholds indeed revolt later on average. However, except for people with threshold 1, the revolt time of people with low thresholds has higher variance than those with high thresholds. In other words, the revolt time of people with low thresholds, except of course people with threshold 1, are influenced much more by their network position. This makes sense: for people with high thresholds, all that matters for revolt is being connected to a large mass of people, but for people with low thresholds, whether you are lucky enough to be close to another person with a low threshold makes a big difference.

Empirical studies often make the assumption, implicit in linear regression, that all independent variables (age, social class, ideology, social position) enter linearly into participation. In other words, one's social position affects one's participation equally, regardless of whether one is very disposed to participate (has a low threshold) or is not disposed to participate (has a high threshold). This example suggests that this might not be justified: people with low thresholds would be affected by social position much more than people with high thresholds. In fact, perhaps another empirical strategy for verifying if social position is a significant influence on participation is to check whether the adoption time of people who seem to have low thresholds has greater variance than the adoption time of people who seem to have high thresholds.

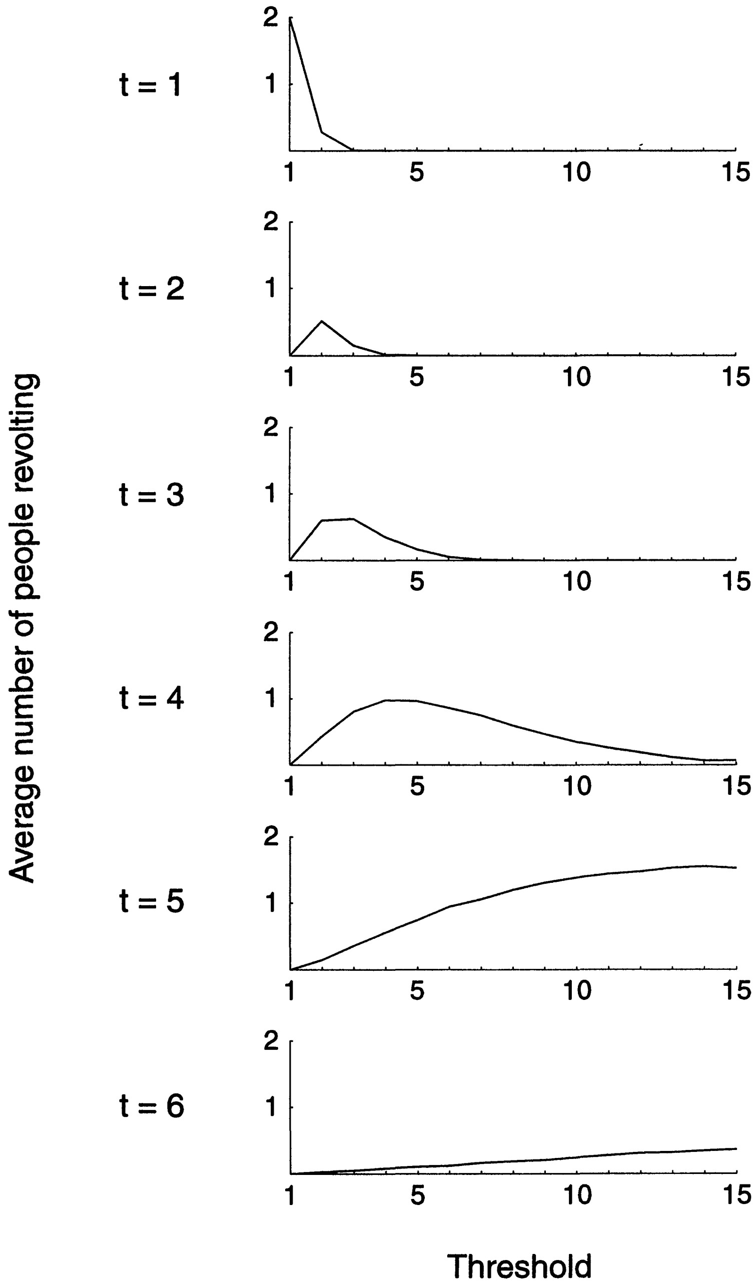

FIG. 6.—Threshold distribution of people who start to revolt at a given time

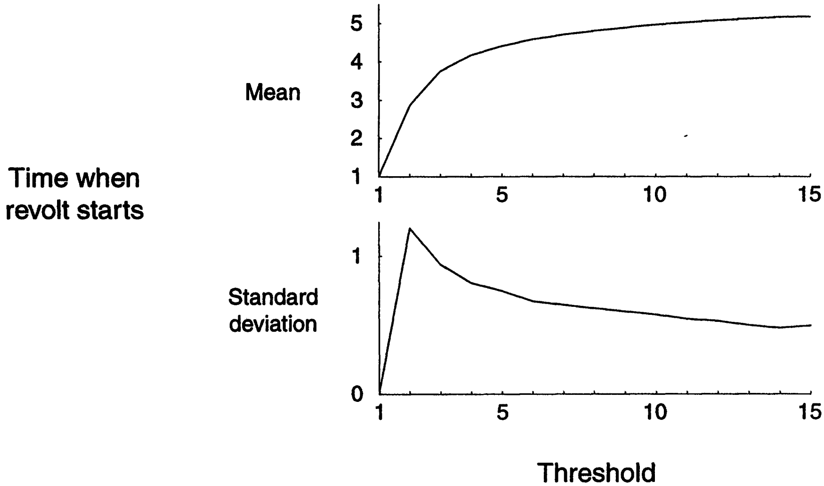

FIG. 7.—Mean and standard deviation of time period in which a person of a given threshold starts to revolt.

