Edited to include Erratum correcting printer error in initial publication. Erratum Published in vol 94 No 4 Jan 1989

Social Networks and Collective Action: A Theory of the Critical Mass. III[1]

Gerald Marwell and Pamela E. Oliver
University of Wisconsin—Madison

Ralph Prahl
Public Service Commission of Wisconsin

Most analyses of collective action agree that overcoming the free-rider problem requires organizing potential contributors, thus making their decisions interdependent. The potential for organizing depends on the social ties in the group, particularly on the overall density or frequency of ties, on the extent to which they are centralized in a few individuals, and on the costs of communicating and coordinating actions through these ties. Mathematical analysis and computer simulations extend a formal microsocial theory of interdependent collective action to treat social networks and organization costs. As expected, the overall density of social ties in a group improves its prospects for collective action. More significant, because less expected, are the findings that show that the centralization of network ties always has a positive effect on collective action and that the negative effect of costs on collective action declines as the group's resource or interest heterogeneity increases. These nonobvious results are due to the powerful effects of *selectivity,* the organizer's ability to concentrate organizing efforts on those individuals whose potential contributions are the largest.

Social ties are important for collective action. For example, it is widely agreed that participants in social movement organizations are usually recruited through preexisting social ties and that mobilization is more likely when the members of the beneficiary population are linked by social ties than when they are not (e.g., Tilly 1978; Oberschall 1973). But exactly how and why social ties are important is less well established. What kinds of ties are most important for collective action? What features

[1] This material is based on work supported by the National Science Foundation under grant SES-8408131. We wish to thank Charles Halaby and Robert Mare for advice on the statistical analysis and David Weakliem for supplemental research assistance. Requests for reprints should be sent to Gerald Marwell, Department of Sociology, University of Wisconsin, Madison, Wisconsin 53706.

0002-9602/89/9403-0002$01.50

of social ties are especially relevant? Are social ties important because they make mobilization less costly, or because they connect more people? And how can we distinguish these explanations?

This paper develops some analytic theory to begin to answer these questions by extending the "theory of the critical mass," our microsocial model of the process by which collective action emerges in groups. (See Oliver, Marwell, and Teixeira [1985] for an initial explication, and Oliver and Marwell [1988] for an application to the problem of "group size" in collective action.) We extend this theory here to include the overall frequency or density of social ties in a group, the extent to which the ties are centralized in a small number of group members, and the cost of using a tie for communication or coordination. We consider only the direct ties between potential organizers and potential participants, not the indirect ties that have often been investigated by social network theorists (Alba 1981; Arabie 1977; Burt 1978, 1980; Doreian 1974; Freeman 1979; Marsden 1981, 1983; White, Boorman, and Breiger 1976). We also do not distinguish the "strength" of the tie (Granovetter 1973). In our conclusion, we discuss possible relations between our findings and other features of social networks.

Like most formal models, ours requires some rather extreme simplifications. Even with these simplifications, the models are often too complex for analytic solutions, so we supplement our analysis with computer simulations, including Monte Carlo methods. Although we create the simulations, their outcomes often surprise us and force us to develop new theory to explain them. In fact, the most important finding of this research was wholly unanticipated. This is the finding that network centralization is beneficial for collective action, a finding that we show arises from the ability of organizers to be selective in choosing the targets of their organizing efforts.[2]

Of course, our goal is not just to understand our artificial simulation, but to generate theory that may apply to natural groups and actions. Therefore, throughout the discussion we attempt to be clear about the scope conditions or limits of generality of our results, and in the conclusion we discuss their substantive applications.

It will be helpful to ground our often technical discussion in an example. Many cities have passed ordinances requiring their employees to be city residents. Consider a number of cities that have just passed such

[2] Although one reviewer paid us the compliment of believing that we knew about selectivity when we planned our simulations, we actually did not. We developed the concept of selectivity from struggling with simulation results that we did not understand. We believe there is value in showing readers some of the steps of that process, especially since many sociologists believe that simulations are incapable of producing any "new" knowledge that is not apparent in the initial assumptions.

an ordinance, effective next year. None now exempts current employees from the requirement, but an amendment providing such an exemption is possible if a critical mass of employees can be assembled to mount a visible campaign. Assume that each city happens to have 400 employees who live outside its limits and that the amount of effort required to gain passage of the amendment is the same in all cities.

In the formal language of collective action theory, the 400 affected employees in each city are an *interest group*. The amendment is a *collective good:* all affected employees will benefit from it, regardless of whether they help to get it passed (Olson 1965).[3] Most formal collective action theory assumes that individuals make isolated, independent decisions whether to contribute to collective action and predicts that individuals will usually try to "free-ride," that is, to reap the benefits of a collective good provided by the contributions of *other* group members while paying none of the price. Thus, each suburb-dwelling city employee hopes her similarly affected colleagues will spend their time picketing and lobbying for the amendment, while she goes bowling. However, since everyone has the same hope, no one pickets (and the bowling alleys are full).

Of course, the assumption that individuals act in isolation is usually wrong. In most decision situations, people are at least generally aware of what others are doing, and often they have social relations that make influence, or even sanctions, possible. Empirically correct decision equations must take account of these interdependencies, and, when interdependency is considered, individuals can find it "rational" to participate in collective action (see Oliver, Marwell, and Teixeira 1985; Hardin 1982).

[3] Although the sociological literature on formal collective action theory is relatively small (see Marwell and Oliver 1984; Marwell and Ames 1979; Oliver 1980; Oliver et al. 1985; Oliver and Marwell 1988; Oberschall 1973, 1980; Fireman and Gamson 1979; Smith 1976), the literatures in economics, political science, and psychology are much too large to review adequately here. The "public goods problem" is long-standing in economics, although Samuelson (1954) and Bator (1958) are usually seen as the classic statements; Head (1974) provides a clear treatment of the fundamental economic issues, Schelling (1973) ties public goods to the more general issue of externalities, and Brubaker (1982) examines some empirical evidence in experimental economics. Olson (1965) is important for having exported the economic ideas to other social sciences but does not provide a definitive formal statement of the problem. Political scientists who have refined these ideas and applied them to collective action dilemmas include Chamberlin (1974), Frohlich and Oppenheimer (1970) and Frohlich et al. (1975), Hardin (1971), Schofield (1975), and Van de Kragt, Orbell, and Dawes (1983). Perhaps the most thorough and wide-ranging treatment of the issues of collective action is Hardin (1982). Psychologists generally explore experimentally the dynamics of social dilemmas in relatively small groups (Bonacich et al. 1976; Brewer 1985; Dawes 1980; Dawes, McTavish, and Shaklee 1977; Dawes and Orbell 1982; Kelley and Grzelak 1972; Messick and McClelland 1983; Messick et al. 1983).

Social networks matter when decisions are interdependent. But just how do they influence the prospects for collective action? What characteristics of their networks determine the ability of the affected employees to mobilize in support of the amendment? We here investigate three prominent features of social networks of interest groups: their density, their centralization (or concentration), and the cost of communicating across a tie.

Density

Most discussions of social networks and collective action consider rather stark contrasts, like that between the cities Alpha and Beta. Alpha is fairly isolated, and most of its affected employees live in a single suburb, Centauri. They attend Centauri's churches; their children attend Centauri's public schools; they belong to Centauri chapters of social and service clubs; they all are served by the same local telephone exchange. In contrast, Beta is part of an ethnically diverse two-state megalopolis, and its affected employees are scattered across a dozen different suburbs in two states and four counties. Thus, they rarely see one another after work. They go to different churches, send their children to different schools, read four or five different newspapers, and pay toll charges for telephone calls between many of the suburbs.