STRONG AND WEAK LINKS

The distinction between "strong" and "weak" links is an early insight of social network theory (Granovetter 1973). Roughly speaking, a strong link joins close friends and a weak link joins acquaintances. A simple empirical finding (Rapoport and Horvath 1961; see also White 1992) is that strong links tend to traverse a society "slowly": start with an arbitrary person, find two of her close friends, then find two close friends of each of these two people, and continue in this manner. As you iterate, the group increases slowly because often no one new is added: the close friends of my close friends tend to be my close friends also. If instead you successively add two acquaintances, the group grows quickly: the acquaintances of my acquaintances tend not to be my acquaintances. Weak links traverse a society "quickly": a demonstration suggests that any two people in the United States can be connected by as few as six weak links (Milgram 1992; see also Kochen 1989). Weak links tend to scatter widely, while strong links tend to be involuted.

To connect a large society, then, weak links are more important than strong links; weak links are more important for spreading information (Granovetter 1995; see also Montgomery 1991). For collective action, however, the importance of strong versus weak links is unclear. Data from volunteers in the 1964 Mississippi Freedom Summer, for example, show that the presence of a strong link to another potential participant correlates strongly and positively with participation, while the presence of a weak link has no correlation (McAdam 1986; McAdam and Paulsen 1993; see also Fernandez and McAdam 1988). In three classic "diffusion" studies, rates of adoption are actually negatively correlated with the presence of weak links (Valente 1995, p. 51).

Our model shows how strong links can be better for participation when thresholds are low and weak links can be better when thresholds are high. This is because the involutedness of strong links, their tendency to form small cliques, is exactly what is needed to form common knowledge at a local level; if my friend's friend is my friend also, then common knowledge among the three of us is formed quickly. If thresholds are high, common knowledge must be formed among a large group of people; then, weak links are better simply because they speed up communication.

Again, consider networks of 30 people, each person having three neighbors. To randomly generate strong and weak networks, we use the following simple procedure (similar to that of Hedström [1994]; see, e.g., Fararo and Skvoretz [1987] and Frank and Strauss [1986] on biased networks in general and specifically Skvoretz [1990] on the difficulties of randomly generating biased networks). We first give each person a location, selected randomly and uniformly, on a unit square. A person's neighbors are selected in the following manner: with probability p, a "bias event" occurs and a neighbor is selected randomly among the closest three neighbors; with probability $1 - p$, a neighbor is selected randomly from the entire population. Hence, if $p = 1$, then each person's neighbors are simply the three closest people; if $p = 0$, then we have an unbiased random graph as discussed in the previous section.

Figure 8 shows some examples. When $p = 1$, the network is composed completely of strong links between geographically close people. When $p = 0.8$, on average 80% of the links are strong and 20% are weak links that scatter widely. When $p = 0.6$, there are more weak links, and when $p = 0$, locality is irrelevant and all links are weak.

For each of these four values of our "bias parameter" p, we generate 200 random networks. One way to check that these networks approximate what is normally considered strong versus weak is to plot a "tracing" for each value of p (Rapoport and Horvath 1961), as shown in figure 9. Figure 9 shows the number of people within a certain number of links (averaged over all 30 people in a graph and over all 200 networks), for each of the four values of p. For example, for all p, starting from a given person, there are exactly four people within one link (the person herself and her three neighbors). The average number of people within two links of a given person (the person herself, her neighbors, and her neighbors' neighbors) is 11.3 for the $p = 0$ graphs and 6.7 for the $p = 1$ graphs. The lower p is, that is, the more weak links in a network, the faster the tracing increases; weak-link networks are better at connecting large groups of people.

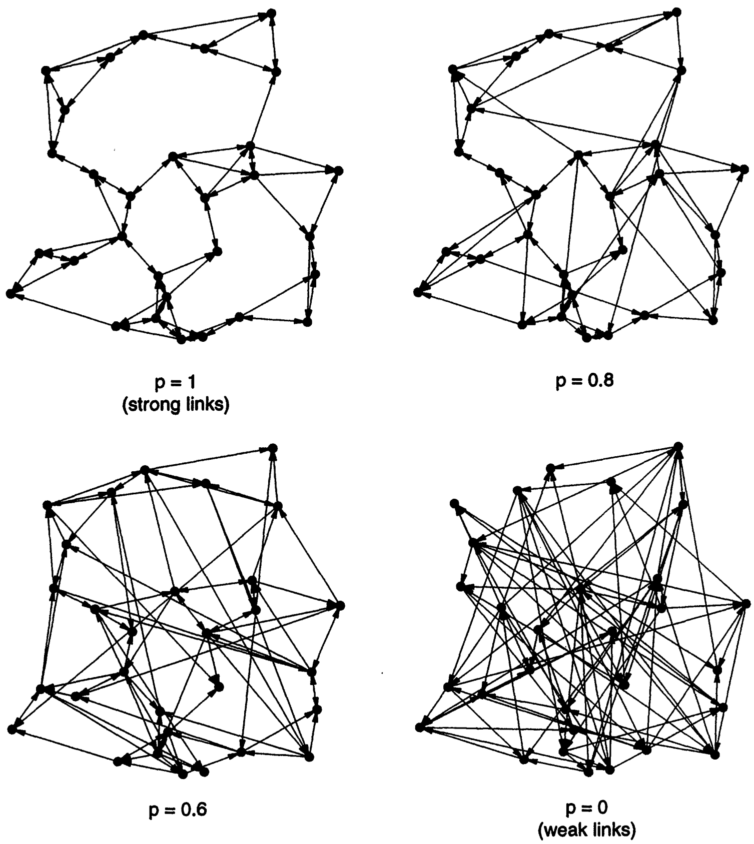

FIG. 8.—Strong and weak links

Another way to see that our choices of p correspond to strong versus weak is to calculate transitivity measures. Strong-link networks are supposed to be more transitive: in a strong-link network, it is more likely that a friend of a friend is my friend also. The simplest transitivity measure is just the number of transitive triads: the number of triples (i, j, k) (where i, j, k are all distinct) such that $i \rightarrow j$, $j \rightarrow k$, and $i \rightarrow k$. For the $p = 0$ weak-link networks, the number of transitive triads has mean 18.5 and standard deviation 4.3, and we can use this as a benchmark. The $p = 0.6$ networks have on average 36.0 transitive triads, roughly 4 standard deviations away from the benchmark, the $p = 0.8$ networks have on average 61.5 transitive triads, roughly 10 standard deviations away from the benchmark, and the $p = 1$ networks have on average 105.4 transitive triads, roughly 20 standard deviations away from the benchmark. Our

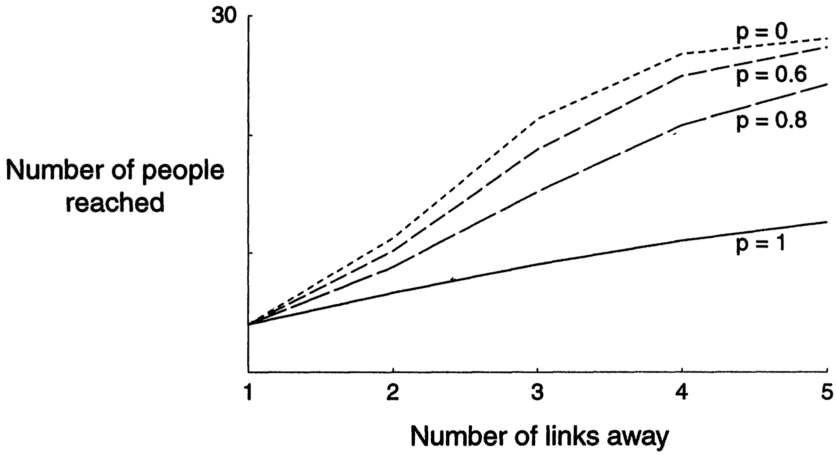

FIG. 9.—Tracings for $p = 1$, $p = 0.8$, $p = 0.6$, and $p = 0$

strong-link networks are real "outliers" when compared to the benchmark set of unbiased networks, but in this way are not completely unrealistic. For example, in their survey of 408 actual social networks, Holland and Leinhardt (1975) find the average number of transitive triples to be more than 5 standard deviations higher than what would be expected assuming a random "unbiased" distribution.

Assume that everyone has the same threshold θ. When everyone has the same threshold, it is possible to show (Chwe 1999) that a person revolts at time t if and only if his neighborhood contains a t-clique of size θ (a t-clique is a set of people in which each person is within t links of every other). In other words, when everyone has the same threshold θ, revolt is simply a matter of forming t-cliques of size θ. Figure 10 shows the average number of people who revolt over time for each choice of p and for θ ranging from 3 to 7. For example, in the first graph, everyone has threshold 3; here, the strong-link networks are very advantageous for revolt in the first time period. This is simply because the strong-link networks have many more cliques of size 3, as shown in figure 8. When everyone has threshold 4, the results are similar: strong links are advantageous (now in the second period); again, the transitivity or "involutedness" of strong-link networks helps form small cliques quickly.

When everyone has threshold 5, strong-link networks are better early on but worse in the long run. Here, strong-link networks form some cliques of size 5 quickly, but in the long run, the weak-link networks are better because of their greater "reach," as illustrated in the tracings in

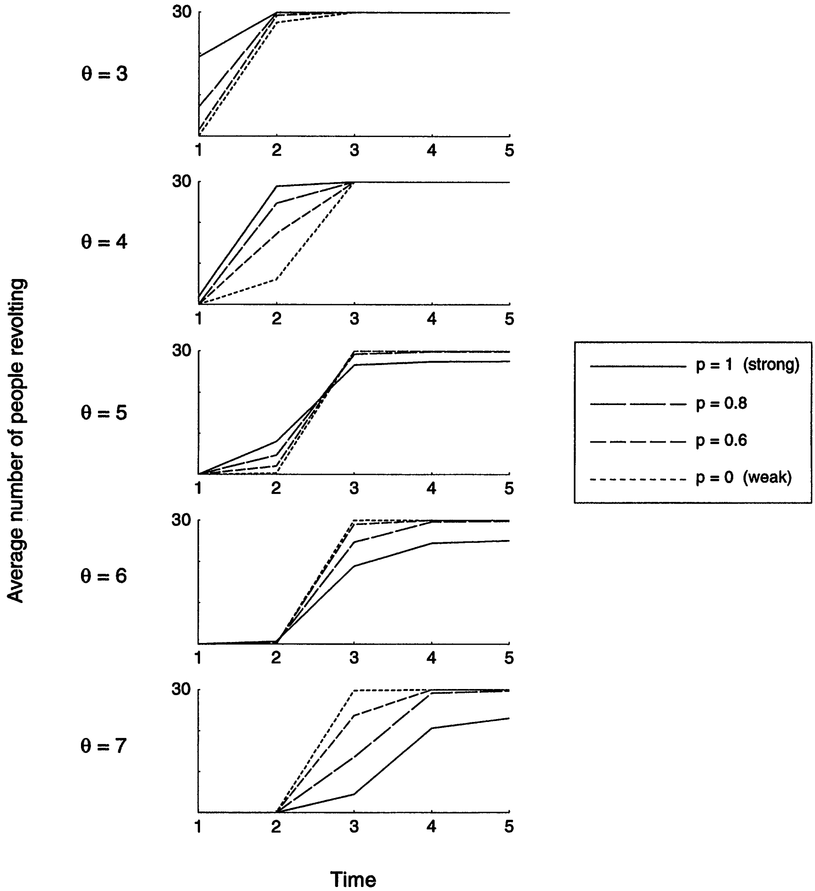

FIG. 10.—The number of people revolting when everyone has the same threshold θ, where $\theta = 3, 4, 5, 6, 7$.

figure 9. When everyone has threshold 6 and everyone has threshold 7, weak-link networks start to clearly dominate.