We would find virtual unanimity among sociologists in predicting that the employees in Alpha are more likely than those in Beta to act collectively. Other things being equal, the large number of social interconnections in Alpha provides a solid basis for interdependence, while their virtual absence in Beta is surely a hindrance. It is practically a truism among social movements theorists that social networks are important for recruiting participants (see, e.g., Oberschall 1973; Tilly 1978; and Fireman and Gamson 1979), and these theorists are supported by a fair amount of empirical evidence (see Snow, Zurcher, and Eckland-Olson [1980] and Walsh and Warland [1983] on participation in movement organizations, or Cohn [1985] on participation in a clerical workers' strike). The basic argument was important even to Marx, who reasoned that employees who are in regular contact with one another will develop a "habit of cooperation" and be more likely to act collectively than employees who work in isolation. Such generalizations, which are almost certainly true and of great practical importance, concern the overall number, or *density,* of social ties[4] in the group.

[4] For network specialists, the term "density" often refers to the extent to which the people known by one person tend to know each other. We are using the term in the nonspecialist's sense of the number of actual social ties in a group divided by the number possible. This seems the only reasonable term for such a concept at the group

Centralization

But there is more to know about an interest group's social ties than simply how many there are. We also need to know how they are structured—more specifically, the extent to which they are *centralized* or concentrated in a few individuals rather than being spread more evenly across the whole group. Network centralization is the key independent variable for this analysis because of its complex and so far unexplored relationship to collective action. In contrast, although density is obviously important, its effects are generally just that—obvious.[5]

To illustrate the effects of centralization, we may consider the cities Delta and Epsilon. The affected employees of Delta have the same total number, or density, of social ties as the employees of Epsilon. However, these ties are distributed quite differently. Delta has annexed all its potential suburbs. Its affected employees are therefore not suburbanites but residents of Delta's "exurbs." They are lightly sprinkled around a sparsely settled rural fringe. They shop in the city's many shopping centers and attend city churches. In fact, they share only one important local institution—the county primary and secondary schools, which serve the entire rural area. The one affected city employee who is on the county school board therefore accounts for a very high proportion (let us say 50%) of all the ties among affected employees.

Social ties among affected employees in Epsilon are much less centralized. Epsilon is the central city of a small metropolis, and affected employees live in a number of contiguous suburbs with diffuse and indistinct social and economic boundaries. There are a few dozen churches that serve larger or smaller catchments, depending on the denomination. Service clubs have loose geographical bases and draw in residents of adjacent suburbs. There are eight municipalities, four school districts, and three suburban newspapers with distinct but partially overlapping territories. Different employees see each other in different contexts.

To make a prediction about the prospects for Delta and Epsilon, we have to assess the effect of the different levels of *centralization* of network

level. When group size is fixed, as it is in our example and simulations, there is an equivalence among the total number of ties in a group, the density (total ties divided by number of ties possible), and the average number of ties per member (total ties divided by number of members). Thus, we often use these three terms somewhat interchangeably, although, of course, when group size is not held constant, they are not equivalent.

[5] Again, because we consider only direct asymmetric ties, our use of the term "centralization" is different from the more complex usage of network theorists, e.g., Freeman (1979). We are referring simply to the tendency for a few people to account for most of the ties, while the network concept of centrality refers to the extent to which any particular person is an important link in the indirect network relations of others.

ties. We have to decide whether collective actions are more likely to thrive where a small number of people know many others (while most know almost no one), or where many people each know some others, but no one knows a large number. The centralization in the school board in Delta suggests the possibility of efficient coordination of action by the affected employee, who knows everyone else. But it also means that if the school board member is incompetent or is just not interested in spearheading the collective action, no alternative affected employee has much of a chance to get things going. On the other hand, the crisscrossing social ties and overlapping memberships in Epsilon evoke images of more diffuse potential loci for action. If one person does not show initiative, perhaps another one will. The problem in Epsilon is that no one person is in a social network position to gather together efficiently some major proportion of all the affected workers. There are advantages and disadvantages in either situation, and, in the absence of good theory, it is not obvious how these will balance out.

Cost

Our third independent variable is organizing *cost,* that is, the cost of using one of the social ties for some purpose, such as asking the person to participate in a collective action. This variable is often confounded with network density in work on social ties. To some extent, we committed this confounding in our comparison of Alpha and Beta. Factors that make it cheaper to communicate across social ties, such as free local versus costly message-unit telephone rates, may also help determine whether such ties exist at all. Empirically, the two are often correlated. However, a *theoretical* distinction between the presence or absence of ties and the cost of using them is needed because each factor has a different effect on the prospects for collective action.

To illustrate this issue, we need a three-way comparison, such as among cities Mu, Sigma, and Rho. Suppose Mu's affected employees average 10 ties to other affected employees and generally have to pay message-unit charges for calls to one another. Sigma's employees also average 10 ties, but all the suburban areas are in the same extended calling area. Finally, Rho's employees average 15 ties but have the same message-unit charges as Mu's. Clearly, we would expect the prospects for collective action to be worse in Mu than in either Sigma or Rho. But what prediction do we have about the prospects for success in Sigma compared with those in Rho? Another way of asking this is, If one were in Mu, would there be a bigger effect from increasing the number of employees who know each other or from devising a cheaper mode of communication among those who already have social ties? A real group of employees might have to decide if its efforts to generate collective action would be helped more by social gatherings and other efforts to increase the group's density of social ties or by finding ways to reduce the cost of communication among people who already know each other, perhaps by printing a directory of home addresses and telephone numbers. This is really a question about comparative rates of change, and it does not have a readily apparent answer.

Group Heterogeneity

We expect the effects of the network variables to interact with the degree of resource and interest heterogeneity in the interest group. The fact that a good is collective or shared does not mean that all those individuals who are interested in it are the same. For example, although all present employees who live outside the city share an interest in an exemption, they differ in the magnitude of that interest, depending on such factors as how easy it would be for them to get other jobs, whether they own or rent their homes, and whether their spouses are employed. They may also vary in the amount of time and money (i.e., resources) objectively available to them. We have previously shown that groups that are quite heterogeneous differ in their patterns of collective action from groups that are more homogeneous (Oliver, Marwell, and Teixeira 1985).

The direct effects of group heterogeneity on collective action can be briefly summarized. When groups are homogeneous, everyone is interchangeable, and the collective action outcome is a simple function of *how many* people participate. In contrast, in a heterogeneous group, it matters *who* is organized as well as how many since one person may be willing and able to contribute much more than another.

To understand the pure (ceteris paribus) heterogeneity effect, we may imagine that a group's resource and interest heterogeneity (i.e., standard deviation) could be varied while the mean is held constant. Changes in heterogeneity would have two effects, which we may call the "mean level" effect and the "conjunction of probabilities" effect. The mean level effect refers to the relation between the mean of a distribution of resources or interest and the level necessary for success. Briefly, if a distribution's mean is sufficiently high, more heterogeneity hinders collective action; if the mean is too low, heterogeneity helps.[6] For example, if the mean amount of employees' free time is so low that no group members can spend three hours at a protest meeting, increasing the heterogeneity of free time could produce a few who could attend. They would be produced

[6] For a variable for which high values are bad for collective action, heterogeneity is beneficial when the mean is high and harmful when the mean is low.

statistically by reducing further others' free time, but this will not affect collective action because those people could not participate in any case. Conversely, when the mean member has the three hours free already, increasing the heterogeneity reduces the proportion of the group that has three or more hours free by increasing the proportion below three.[7]

The second effect, the "conjunction of probabilities," arises from the fact that the joint probability of two or more independent events must be lower than the marginal probability of any one of them. Thus, increasing the heterogeneity of two or more variables at the same time tends to have a negative effect on the outcome. For example, the risk goes up that the employees who have the most time available (high resources) will not be the same ones who are most concerned about losing their jobs (high interest).[8] Thus, the "mean level" principle (which produces a positive heterogeneity effect) and the "conjunction of probabilities" effect jointly create an optimization problem in which some heterogeneity generally improves outcomes but too much heterogeneity usually worsens them.