That strong as opposed to weak links are important in recruiting for Freedom Summer is interpreted by McAdam (1986, p. 80) as the links having different functions: "although weak links may be more effective as diffusion channels, strong ties embody greater potential for influencing behavior." This is of course reasonable, but our model suggests it is not necessary. In Freedom Summer, a person might participate if a few friends also participate, which is like having a low threshold (strictly speaking, in our model only the number of fellow participants matters, not whether they are close friends). If you and I are potential participants connected by a strong link, your friends are likely to be my friends, and participation among our group of friends would be common knowledge among us. If you and I are connected by a weak link, your friends and my friends do not know each other, and hence there is no common-knowledge group to which we both belong. In other words, the idea that weak links are always better for communication relies on the assumption that communication is about knowledge only and not "metaknowledge," knowledge of what others know. Strong links are better for forming common knowledge at a local level, and when thresholds are low, local mobilization is all that is necessary (see also Marwell, Oliver, and Prahl 1988, p. 532).

To summarize, the "structural" question of whether strong links or weak links are better for collective action cannot be answered without considering the "strategic" situation, the distribution of thresholds. Structural and strategic considerations interact in interesting ways.

LIMITED COMMUNICATION AND THE FRAGILITY OF COLLECTIVE ACTION

One of the earliest and simplest explicit models of a coordination process was introduced by Schelling (1978): first, only people with low thresholds participate, but their participation makes people with slightly higher thresholds want to participate. As the number participating grows, people join in successively, in a "snowball" or "bandwagon" effect. I call this the "simple model."

This simple model can be represented as a special case of our model; our model might thus be considered its generalization to arbitrary networks. Doing this shows that the simple model implicitly assumes that communication possibilities are extremely restricted: a person only gets information about people who have lower thresholds, making bilateral communication impossible. One compelling result of the simple model is that collective action depends crucially on the thresholds of people "early" in its growth (Granovetter 1978). I show that this fragility depends heavily on the assumption of extremely restricted communication.

To take an example, say that there are five people: one with threshold 1, one with threshold 2, and so on up to threshold 5. The simple model goes like this: first, the threshold 1 person willingly revolts by himself. When the threshold 2 person sees one person revolting, that is enough to make him want to revolt, and he joins in. Then the threshold 3 person joins in, and so forth until everyone joins in. Note that in this process, there is no explicit communication and indeed no need to learn anyone else's threshold: each person makes their own independent decision given the number of people already revolting.

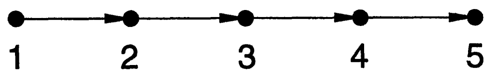

FIG. 11.—"Bandwagon"

To put this in our framework, one simply considers the network shown in figure 11, where the numbers 1, 2, 3, 4, 5 indicate each person's threshold. Our model has quite different premises: people do not observe each other's actions but learn their thresholds. Still, the two models give the same conclusion. In the first period, the threshold 1 person naturally revolts and the threshold 2 person revolts because he knows that there is a threshold 1 person. In the second period, the threshold 3 person revolts because he learns that there is a threshold 1 person and a threshold 2 person, and so on.

Note that the communication network that adapts our model to the simple model is very sparse; in fact, it could not be made any more sparse without disconnecting the group. Communication is ample in that a person eventually learns about everyone with a lower threshold (Braun 1995), but restricted in that a person never learns about anyone with a higher threshold; communication is in only one direction, never reciprocal.

The effects of this sparse network are clear in an example. Start with the same network and the same five people. Now change the first person's threshold from 1 to 2 so thresholds are now 2, 2, 3, 4, 5. Now, no one ever revolts. The first person no longer wants to jump in all by himself, and hence the second person does not follow. Participation is vulnerable to small perturbations of people "early" in the bandwagon.

However, if the first and the second person could talk reciprocally, as in the network in figure 12, the two would jump in together and the bandwagon would go on as before. In other words, the breakdown of revolt when going from 1, 2, 3, 4, 5 to 2, 2, 3, 4, 5 depends entirely on the assumption that communication is extremely limited, never reciprocal.

To illustrate this more generally, again consider networks of 30 people in which each person has three neighbors. Again, say two people have threshold 1, two people have threshold 2, and so on up to 15. To construct the social network, we first place the people in a line, ordered by their threshold, like this: 1, 1, 2, 2, 3, 3, . . . , 15, 15. A person's neighbors are selected in the following manner: with probability q, a "bias" or "skew"

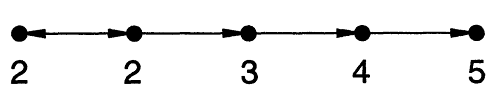

FIG. 12.—If the network is less sparse, the bandwagon is less fragile

event occurs and a neighbor is selected randomly from a person to the left (if there is no one to the left, no one is selected); with probability $1 - q$, a neighbor is selected randomly from the entire population. Hence, if $q = 1$, communication only flows in one direction, from left to right, as in figure 11. If $q = 0$, we have an unbiased random graph, as discussed earlier.

Say we set $q = 1$ and generate a "maximally skewed" random graph in this manner. Revolt then starts with the two threshold 1 people and then spreads to people with higher thresholds. However, if we perturb the situation by letting the first two people have threshold 2 instead of threshold 1 (in other words, instead of thresholds being 1, 1, 2, 2, 3, 3, . . . , 15, 15, they are 2, 2, 2, 2, 3, 3, . . . , 15, 15), revolt collapses completely, as in the simple model. Again, this is because when communication is never reciprocal, threshold 1 people are necessary to make any revolt at all possible.