THE MODEL

Any formal analysis requires a series of definitions, simplifications, and assumptions that define its theoretical model of the complex empirical world. In this section, we describe the elements of our model.

All-or-None Contracts

Networks are irrelevant to isolated independent actors. They affect collective action only when there is some "social resolution" to the collective action dilemma. A variety of assumptions modeling this "social resolution" are available, but the specific form of this assumption at the individ-

[7] This "mean level" principle requires a caveat for skewed distributions. In a skewed distribution with a lower bound of zero (like the lognormal), increases in heterogeneity come from a few people with very high values balanced by many people with values close to zero. If the mean is too low to support collective action, increases in heterogeneity first increase the proportion with a high enough value to act but then decrease that proportion as the distribution comes to consist of fewer and fewer individuals whose interest or resource level gets progressively higher, much higher than necessary, balanced by the vast majority whose interest is much too low. (This does not arise with symmetric distributions, for which increasing the heterogeneity is always beneficial if the mean is too low.) At the extreme, high interest heterogeneity would come from one person who would lose everything balanced by everyone else who would lose nothing. Unless this one person is also extraordinarily resourceful, there cannot be a critical mass for collective action.

[8] This risk is increased when distributions are skewed so that most of the group have very low values on each of the independent variables.

ual level turns out to have little impact on the effects of the network variables as predictors of collective action at the group level. The particular model we have chosen to use in this analysis is called the "all-or-none contract."

The terms of an all-or-none contract are that all parties to the contract contribute their agreed amounts, or no one does, and there is some mechanism for ensuring that the terms of the contract are carried out. Individuals are asked whether they will agree to bind themselves to the contract, with the proviso that they will actually make their contribution only if the sum of all contributions will exceed a specified level. We assume that, in this situation, individuals weigh their own costs against their personal payoffs from the *total* change in the collective good produced by all the contracted contributions. Individuals may therefore find it profitable to bind themselves to such a contract even when they would not make independent contributions (Oliver, Marwell, and Teixeira 1985).[9] Individuals' contributions to these contracts can be negotiated, and different parties to the contract can contribute different amounts. Formally negotiated all-or-none contracts are rather common resolutions to some problems, such as providing office microwaves. The model of the formal contract is a reasonable approximation of more informal processes relying on convention (Hardin 1982, pp. 155–230), normative pressure, or solidarity (Fireman and Gamson 1979).

Organizers

As common as formal or informal contracts are, they rarely happen spontaneously. They entail communication and enforcement costs that are never uniformly distributed across parties to the contract. Empirically, these kinds of contracts generally have well-defined points of origin, that is, some specific person (or small cadre) who incurs the organizing costs. In short, they have organizers. Organizers, just like everyone else, have available resources and an interest in the collective good, but they use their resources to organize a contract that will make others willing to contribute to the collective good.

[9] We assume that people will contribute if their contributions will produce a profit—i.e., if the total benefit exceeds their own costs—and that they ignore the possibility that someone else will provide the good. If actors do attend to the possibility that someone else would also find it profitable to contribute, we must have information about their subjective probability estimates about others' actions and employ a more complex model, which is beyond the scope of this paper. This gaming dilemma is acute when independent decisions are modeled, but it is less acute when social resolutions are modeled, for it is plausible to assume in such circumstances that all the high-probability actors are parties to the social resolution.

Organizer Networks

We assume that each person in the group has an asymmetric ego network (Burt 1980, p. 91) or "organizer network" of other group members he or she can organize. In this analysis, we ignore indirect links and consider only a direct asymmetric "send" link: *A* is linked to *B* if it is physically and socially possible for *A* to ask *B* to participate in a given collective action. We assume that organizers can recruit only those individuals with whom they have some prior social connection. It may be easy for a policeman to recruit other members of the force to support the amendment but socially impossible for him even to ask secretaries from the mayor's office to join the pickets. We assume asymmetry because we believe that stratification and other social realities often prevent a person from approaching someone who could approach him.[10] A police lieutenant might recruit patrolmen, who would be wary of initiating discussions with him. We treat network ties as present or absent without considering the intensity or nature of the connection; we treat the cost of communicating across a tie as a separate variable from the existence of the tie itself.

These assumptions greatly simplify our treatment of networks. Let t_i be the number of ties (i.e., the "out-degree") for individual i, and let T be the total number of ties in the group. If the total interest group size is N, the number of possible ties is simply $N(N - 1)$, and the overall interest group density, D, is $T/N(N - 1)$; the density, d_i, of each individual's ego network is $t_i/(N - 1)$. In general, therefore, the mean of the individual network densities equals the overall group density.

Given a specific group density, the *centralization* of an interest group's social network may be defined as the standard deviation, σ_t, of the number of ties across the ego networks of all group members. When all the group's ties are concentrated on one individual (and everyone else has none), the group's network is fully centralized, and σ_t is very large. When all group members have the same number of ties, the network is completely decentralized, and σ_t equals zero. The standard deviation of the ego-network densities, σ_d, equals $\sigma_t/(N - 1)$ and may be interpreted similarly. When N is treated as fixed, as it is in this article (where $N = 400$ for all groups), there is a direct equivalence between the standard deviation of ego-network densities and the standard deviation of numbers of ties since the former is just the latter divided by $N - 1$.

[10] These asymmetric ties may be mutual. The link of *A* to *B* does *not* preclude *B*'s link to *A;* it just does not entail it. In the simulation, the two links (*A* to *B* and *B* to *A*) would have independent probabilities of occurrence. Our analysis assumes that the mutuality of ties has no effect, which is true in our simulation.

Organizing Costs

Organizers may or may not be able to communicate with all the group members they know, depending on communication costs. Face-to-face contacts take a certain amount of time per contact. Letters and telephone calls usually require less time per contact but more money. Local telephone calls cost less than long distance. Mass media (television, radio, newspapers) or diffuse media (posters, leaflets) have certain costs, which include the "waste" of communicating with people who are not potential contributors. It is much less expensive for employees to contact each other if they all work in the same building than if they are dispersed all over town, and less expensive if they all live in the same suburb than if they are spread thinly across the countryside. Interdepartmental mail is much cheaper to use than the U.S. mail. The surveillance or enforcement mechanisms for the contract may also have some material or social cost.

The effects of different kinds of organizing costs are complex enough to require separate analysis. In the present work, we make the very simple assumption that organizing costs are exactly proportional to the number of parties to the all-or-none agreement and that this constant cost per participant is a characteristic of the interest group that does not vary with the organizer or the particular individuals involved. Thus, average costs may be higher in one city than another because of telephone cost differences, but Fred Jones and Ann Smith, who are both employed by the same city, are assumed to face identical costs per network member recruited.

Individual Decision Model

We assume that individuals' decisions about contributing to the all-or-none contract are determined by weighing the costs and benefits of such a contribution. We begin with a group of individuals defined by the fact that they will all be positively affected by some collective good, although they differ in their resources (r) and in the level of their interest in the good (I). Each person in that group is both a potential organizer and a potential contributor to a contract organized by someone else. The "production function" $p(r)$ specifies the relation between inputs of resource contributions (r) and outputs of levels of the collective good (p). For dichotomous goods, such as whether the exemption amendment passes or not, p represents the probability of obtaining the good.