What happens as we move q away from 1, so our networks are not completely skewed? Let $q = 1, 0.9, 0.8, 0.7, 0.6, 0$, and randomly generate 200 sample graphs for each value of q. Figure 13 shows the average number of people revolting over time, where the black line corresponds to the original thresholds 1, 1, 2, 2, 3, 3, . . . , 15, 15 and the gray line corresponds to the perturbed thresholds 2, 2, 2, 2, 3, 3, . . . , 15, 15. Note that for all values of q, revolt in the first period is much greater with the original than with the perturbed thresholds; this is simply because the original thresholds include two people with threshold 1. When $q = 1$, with the perturbed thresholds, revolt completely collapses. When $q = 0.9$, however, by the third period, revolt with the perturbed thresholds is half or more of that of the original thresholds. As q decreases from $q = 0.8$, the perturbation quickly makes relatively little difference. Another interesting thing to note is that there is an "optimal" level of q; for both the original and perturbed thresholds, revolt increases faster when $q = 0.6$, for example, than when either $q = 0$ or $q = 1$. In other words, there is an advantage to having information flow from people with lower thresholds to people with higher thresholds, but too much skew is bad because it reduces reciprocal communication.

To summarize, the simple model greatly overestimates the fragility of collective action; even when only 10% or 20% of the links are random and not "skewed," revolt does not collapse with a small perturbation of thresholds. Putting it another way, the less skewed the network and the more possibilities for reciprocal communication, the less fragile is revolt.

One can defend the simple model as applying to a spontaneous, unorganized process, explicit communication being unnecessary or unavailable. Communication is not necessarily automatic or easy, but it is hard to believe that a few words ("I'll do it if you do it") or even eye contact between

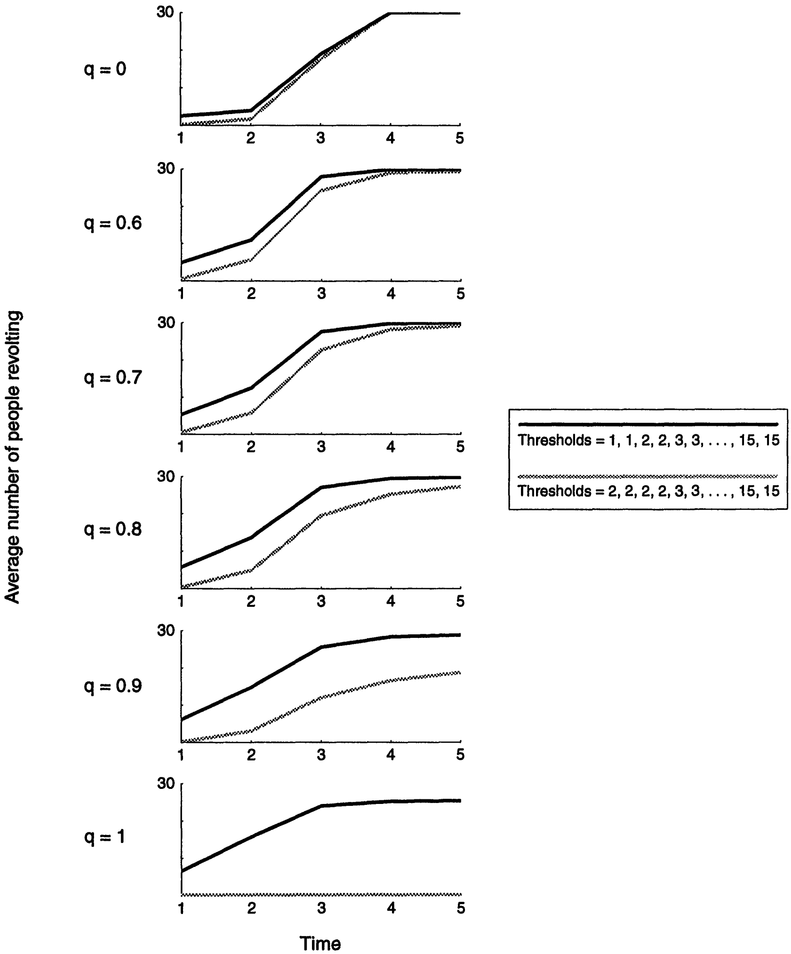

FIG. 13.—If the network is less skewed, the bandwagon is less fragile

two people at least could not arise spontaneously. "Research on participation in a variety of prosaic gatherings, in political, religious, and sports demonstrations, and probably in riots, establishes that most participants are neither alone or anonymous; rather they assemble with or soon encounter family members, friends, or acquaintances" (McPhail 1991, p. 14). In any case, the simple model does not approximate the "general" case; when coordination takes place under the simple model, it takes place under communicative conditions so limited that no pair of individuals can reciprocally communicate.

CONCLUDING REMARKS

The broad message of this article is that there is no necessary opposition between structure and rationality; our model incorporates both explicit network structure and complete rationality and derives results not obtainable from either perspective alone. This supposed opposition, which has a long and undoubtedly complex historical basis, is evident in current research: almost all models of the effects of social structure on individual action make adaptive, bounded rationality, or behavioral assumptions. This is true not only in sociology (e.g., Gould 1993; Macy 1991, 1995; but see also Marwell and Oliver 1993) but also in recent work in economics (e.g., Akerlof 1997; Anderlini and Ianni 1996; Blume 1993, 1995; Ellison 1993; Goyal 1996; Mailath, Samuelson, and Shaked 1997; Morris 1999; Temzelides 1997; Tesfatsion 1997; Young 1996, 1998). Sometimes this supposed opposition is drawn very sharply; for example, in his discussion of the Schelling-Granovetter threshold model, Michael Macy (1991, p. 735) finds that "structural analysis is badly handicapped by Granovetter's rational choice assumption . . . that individual participation is triggered by the contribution rate of the group as a whole. The choices of those in one's immediate circle have no special relevance. . . . The structure of social ties is much more relevant if the actions of others shape individual behavior directly rather than as a by-product of an underlying interest in marginal utility." Here, adding structure seems to require the immediate jettisoning of rationality, as opposed to any other assumption (see, e.g., Krassa 1988; Valente 1996), and seems to lead directly and unarguably toward a model in which people do not make choices but rather are influenced directly by others.