If decisions are made independently, we may write a simple decision equation for the net gain from contributing: $G = p(r)I - r$. An individual acting independently finds that her individual contribution produces a profit if $G > 0$, that is, if $p(r)I > r$, meaning that the individual's interest times the probability of obtaining the good exceeds the size of her contribution. An algebraically equivalent expression for this condition is $p(r) > r/I$. Although independent contributions to collective goods are rare, they are not impossible. Situations may exist in which there is at least one "large" potential contributor whose interest is sufficient to make him willing to provide the collective good himself. This occurs because collective goods often have high jointness of supply, which means that the cost of providing the good does *not* rise with the number who enjoy it (see Hardin 1982; Oliver and Marwell 1988). If individuals are willing to contribute without organizing, there is no need for an organizer and no need for anyone to bear the organizing costs. (These large contributors might find it cheaper to organize others' interdependent contributions than to make their own independent contributions, but such considerations of alternative strategies are beyond the compass of this paper.) Thus, one scope condition for our theory of organizing is $r/I > p(r)$, that is, that individuals will *not* find it rational to make independent contributions.

The individual decision rule for contributing to an *interdependent* all-or-none contract makes one crucial change in the previous equation. The *total* contribution by all parties to the contract, t, replaces r, the *individual's* contribution level, in the production function. This gives us the revised decision equation $G = p(\Sigma r)I - r$. Contributing to the contract produces a profit if $p(\Sigma r) > r/I$, that is, if the total payoff from *all* contributions to the contract exceeds the individual's r/I ratio. Obviously, this is much more conducive to a positive decision.

Organizers' Decisions

A group member's decision to organize is more complex than the decision to contribute. Organizing will be profitable if the $p(\Sigma r)$ produced by the contributions of the *other* group members exceeds the *organizer's* r/I. A group member's interest (I) is the same regardless of whether she is organizing or contributing. To find the cost of organizing, we make the simplifying assumption that organizing costs are linear with the number of contributors. If C is the cost per contributor and k is the number of contributors, then the *organizer's* cost, or contribution, equals kC. If we let Σr_k represent the total contributions of the k individuals the organizer contacts, then the organizer's gain from organizing is $G_0 = p(\Sigma r_k)I - kC$. The payoff will be greater than zero, and organizing will be profitable, if the potential organizer knows and can afford to contact a group of k individuals, such that $p(\Sigma r_k) > kC/I$.

The Process of Forming a Contract

We simplify our model greatly by assuming that every group member has perfect information about everyone's interest and resource levels and that everyone's behavior is a perfectly determinate function of his own interests and resources. This means that when social worker Fred Jones, organizer, approaches patrolwoman Ann Smith, group member, with the proposition, "Will you agree to spend five hours fighting for an exemption if you are guaranteed that a total of fifty hours will be contributed to this project?" Fred knows the answer before he asks. He can mentally review a list of his acquaintances and determine whether he knows and can afford to contact a critical mass of people who are interested and resourceful enough that their combined contributions would produce a big enough impact to make everyone willing to contribute. He spends his resources to pay the organizing costs only if he is sure he will be successful. In this simplified world, there are no unsuccessful organizing drives.

Group Outcomes and Multiple Organizers

Two (or more) organizing campaigns going at the same time could reinforce each other, but they could also hurt each other if they compete for the same contributors. The complexities of multiple mobilizations are beyond the scope of this paper. We hope to address this issue in future analyses. Here, we will only identify the *best* contract possible in a given group. This implies that mobilization is preemptive, that is, that one organizer's efforts preclude another's. This is often a realistic assumption. Organizing has a decelerating production function (Oliver, Marwell, and Teixeira 1985), and many people will not organize if they know someone else is (see Oliver 1984). Preemption is especially common when the interest group already has some "official" representative (such as a city employees' union). In real life, the preemptive organizer is not necessarily the one who would organize the best possible contract.

The probability that an interest group will contain at least one successful organizer is much higher than the probability that any particular individual can be a successful organizer. When preemption is assumed, the probability that at least one of N potential organizers will be successful is given by $1 - (1 - q)^N$, where q is the probability that any particular organizer will be successful. Very small probabilities of success for individual organizers produce very large probabilities of at least one success in a group. For example, in a 400-person group, if the probability of any individual's being a successful organizer is .01, the probability of there being at least one successful organizer is over .98. For there to be a better than even chance of the group's having at least one success (i.e., the group probability is greater than .5), the individual probability can be as low as .002.

THE SIMULATION

Some of the effects on collective action of network density and centrality, organizing costs, and the heterogeneity of resources and interests may be shown with formal analysis. Their interactions, however, are more opaque. We need some kind of data for testing our ideas about these more complex relationships. Since no real data are available with the kind of information we need, this section presents an analysis of artificial data constructed using computer simulation techniques. In the next section, we use the results of this analysis as a point of departure for a more general theoretical discussion.

Our analysis builds from the simple to the complex. We first consider groups that differ only in their organizing costs. Next, we vary the density of network ties, and then their centralization. The results to that point are essentially analytic, although we create a data set to compare the relations among these variables with those from the simulation. The last step involves varying the dispersion of interest and resources around their means. This problem does not permit analytic solutions and is investigated with data produced by a rather complex simulation of the decision-making processes involved.

Fixed Parameters and Algorithms

We hold some factors constant. All groups have 400 members. The collective good is dichotomous: the exemption either passes or fails, with contributions of resources affecting the probability that the good will be provided. Individuals have resources that may be spent either for organizing or for making a contribution to a contract organized by another. Organizing costs are paid wholly by the organizer, are proportional to the number of group members who contribute to a contract (regardless of the sizes of their contributions), and do not vary among individuals in a group, although they vary between groups. As a further simplification, we do not permit organizers to make direct contributions to the collective good even if they have resources left over after organizing.

A standard metric is created by setting the total resource contribution necessary to provide the good with certainty equal to one; that is, $p(1) = 1$. Organizing costs are assumed to have the same metric. An individual's interest level, I, is expressed in the same metric as the payoff he would experience if the good is provided with certainty, so that the individual's benefit from a particular provision level equals $Ip(\Sigma r)$. If the production function is linear, the standardization means that the total amount of resources contributed equals the provision level; that is, $p(\Sigma r) = \Sigma r$.

Rather than being linear, the production function we use in this analysis is slightly accelerating: $p(\Sigma r) = t^{1.2}$; $0 < t < 1$. We have argued elsewhere that contractual resolutions to the collective dilemma are especially important when the production function is accelerating, that is, has increasing marginal returns (Oliver, Marwell, and Teixeira 1985). In this case, individuals are positively interdependent, so the value of any individual contribution is *greater,* the more others contribute, and there are initial periods of low returns (i.e., start-up costs) that are difficult to overcome without interdependence. The particular choice of function is arbitrary.

The total resource and interest levels of the group as a whole have an enormous influence on the prospects for collective action. Because we are interested in other independent variables, we must constrain these factors to ranges that neither guarantee nor prohibit contributions. We set the *mean* of the distribution at .06 for resources and .15 for interest. The interest mean is substantially higher than the resource mean so that many group members will meet the criterion $p(\Sigma r) > r/I$, which permits contracts that produce probabilities between zero and one.

Dependent Variable

The *group* outcome is the best possible outcome across all possible organizers in the group. Each group member is a possible organizer, and we calculate the total contribution from his best possible contract, given the specific others to whom he has social ties. We then select the one organizer in the group whose contract commands the highest total contribution. His contract becomes the *group's* contract, and the group payoff becomes this total contribution raised to the 1.2 power. Since payoff and contribution are perfectly correlated, we use the latter in our regression analyses. In other analyses, we transform this continuous contribution/payoff variable in three ways: into the dichotomy between achieving certainty (contributions total 1.0 or more) and not; into the dichotomy between failing completely (contributions total zero) and not; and into the trichotomy among total failure, partial success, and achieving certainty.

Independent Variables

For each independent variable, we define a range of values of interest. The actual values assigned to each case on each variable are systematically generated or randomly selected from the uniform distribution defined by that range. These uniform distributions do not mean anything substantive; they are continuous analogs of the equal cell *N*s in a standard fixed-effects analysis of variance design.

Individuals' personal network sizes (or out-degrees), resources, and interests are distributed lognormally in each group (Aitchison and Brown 1963). It is most plausible to assume these distributions are skewed, both because they are constrained to be nonnegative and because one would generally expect to find a large number of group members low and only a few high on each variable. In most groups, there are many more people who are very poor than are very rich. The choice of the lognormal, rather than some other skewed distribution, is largely practical.

Interest and resource *heterogeneity* is operationalized as the *standard deviations* of lognormal distributions. Among groups, these range from .01 to .06 around the resource mean of .06, and range from .025 to .15 around the interest mean of .15.

Organizing cost is defined as the cost per each member of a contract. It varies among groups from .001 to .010 in increments of .001. Organizing is never costless, but its cost ranges from insignificant to prohibitive, given organizers' resource levels. Each of a group's 400 members faces the same organizing cost.

The density and centralization of group networks are operationalized in a nonstandard way because these numbers must be integers. We operationalize a group network's density as the *mode* of a lognormal distribution of an individual's personal network sizes in the group, and the group's centralization or dispersion as the difference between the mode and the "largest possible" personal network, where the "largest possible" personal network is the one above which less than .0005 of a lognormal distribution with the given mode falls. In these data, the correlation between the mean and the mode of the distribution of personal network sizes is .99, while the correlation between the standard deviation and the measure of dispersion we used is .96.

The range among groups for modal personal network size was set as low as possible, ranging from 1 to 10. Dispersion was chosen to range from 2 to 20, a range that seemed on its face to provide extensive variation among groups in the degree of centralization of network ties. The "largest possible" group size is the sum of mode and dispersion and therefore ranges from 3 to 30 with a flat bell shape. The range of modes is extremely unfavorable for successful contracts, but the range of "largest possible" sizes includes a good proportion that can achieve certainty.

We used a computer algorithm for actual assignment of network sizes to individuals in the groups to produce the frequency distribution of integers that is the best approximation to the lognormal distribution defined by the mode and the "largest possible" group size. Because the "largest possible" group size is the level at which one in 2,000 networks would obtain—although the groups have only 400 networks—the actual size of the largest group in these frequency distributions is a little lower, especially when the dispersion is high (but the correlation between the two is .98).

Ties are randomly assigned to group members in accordance with this frequency distribution. For example, an individual might be randomly determined to be able to organize (i.e., "send to") eight others. The particular eight would be randomly chosen from the 399 other group members entirely independently of the presence or absence of other ties. That is, the fact that *A* can organize *B* has no effect on the probability that *B* can organize *A*, and the fact that *A* and *B* can both organize or be organized by *C* has no effect on the probability of a link between *A* and *B*.

RESULTS

Although we are ultimately interested in the interactions among all the variables at once, we introduce them in three stages so that the results for the simpler cases may be used as baselines for interpreting the more complex.

Case 1: Homogeneous Groups

We begin our analysis by considering interest groups that have no *within*-group variance on any important characteristic. Each member therefore has .06 in resources, an interest of .15, and some constant number of other group members she could organize.

Under these circumstances, all group members, and the networks of those they could organize, are completely equivalent. The total group's contribution is simply proportional to the number organized. In turn, the number organized is determined by the least favorable of two constraints: the number of others each group member has a tie to that permits organizing (i.e., organizer network size) and the number he can afford to contact (i.e., r/C, available resources per person divided by the per-capita organizing cost). The outcomes for the ranges of variables used in our simulation are easy to calculate and are summarized in table 1. These results are not intrinsically interesting because they are a direct product of the ranges selected for organizing costs and network sizes. They provide the baseline information that the simulation has been parameterized so that homogeneous groups cannot provide the good with certainty and have only a one in five chance of being able to provide an intermediate level of the good. As it should, the model indicates that groups with more

TABLE 1

SUMMARY OF THE ANALYSIS OF THE HOMOGENEOUS CASE

DESCRIPTION	OUTCOME*		
	Certain	Partial	Fail
Resources per person (r) (assumed constant)	.06	.06	.06
Minimum necessary total contribution to achieve specified outcome (t)*	1.0	.47	0
Minimum network size for minimum total contribution ($n' = [t/r]$)	17	8	0
Maximum organizing cost permitting contact of n' others ($C' = r/n'$)	.0035	.0075	any
Proportions having $n \geq n'$, for groups in simulation with (modal) network sizes drawn from uniform distribution between 1 and 10	0	.3	.7
Proportions having $C \leq C'$ for groups in simulation with organizing cost drawn from uniform distribution between .001 and .010	.3	.4	.3
Proportions constrained to each outcome by the least favorable of n and C for the distributions used in the simulation	0	.21	.79

* A total contribution of one achieves certainty because of the standardization. A contract for total contributions less than one is viable if $r/I < p(t)$. For $r = .05$ and $I = .15$, the constraint is $p(t) > .4$. For the production function $p = t^{1\,2}$, the solution is $t > .47$.

social ties among their members and lower organizing costs should support more collective action.

Case 2: Heterogeneity in Social Ties

The next step in our analysis is to permit heterogeneity in organizers' network sizes, that is, variation within the group in the degree of network centralization. As long as interests and resources remain invariant within the group, the group's total contribution is still solely a function of the *number* of individuals organized. The organizing cost constraint remains the same, so that every group member can afford to contact exactly r/C others. The only difference from the first case is that individuals differ in the sizes of their personal networks. The key to analyzing case 2, therefore, is to recognize that the best *group* outcome is simply the outcome for the group member who would make the best organizer, and that the best organizer is always the one with the largest organizer network.

Thus, when resources and interests are homogeneous in a group, analysis of the effects of degree of centralization is simply a matter of determining the *largest* organizer network. The density and centralization of a group's network ties jointly determine the largest network size in that group. The higher the average number of ties (density) in the group is, the larger will be the largest network. For a given density, greater centralization leads to larger largest networks since available linkages become concentrated on a few sociometric stars, such as the Beta school board member. Of course, as centralization increases, the sizes of most group members' personal networks decrease. One's intuition might suggest that this decrease in the "typical" personal network size would have some harmful consequences, but for organizer-centered actions in homogeneous groups, this intuition is false.

Table 2 shows the outcome distributions for the ranges of variables used in our simulations. The computations of minimum necessary organizer network sizes and maximum possible organizing costs are the same as for the homogeneous case (table 1) and are not reproduced. The proportions constrained to each outcome by the organizing cost distribution is unchanged. The difference in outcomes is due entirely to replacing the distribution of the *modal* organizer network size in table 1 with the distribution of the actual *largest* organizer network size in table 2.

Just as for the homogeneous case, the particular proportions shown in table 2 are not intrinsically interesting, as they are directly determined by the parameter ranges chosen. It is clear, however, that network centralization (heterogeneity of network sizes) in the group permits much greater rates of success. The explanation for this is straightforward since a successful outcome depends only on the size of the *largest* organizer network, and heterogeneity around a mean increases the proportion falling above some value greater than the mean. This constraint interacts with the organizing cost constraint just as in the homogeneous case, with the outcome determined by the less favorable of these two constraints.

The substantive finding that centralization in a group's network promotes collective action has *not* been recognized in the literature. Although the importance of well-connected organizers is understood, there has been little recognition in theoretical discussions that the desirable properties for organizers imply group-level characteristics that give rise to them.