One advantage of explicitly modeling individual choices, as opposed to letting behavior directly affect behavior, is that one is led to recognize the wide variety of strategic situations in which networks play a role and to develop models that are sensitive to this variety. For example, Hedström (1994, p. 1162) develops a model in which a person's probability of participating is a linear function, weighted by social distance, of the participation of others and readily admits that, for the sake of simplicity, the model "assumes that the decision to join an organization for collective action resembles the decision whether or not to buy a private good." But surely these are quite different decisions; the entire theory of public goods and the "free-rider problem," for example, depends on this distinction. Instead of effacing this distinction, a network model could more interestingly interact with it: networks that are good for the diffusion of private goods purchases might be quite bad for the spread of collective political action, for example.

Similarly, what exactly people do with the network can also be made explicit. It is commonly argued that a social network is used to communicate information, but information is almost never modeled explicitly (the "state space" model, used here, is one of several ways to model it). Also, it is often unclear what the information is about. As in the Schelling-Granovetter model, a person usually is assumed to react to the actions of her neighbors in the network. In my model, a person learns about her neighbors' preferences, their willingness to revolt, and does not directly respond to their actions. This distinction is not merely conceptual but can be found in "real life," for example, in Dingxin Zhao's (1998) analysis of the April 27, 1989, student demonstration in Beijing. A crucial "ecological" factor for participation was that students lived close together in university dormitories, allowing them to discuss and "share their anger." This communication process, which enabled students to know that enough others felt the same way, is what my model tries to describe. Once students marched on the street, however, students on bicycles rode back and forth between separate groups, thereby informing each group that other groups were already marching, and there was a "snowball" effect as in the Schelling-Granovetter model.

One specific contribution of my model is that it allows a person's choice to depend on the relationships among his neighbors. I make the strong assumption that a person knows whether two of his neighbors know each other, and as illustrated in the many examples above, what a person knows about the knowledge of other people is crucial. The extent to which this assumption is empirically valid is not yet settled: some empirical researchers (notably Bernard, Kilworth, and Seiler 1980, 1982) find that a person's knowledge of specific instances of social interactions among the people around her is typically quite poor, while others (e.g., Freeman, Romney, and Freeman 1987; Freeman, Freeman, and Michaelson 1988; see also Romney and Faust 1982) find that a person is typically fairly good at understanding the long-term patterns of interaction around her. Presumably, one reason this question is interesting is because it might affect individual and thus social behavior. However, in almost all existing models, including models in which a person's action is a linear function of the actions of his neighbors (e.g., Gould 1993; Hedström 1994), learning models (e.g., Macy 1991), and models in which people can explicitly organize each other and sign all-or-nothing contracts (Marwell and Oliver 1993), the question is irrelevant. In all of these models, a person's action depends only on her neighbors and their actions and not on the connections among her neighbors; connections among neighbors do affect the long-run evolution of society-wide behavior but do not affect an individual's action at a given moment. Whether the two friends I get information from know each other or not should have an effect on my decision. Although this effect depends on me knowing whether my two friends know each other, it is also unwarranted to a priori rule out this effect completely, as do existing models.

In this way, existing models are somewhat limited in looking at network transitivity; for example, cliques are often found to be bad for collective action, despite the intuition that collective action starts in small "subcultures," which are commonly thought of and "detected" as cliques. To take another example, there is some evidence that people's beliefs are systematically biased toward transitivity: people tend to think that the networks they belong to are more transitive than they really are (see, e.g., Freeman 1992). To explore the implications of this bias for individual and social action, one needs a model in which transitivity directly matters. Finally, there is even an immediate implication for survey methods: when Opp and Gern (1993), for example, surveyed participants in the demonstrations that led to the collapse of East Germany, they asked each person whether he had friends who participated and found that this was a significant variable in predicting his participation. Our model suggests that each person should also be asked if his friends who participated knew each other. If this variable is significant (of course after controlling for spurious correlations), then this would be evidence for a communication model as in this article.

In sum, structure and strategy can work together in various interesting and unexpected ways: for example, the game-theoretic model here, in which knowledge of other people's knowledge is crucial, allows transitivity to affect individual action more directly than bounded rationality or behavioral models seemingly more congenial to structural considerations. The structural approach and the strategic approach to collective action can surely find more insights together than separately.

APPENDIX

Here I specify explicitly people's utilities, knowledge, and strategies, and then consider equilibria. First, define utility: given person i's own threshold θ_i and everyone's actions $a = (a_1, \ldots, a_n)$, her utility $u_i(\theta_i, a_1, \ldots, a_n)$ is given by

$$u_i(\theta_i, a_1, \ldots, a_n) = \begin{cases} 0 & \text{if } a_i = s \\ 1 & \text{if } a_i = r \text{ and } \#\{j \in N : a_j = r\} \geq \theta_i \\ -z & \text{if } a_i = r \text{ and } \#\{j \in N : a_j = r\} < \theta_i, \end{cases}$$

where $-z < -(n + 1)^n$. In other words, a person always gets utility 0 by staying at home. If he revolts, he gets utility 1 if the total number of people revolting is at least θ_i, his threshold. If he revolts and not enough people join him, he gets the very large penalty $-z$.

I represent person i's knowledge as a partition of the set of all possible states of the world. Since the uncertainty here is over everyone's thresholds, the set of all possible states of the world is $\Theta = \{1, 2, 3, \ldots, n + 1\}^n$. For simplicity, I assume that thresholds are independently and identically distributed, and that each threshold is equally likely: thus the probability distribution $\pi{:}\Theta \rightarrow [0,1]$ over Θ is given by $\pi(\theta) = 1/(n + 1)^n$. Since people are also assumed to have incomplete information about the network (person i only knows the network structure within his neighborhood $B(i)$), I should also explicitly model uncertainty in the network. However, given our other particular assumptions, doing this would not change any of the results.