Case 3: Resource and Interest Heterogeneity

Understanding the effects of the group's network variables on collective action becomes much more difficult when groups are allowed to be heterogeneous in interests and resources as well as personal network sizes. Under these circumstances, outcomes are not determinate but are due to the conjunction of several probabilistic events. For collective action to occur, the group must contain at least one organizer network with enough resourceful people that the sum of their contributions forms a workable contract. That same network must also have an organizer who

TABLE 2

SUMMARY OF THE ANALYSIS OF THE CASE OF HETEROGENEITY IN NETWORK SIZES AND HOMOGENEITY IN INTERESTS AND RESOURCES*

	OUTCOME		
DESCRIPTION	Certain	Partial	Fail
Proportions having $C \leq C'$ for groups in simulation with organizing cost drawn from uniform distribution between .001 and .010	.3	.4	.3
Proportions with $n_{\text{largest}} \geq n'$ for groups in simulation with largest network sizes determined by modal network size uniformly distributed between 1 and 10 and spread between mode and largest possible network size uniformly distributed between 2 and 20	.34	.56	.10
Proportions constrained to each outcome by the least favorable of n and C for the distributions used in the simulation	.11	.52	.37

* Computations of minimum necessary network size and maximum possible organizing cost are the same as for table 1 and are not shown.

can afford to contact enough people to form the contract. In our symbolic language, $\Sigma r_k \geq 1$, and $r_{org} \geq kC$. These factors alone mean that there are three probabilistic elements: the organizer's resources, the size of the organizer's network, and the resources of that network's members. In addition, the members and organizer are constrained by their interests because those who could make especially large contributions, or organize a large network, will not do so if their interests are too low. In symbolic terms, the constraint for each individual is $p(\Sigma r_k) = (\Sigma r_k)^{1.2} > r/I$. These simultaneous inequations are difficult to solve because Σr_k is recursively determined. Individuals' willingness to contribute depends on how much others are contributing.

When resources and interests vary within the group, it is not necessarily true that the organizer network with the largest number of members will make the largest contribution of resources. There are sampling fluctuations in the group in the composition of each organizer network and, when network sizes are relatively small (e.g., 3–27 in our simulations) and the independent variables are lognormally distributed, these fluctuations may be very large. Within most of the ranges of our simulation, the outcome from a particular set of parameters varies widely, depending on the "luck" of the joint distributions of the variables.

We use Monte Carlo simulations to explore these interrelations. The complexity of the case is reflected in the complexity of the computer program that simulates it. The indirect way interest constrains resources requires an iterative algorithm to search for the best contract for a particular organizer with a particular network, and this algorithm must be repeated across all the organizers in a group. In brief, for each organizer, the algorithm searches for the "best" set of contributors from his network, given the organizer's resource constraint. When the resource and interest conditions neither assure certainty nor force failure, numerous iterations may be required to find the best "partial success" contract or to determine that such a contract is impossible. The program implementing this process takes a long time to run. For this analysis, we have generated 2,794 groups, with each group constituting a single case.

This material replaced by erratum

Effects of Resource and Interest Heterogeneity

Our central interest is in the effects of a group's network characteristics, but we begin with a brief summary of the effects of resource and interest heterogeneity to provide a proper context for those results. As table 3 shows, resource and interest heterogeneity increases the *dispersion* among groups in rates of collective action. A comparison of the first line (taken from table 2) with the second shows that resource and interest heterogeneity increases *both* the proportion of all groups achieving certainty *and* the proportion failing completely.

The rest of the table breaks groups into the outcomes that would be expected from the group's "network" variables (i.e., those in table 2) alone and shows the effects in these groups of greater heterogeneity separately for resources and interest. Inspection of these subtables reveals that resource heterogeneity has clear positive effects if the group network variables predict failure for homogeneous groups, clear negative effects if the network variables predict achieving certainty, and stronger negative than positive effects if the network variables predict a partial success. The effect of interest heterogeneity is more irregular: it is clearly negative if the group network variables predict achieving certainty, somewhat negative if they predict partial success, and small and irregular if the network variables predict failure. Overall, these data show the importance of the "mean level" effect of heterogeneity. This effect is more clear-cut for resources because the mean resource level is set relatively low, while the mean interest level is set relatively high.[11]

[11] Recall that the mean interest level is set at 2.5 times the mean resource level. Increases in interest heterogeneity increase the proportion of group members whose interest levels are below their resource levels, i.e., too low to contribute. Only when resource heterogeneity is very high, meaning that there are a few group members with very high resource levels, is interest heterogeneity potentially beneficial.

TABLE 3

EFFECT OF RESOURCE AND INTEREST HETEROGENEITY ON OUTCOMES, BY OUTCOME NETWORK VARIABLES WOULD PRODUCE IN A HOMOGENEOUS GROUP

CASES	PROPORTIONS					
	Fail	Partial	Certain	N	MEAN	SD
Homogeneous (analytic results)	.37	.52	.11	1,900	.46	.38
All heterogeneous (simulation)	.50	.31	.19	2,794	.38	.42
Failure if homogeneous	.69	.23	.08	1,007	.19	.33
Resource deviation:						
Low 1	.71	.29	...	252	.16	.26
2	.67	.31	.02	268	.19	.32
3	.67	.21	.12	243	.20	.35
High 4	.72	.09	.19	244	.21	.38
Partial success if homogeneous	.45	.39	.17	1,537	.41	.41
Resource deviation:						
Low 1	.37	.53	.11	395	.48	.40
2	.42	.41	.17	398	.45	.40
3	.45	.35	.20	366	.37	.41
High 4	.56	.25	.20	378	.34	.42
Certainty if homogeneous	.01	.16	.83	250	.95	.17
Resource deviation:						
Low 1	...	.11	.89	61	.99	.03
2	...	.10	.90	63	.97	.13
3	.05	.20	.75	65	.91	.22
High 4	...	.23	.77	61	.92	.21
Failure if homogeneous	.69	.23	.08	1,007	.19	.33
Interest deviation:						
Low 1	.71	.18	.11	267	.22	.37
2	.74	.18	.09	224	.18	.34
3	.68	.22	.10	249	.19	.33
High 4	.64	.33	.03	267	.17	.27
Partial success if homogeneous	.45	.39	.17	1,537	.41	.41
Interest deviation:						
Low 1	.40	.34	.26	378	.50	.44
2	.50	.34	.16	364	.40	.42
3	.45	.40	.15	392	.40	.40
High 4	.43	.47	.10	403	.36	.36
Certainty if homogeneous	.01	.16	.83	250	.95	.17
Interest deviation:						
Low 1	...	.03	.97	58	.99	.06
2	...	.19	.81	66	.96	.15
3	.02	.12	.86	65	.94	.19
High 4	.03	.33	.64	61	.90	.22

Effects of the Group's Network Variables

Treating groups as the units of analysis, we want to assess the effects of organizing costs and the density and centralization of network ties and to determine how these effects interact with the level of resource and interest heterogeneity. As table 3 indicates, however, the distribution of the outcome variable is quite unusual. Half of the cases fall at the low extreme of zero, that is, complete failure; nearly 20% fall at the other extreme of one, achieving certainty. The distribution of the 30% that fall between these extremes has a rough bell shape around a mean of .61; it is a bit flat (kurtosis = −.417) and has a small negative skew (−.382).[12] To be sure that the mode of analysis does not distort the results, we analyzed three dependent variables (total contribution, achieving certainty, avoiding failure) two ways (OLS regression, probit analysis). All yield the same basic patterns of results. We report only the standard OLS regression using the total contribution level.

In table 4, we show the linear regression of our outcome variable, the contribution of the most successful organizer network in the group, on the group's three network variables (density and centralization of networks, and organizing cost) separately for groups with different levels of interest and resource heterogeneity. In the first row of the table, a comparable regression is reported for data from case 2 (resources and interest are homogeneous) to provide a baseline for comparison.[13] For case 2, the total group contribution is exactly determined by the three independent variables but in a way that is only approximated by a linear regression, as indicated by R^2 less than one. We do not know what the best equation relating the three group network variables to the outcome is for case 3, but it is helpful to compare the linear approximation with that for case 2.

The variances of centralization, density, and cost differ, so we require standardized coefficients for comparisons of the effects of different variables in equations, along with the unstandardized coefficients for comparisons of the effects of the same variable between equations. For each independent variable, we report the result of a significance test for the

[12] If interest were always exactly 2.5 times the resource level, there would be no outcomes falling between 0 and .47, but because both vary, interest is sometimes greater than 2.5 times the resource level, and there can be partial successes below the .47 level.

[13] To provide an appropriate baseline for comparison, we created a data set consisting of 1,900 cases, one for each possible combination of 10 modal network sizes, 19 largest network sizes, and 10 levels of organizing costs. The outcome for each of these cases, when each individual has resource level .06 and interest level .15, was determined analytically and added to the data set. (The outcome is entirely determined by the three manipulated variables, so there is nothing probabilistic in these data.) The same regressions were performed on these data as on those generated for the heterogeneous case.

TABLE 4

LINEAR REGRESSION OF TOTAL CONTRIBUTION LEVEL ON CONTACTING COST AND THE DENSITY AND CENTRALIZATION OF NETWORK TIES, SEPARATELY FOR DIFFERENT LEVELS OF RESOURCE AND INTEREST HETEROGENEITY

	UNSTANDARDIZED			STANDARDIZED			
	Cost	Density	Central	Cost	Density	Central	R^2
Homogeneous	−.10	.03	.02	−.78	.22	.27	.72
All heterogeneous	−.05	.08	.03	−.29	.56	.45	.60
Resource deviation:							
Low 1	−.06	.08	.03	−.36	.54	.47	.64
2	−.05	.08	.03	−.33	.58	.46	.64
3	−.05	.08	.03	−.28	.54	.45	.59
High 4	−.03	.08	.03	−.18	.56	.42	.56
Significance of difference*	<.0001	N.S.	N.S.				
Interest deviation:							
Low 1	−.05	.04	.08	−.29	.51	.51	.60
2	−.05	.08	.04	−.28	.55	.47	.62
3	−.05	.09	.03	−.28	.59	.42	.62
High 4	−.05	.08	.03	−.31	.60	.39	.59
Significance of difference*	N.S.	.29	<.0001				

* This is the p for an F-test on the difference of slopes across the four heterogeneity levels for each independent variable taken one at a time. See n.14 in text for model specification.

difference in the unstandardized coefficient, depending on the level of interest and resource heterogeneity.[14]

In examining table 4, we see three important patterns emerge quite clearly.[15] First, and perhaps foremost among these, all three independent variables are strong predictors of the group's level of contribution. Only

[14] Table 4 shows the p values for six F-tests that tested each of the three independent variables separately for interaction with resource and interest heterogeneity. The restricted model in each case was $Y = b_0 + b_1M + b_2S + b_3C + \Sigma_{bk}h_k$, where b_0 is the intercept, M is the modal network size, S is the spread of network sizes, C is the organizing cost, and the h_k is three dummy variables to represent the four levels of heterogeneity of either resources or interest. The full model in each case is $Y = b_0 + b_1\mathrm{M} + \mathrm{b}_2S + b_3C + \Sigma_{bj}h_j + \Sigma b_kh_kX$, where X is either M, S, or C. That is, the full model allows the independent variable in question to have a different slope for each heterogeneity level. To aid interpretation, table 4 shows the coefficients for the independent variables for separate regressions for each heterogeneity level.

[15] In general, results from analyses of the variables of avoiding failure and achieving certainty closely replicate those using the group's total contribution. The effect of organizing cost on avoiding failure declines, but not significantly, with increasing resource heterogeneity. The effect of density (mode) on achieving certainty increases with resource heterogeneity but decreases with interest heterogeneity. Otherwise, all three dependent variables show the same results.

for organizing cost under conditions of high resource heterogeneity is there a standardized coefficient (beta) less than .2. As we have noted, we expected that the overall density of ties and organizing cost would be strong predictors of collective action. However, the consistently positive effect of centralization is a result not anticipated in the literature.

The second significant pattern in table 4 is that the importance of organizing cost as a constraint on collective action declines with resource heterogeneity. In the "homogeneous" data, the coefficient for organizing cost is much larger than the coefficients for either the mode (density) or the spread (centralization) of group network ties. This relationship *reverses* for the "heterogeneous" data, where the coefficients for both mode and spread are much higher than for organizing cost, and the coefficient for organizing cost declines significantly as resource heterogeneity increases. The third important pattern in table 4 is that the coefficient for centralization declines significantly as interest heterogeneity increases.

In the next section, we provide a general theoretical discussion of the importance of organizing costs and the density and centralization of network ties under conditions of resource and interest heterogeneity. This discussion ties these specific results to general principles that have broader applicability.

DISCUSSION

No theory or simulation that hopes to reflect the most commonly held understandings of the empirical social world could fail to find that the sheer density of an interest group's network and its organizing costs are important predictors of collective action. Our "news" is the strong positive effect of the group network's centralization on collective action, and the interaction of all three independent variables with interest and resource heterogeneity. In this discussion, we identify the underlying principles that produced these nonintuitive results, thus demonstrating their applicability beyond the bounds of our simulation.

Centralization as Network Heterogeneity

For a given level of group network density, the centralization of ties means that group members are more heterogeneous in the number of ties each has. Some of our results can be explained by the general principles determining whether heterogeneity promotes or impedes collective action. The overall positive effect of centralization is due in part to the "mean level effect." The modal personal network size in the simulation is always too small to achieve certainty and usually too small to support even partial success, so heterogeneity around the mode is necessary for there to be any organizer networks large enough to support collective action. This effect obviously obtains only when the average organizer network is too small. If the average personal network size is large enough to support collective action, we would not expect centralization to be beneficial.

The declining effect of centralization as interest heterogeneity increases may be explained as a "conjunction of probabilities." As centralization and interest heterogeneity jointly increase, the risk grows that the organizer of the single largest network (in a group with mostly very small personal networks) will not be interested enough to organize her network. Although the effect of centralization is always significantly positive in our simulation, even for the highest levels of interest heterogeneity, the decline in the size of the effect might seem to suggest that there may be situations in which centralization would not be beneficial. However, we believe that this suggestion is false, as we show below.

Selectivity

We believe the effect of what we call "selectivity" is the most important theoretical insight to emerge from our analysis. We implicitly built selectivity into our model without recognizing its enormous potential for affecting collective action. Although it is arguable that we should have anticipated its effects, the fact is that we did not. We had to develop the concept to explain two puzzling results. First, proceeding from our understanding of the effects of the conjunction of probabilities, we expected centralization to have a curvilinear effect in interaction with interest and resource heterogeneity, with some moderate level of centralization being optimum. The declining slope in interaction with interest heterogeneity was consistent with our expectation, but we did not understand why no such effect appeared for resources or why the effect of centralization is always positive. Second, we unexpectedly found that the magnitude of the negative effect of organizing costs on collective action decreases greatly as resource heterogeneity increases. We had not expected organizing costs to show any interaction effects. We can explain these results by developing the concept of selectivity.

Consider the decisions being made by an organizer in the field. If she can afford to contact everyone in her network, she does, and there is no selection involved. But suppose organizing costs are high and her resources are low. Then she can contact only some fraction of the people she knows and could potentially organize. Whom will she contact? Past theorizing has often used group means and thus implicitly assumed that organizers randomly pick targets. But random choice would be silly and wasteful. Logically, an organizer operating under resource constraints will approach those individuals whose contributions seem likely to be largest. That is how to maximize outputs of contributions for inputs of resources spent on organizing. Framed this way, the process of selection is so obvious that we programmed it into our algorithm without considering its implications. We wrote code that simply rank-ordered all the members of an organizer's network and chose the top $k = r_0/C$ of them, where k is the number chosen (truncated to the nearest integer), r_0 is the organizer's resource level, and C is the cost per contact.