Person i knows the thresholds of people in the neighborhood $B(i)$. Hence, if the actual state of the world is $\theta \in \Theta$, then person i knows only that the state of the world is in the set $P_i(\theta) = \{(\theta_{B(i)}, \phi_{N-B(i)}){:}\phi_{N-B(i)} \in \{1, 2, 3, \ldots, n + 1\}^{n-\#B(i)}\}$ (we use the notation $\theta_A = (\theta_j)_{j\in A}$). Taken together, the sets $\{P_i(\theta)\}_{\theta\in\Theta}$ form a partition of Θ, which we call $\mathcal{P}_i$.

Person i's strategy is a choice of action given her knowledge of the state of the world. Thus, I define a strategy for person i to be a function f_i: $\Theta \rightarrow \{r, s\}$ which is measurable with respect to $\mathcal{P}_i$, that is, for all $\theta, \theta' \in \Theta$, if $\theta, \theta' \in P$, where $P \in \mathcal{P}_i$, then $f_i(\theta) = f_i(\theta')$. The idea here is that if θ and θ' are in the same element of the partition $\mathcal{P}_i$, then person i cannot distinguish between the two states θ and θ', and hence must take the same action in the two states. Say that F_i is the set of all strategies of person i.

Given strategies $(f_1, \ldots, f_n)$, person i's expected utility is $EU_i(f_1, \ldots, f_n) = \sum_{\theta\in\Theta} \pi(\theta)u_i(\theta_i, f_1(\theta), \ldots, f_n(\theta))$. We say that $(f_1, \ldots, f_n)$ is an equilibrium if for all $i \in N$, and for all $g_i \in F_i$, $EU_i(f_1, \ldots, f_n) \geq EU_i(f_1, \ldots, f_{i-1}, g_i, f_{i+1}, \ldots, f_n)$. In other words, $(f_1, \ldots, f_n)$ is an equilibrium if no individual can gain by deviating to another strategy $g_i \in F_i$.

In general, there exist many equilibria. It is not hard to show that there uniquely exists an equilibrium $(f_1^*, \ldots, f_n^*)$ such that if $(f_1, \ldots, f_n)$ is an equilibrium, then $f_i(\theta) = r \Rightarrow f_i^*(\theta) = r$. In other words, the equilibrium $(f_1^*, \ldots, f_n^*)$ has the greatest possible revolt: if person i revolts in any equilibrium, she revolts in this "best" equilibrium. We assume that this equilibrium is the one which is played. Given the relation $\rightarrow$ and thresholds θ, the set of people who revolt is $\{i \in N{:}\, f_i^*(\theta) = r\}$.

REFERENCES

Akerlof, George A. 1997. "Social Distance and Social Decisions." *Econometrica* 65: 1005–27.

Allport, Floyd H. 1924. *Social Psychology*. Boston: Houghton Mifflin.

Anderlini, Luca, and Antonella Ianni. 1996. "Path Dependence and Learning from Neighbors." *Games and Economic Behavior* 13:141–77.

Aumann, Robert J. 1976. "Agreeing to Disagree." *Annals of Statistics* 4:1236–39.

Bernard, H. Russell, Peter D. Kilworth, and Lee Sailer. 1980. "Informant Accuracy in Social-Network Data IV: A Comparison of Clique-Level Structure in Behavioral and Cognitive Network Data." *Social Networks* 2:191–218.

———. 1982. "Informant Accuracy in Social-Network Data V. An Experimental Attempt to Predict Actual Communication from Recall Data." *Social Science Research* 11:30–66.

Blume, Larry. 1993. "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 5:387–424.

———. 1995. "The Statistical Mechanics of Best-Response Strategy Revision." *Games and Economic Behavior* 11:111–45.

Braun, Norman. 1995. "Individual Thresholds and Social Diffusion." *Rationality and Society* 7:167–82.

Burt, Ronald S. 1987. "Social Contagion and Innovation: Cohesion versus Structural Equivalence." *American Journal of Sociology* 92:1287–335.

Chwe, Michael Suk-Young. 1998. "Culture, Circles, and Commercials: Publicity, Common Knowledge, and Social Coordination." *Rationality and Society* 10:47–75.

———. 1999. "Communication and Coordination in Social Networks." *Review of Economic Studies,* in press.

Coser, Lewis A. 1990. "The Intellectuals in Soviet Reform: On 'Pluralistic Ignorance' and Mass Communications." *Dissent* spring:181–83.

Ellison, Glenn. 1993. "Learning, Local Interaction, and Coordination." *Econometrica* 61:1047–71.

Fararo, Thomas J., and John Skvoretz. 1987. "Unification Research Programs: Integrating Two Structural Theories." *American Journal of Sociology* 92:1183–1209.

Fernandez, Roberto M., and Doug McAdam. 1988. "Social Networks and Social Movements: Multiorganizational Fields and Recruitment to Freedom Summer." *Sociological Forum* 3:357–82.

Frank, Ove, and David Strauss. 1986. "Markov Graphs." *Journal of the American Statistical Association* 81:832–42.

Freeman, Linton C. 1992. "Filling in the Blanks: A Theory of Cognitive Categories and the Structure of Social Affiliation." *Social Psychology Quarterly* 55:118–27.

Freeman, Linton C., Sue C. Freeman, and Alaina G. Michaelson. 1988. "On Human Social Intelligence." *Journal of Social and Biological Structures* 11:415–25.

Freeman, Linton C., A. Kimball Romney, and Sue C. Freeman. 1987. "Cognitive Structure and Informant Accuracy." *American Anthropologist* 89:310–25.

Gould, Roger V. 1993. "Collective Action and Network Structure." *American Sociological Review* 58:182–96.

———. 1995. *Insurgent Identities: Class, Community, and Protest in Paris from 1848 to the Commune.* Chicago: University of Chicago Press.

Goyal, Sanjeev. 1996. "Interaction Structure and Social Change." *Journal of Institutional and Theoretical Economics* 152:472–94.

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

———. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83:1420–43.

———. 1995. *Getting a Job: A Study of Contacts and Careers,* 2d ed. Chicago: University of Chicago Press.

Hedström, Peter. 1994. "Contagious Collectivities: On the Spatial Diffusion of Swedish Trade Unions, 1890–1940." *American Journal of Sociology* 99:1157–79.

Holland, P. W., and S. Leinhardt. 1975. "The Statistical Analysis of Local Structure in Social Networks." Pp. 1–45 in *Sociological Methodology 1976,* edited by D. R. Heise. San Francisco: Jossey-Bass.
Katz, Daniel, and Floyd H. Allport. 1931. *Student Attitudes.* Syracuse, N.Y.: Craftsman Press.
Kochen, Manfred. 1989. *The Small World.* Norwood, N.J.: Ablex Publishing.
Krassa, Michael A. 1988. "Social Groups, Selective Perception, and Behavioral Contagion in Public Opinion." *Social Networks* 10:109–36.
Kuran, Timur. 1991. "Now Out of Never: The Element of Surprise in the East European Revolution of 1989." *World Politics* 44:7–48.
Lewis, David K. 1969. *Convention: A Philosophical Study.* Cambridge, Mass.: Harvard University Press.
Macy, Michael W. 1991. "Chains of Cooperation: Threshold Effects in Collective Action." *American Sociological Review* 56:730–47.
———. 1995. "Beyond Rationality in Models of Choice." *Annual Review of Sociology* 21:73–91.
Mailath, George J., Larry Samuelson, and Avner Shaked. 1997. "Correlated Equilibria and Local Interactions." *Economic Theory* 9:551–56.
Marwell, Gerald, and Pamela Oliver. 1993. *The Critical Mass in Collective Action.* Cambridge: Cambridge University Press.
Marwell, Gerald, Pamela E. Oliver, and Ralph Prahl. 1988. "Social Networks and Collective Action: A Theory of the Critical Mass. III." *American Journal of Sociology* 94:502–34.
McAdam, Doug. 1986. "Recruitment to High-Risk Activism: The Case of Freedom Summer." *American Journal of Sociology* 92:64–90.
McAdam, Doug, and Ronnelle Paulsen. 1993. "Specifying the Relationship between Social Ties and Activism." *American Journal of Sociology* 99:640–67.
McPhail, Clark. 1991. *The Myth of the Madding Crowd.* New York: Aldine de Gruyter.
Milgram, Stanley. 1992. *The Individual in a Social World: Essays and Experiments,* 2d ed. Edited by John Sabini and Maury Silver. New York: McGraw Hill.
Montgomery, James D. 1991. "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis." *American Economic Review* 81:1408–18.
Morris, Stephen. 1999. "Contagion." *Review of Economic Studies,* in press.
Mullen, Brian, and Li-tze Hu. 1988. "Social Projection as a Function of Cognitive Mechanisms: Two Meta-Analytic Integrations." *British Journal of Social Psychology* 27:333–56.
O'Gorman, Hubert J. 1979. "White and Black Perceptions of Racial Values." *Public Opinion Quarterly* 43:48–59.
———. 1986. "The Discovery of Pluralistic Ignorance: An Ironic Lesson." *Journal of the History of the Behavioral Sciences* 22:333–47.
Opp, Karl-Dieter, and Christiane Gern. 1993. "Dissident Groups, Personal Networks, and Spontaneous Cooperation." *American Sociological Review* 58:659–80.
Rapoport, Anatol, and W. J. Horvath. 1961. "A Study of a Large Sociogram." *Behavioral Science* 6:279–91.
Romney, A. Kimball, and Katherine Faust. 1982. "Predicting the Structure of a Communications Network from Recalled Data." *Social Networks* 4:285–304.
Rubinstein, Ariel. 1989. "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge.' " *American Economic Review* 79:385–91.
Schelling, Thomas C. 1978. *Micromotives and Macrobehavior.* New York: Norton.
Scott, James C. 1990. *Domination and the Arts of Resistance: Hidden Transcripts.* New Haven, Conn.: Yale University Press.
Shamir, Jacob. 1993. "Pluralistic Ignorance Revisited: Perception of Opinion Distributions in Israel." *International Journal of Public Opinion Research* 5:22–39.

Skvoretz, John. 1990. "Biased Net Theory: Approximations, Simulations, and Observations." *Social Networks* 12:217–38.

Temzelides, Ted. 1997. "Evolution, Coordination, and Banking Panics." *Journal of Monetary Economics*.

Tesfatsion, Leigh. 1997. "How Economists Can Get a Life." Pp. 533–64 in *The Economy as an Evolving Complex System II,* edited by W. Brian Arthur, Steven N. Durlauf, and David A. Lane. Redwood City, Calif.: Addison-Wesley Longman.

Valente, Thomas W. 1995. *Network Models of the Diffusion of Innovations.* Creskill, N.J.: Hampton Press.

———. 1996. "Social Network Thresholds in the Diffusion of Innovations." *Social Networks* 18:69–89.

White, Harrison C. 1992. *Identity and Control: A Structural Theory of Social Action.* Princeton, N.J.: Princeton University Press.

Young, Peyton. 1996. "The Economics of Convention." *Journal of Economic Literature* 10:105–22.

———. 1998. *Individual Strategy and Social Structure : An Evolutionary Theory of Institutions.* Princeton, N.J.: Princeton University Press.

Zhao, Dingxin. 1998. "Ecologies of Social Movements: Student Mobilization during the 1989 Prodemocracy Movement in Beijing." *American Journal of Sociology* 103:1493–529.