We think it will prove useful to develop the *degree of selectivity* as a formal concept. Each of the n members of the organizer's network has a potential contribution, which is his resource level constrained by interest. The distribution of these contributions has a mean, R, and the total potential contribution is nR. If the organizer can afford to contact only k of those he knows, he will obtain contributions from the *top* k/n of the distribution. This top k/nth of the distribution has a mean, R_k, and the total contribution obtained is kR_k. We can express the degree of selectivity in organizing with the ratio R_k/R and can investigate its properties.

A detailed mathematical and numerical analysis of this ratio is beyond the scope of this paper (see Oliver and Marwell 1987). Briefly, for a constant number k whom the organizer can afford to organize, the selectivity ratio increases as n increases. For a given k/n ratio, the selectivity ratio increases as the dispersion in contribution sizes in the organizer's network increases. In other words, the *average* contribution level of the top tail of a distribution of potential contributors is higher, the smaller the fraction in the tail and the greater the heterogeneity of contribution sizes. The larger the average contribution is, the greater the total contribution from a fixed number of contributors. Thus, organizers with limited resources can produce more total contributions when their personal networks are large and members are more diverse in their potential contributions (around a constant mean).[16]

[16] It might seem strange to view k, the number of contributors, as relatively fixed, while the organizer's network size varies. Even more counterintuitive is that k need not change if we permit the total number of persons in the interest group to vary. One might expect that the number of contributors would have to increase with either of these factors, particularly group size. However, collective goods generally have high "jointness of supply," which means that the cost of providing the good is relatively invariant with the number that enjoys the good. Thus, treating k as relatively fixed is often appropriate. See Oliver and Marwell (1988) for an extended discussion of jointness of supply and the effects of group size on collective action, including the paradoxical finding that the number of contributors (k) may actually be smaller in larger interest groups.

Selectivity, Organizing Cost, and Network Size

Selectivity directly explains why organizing cost declines in significance as resource heterogeneity increases. As the degree of selectivity increases, a given total contribution can be produced by a smaller number of contributors. Thus, when resource heterogeneity is high, the number of others an organizer can afford to organize becomes much less significant as a constraint. If an organizer happens to have two or three big contributors in her network, she can achieve success without spending a lot of resources on organizing.

When selectivity is high, organizer network size is the overwhelmingly important factor determining success, not because it affects how many others can be organized but because it affects the likelihood that the few largest potential contributors are in the same network. Increased network centralization increases the size of the largest organizer network and thus is always beneficial; *who* the organizer at the center of this network is (in terms of the organizer's resources and interest) pales in significance when compared with her network's sheer size.

The organizer is not entirely irrelevant. As resource and interest heterogeneity increases, there is some risk that she will lack even the modest levels of interest or resources that are necessary. The reduced effect of centralization when interest heterogeneity is high captures this effect. It does not appear for resources in our simulation because the same resource distribution is used for group members in their roles as organizers and as members being organized, so the selectivity effect for members masks the effects of this risk for organizers.

General Applicability of the Results

Since our reasoning is highly abstract, it is important to identify the range of situations to which our results should apply. We can begin by stressing that organizer-centered mobilization is central to our theoretical model and our results can be assumed to apply only to such mobilizations. However, we believe that organizer-centered mobilizations are the rule rather than the exception and that *most* collective goods are produced by actions that originate with one person (or a few people) who plans a campaign and purposely seeks to draw others into it. Thus, we expect our results to be relevant to a very large proportion of actual collective action situations. Organizer-centered mobilizations presuppose some sort of interdependent decisions that make people's willingness to contribute dependent on the willingness of others. The explicit or implicit all-or-none contract is one important instance of interdependence, but our findings about centralization and organizing costs should hold true for any strategy of mobilization in which an organizer solicits participation or contributions.

Central to our results is the assumption that organizers try to maximize others' contributions, given their own resource and interest constraints. If others vary in how much they will contribute and if organizing costs are proportional to the number organized, organizers may be expected to exercise selectivity in their choices of targets and to find that the sheer size of their personal networks is crucial to their success or failure. The importance of personal network size for a single organizer translates into the group-level importance of centralization of group network ties, for centralization at the group level increases the size of the largest organizer network. We show the nonintuitive result that the effect of centralization in generating at least one large organizer network is so powerful that it dwarfs other factors that might otherwise predict the emergence of "good" organizers.

Our model assumes that all individuals make instrumental decisions, that is, that they explicitly weigh benefits and costs. However, this assumption about contributors is not necessary for our results. All that matters is that different people will potentially contribute different amounts. Their reasons for doing so are irrelevant, as is the question of whether these reasons differ for different people. In fact, it is quite possible that group members are utterly indifferent to the collective goods dilemma and contribute amounts that are solely determined by their feelings about the organizer or their ideological beliefs with respect to the issue at hand. Organizers will still exercise selectivity, and network centralization will still be salutary.

Two assumptions we make that might significantly limit the generality of our results are that organizing costs are proportional to the number organized and the closely related assumption of perfect information. When information is perfect, it is plausible to suppose that organizers will expend resources only on the "best" possible contributors, thus making organizing costs proportional to the number organized (as we assume), *not* to the total size of the organizer's network.

If the opposite holds and there is zero information, the organizer would either have to recruit randomly from his acquaintances or have to expend resources to obtain more information about their likely contributions. Such an expense would probably be proportional to the organizer's personal network size. Recruiting randomly would eliminate selectivity as a factor and thus would mean that organizer networks larger than the minimum necessary for success would have no additional advantage. Recalling our previous discussions of heterogeneity, we find that, even if recruitment were random, group network centralization (i.e., within-group heterogeneity in personal network sizes) would have a positive effect whenever the *mean* organizer network was too small to support a workable collective action. However, there would be no additional positive benefit through the selectivity effect.

The selectivity effect would be weakened but not eliminated if information were less than perfect but not zero. Partial information can be treated as a quantifiably weaker level of selectivity, which would lead us to expect that centralization would have positive but weaker effects. Following this line of speculation, we might introduce into the model a requirement that the organizer pay for all information about others. This would produce an optimization problem for organizer network size between the information cost and the benefit of selectivity (and require a more detailed analysis than we are able to provide here).

Perhaps the greatest uncertainty in setting the limits of generality of our results arises from our disregard of indirect ties. Such ties are irrelevant for the one-step mobilization process we assume. Our hunch is that centralization will prove to be beneficial for mobilizations with a single point of origin, even if two- or multistep mobilization processes through indirect ties are considered. However, nothing in our present analysis can defend this hunch. At present, we must limit the scope of the result to one-step mobilizations.

To sum up, we predict that, ceteris paribus, heterogeneous groups with more centralized network structures will be more able to sustain organizer-centered mobilizations than those with less centralized network structures and that the strength of this effect will be greater, the more information potential organizers have about the potential contribution sizes of others.

We may draw out some of the implications of our analysis by considering its relation to Granovetter's (1973) important and frequently cited analysis of the effects of strong and weak ties. Granovetter argues that strong ties tend to form cliques, while weak ties tend to bridge cliques and bring everyone into the same network, so that weak ties are a better basis for collective action. The imagery of this argument suggests that decentralized nets and centralized wheels of weak ties would be equally effective network structures for collective action. Our results imply that it is not weak ties, per se, that are useful but their tendency to be centralized. Residents who are all bridged by the same weak tie—to a parish priest, for example—are more likely to be mobilized than those linked by the same number of weak ties distributed more widely through cross-cutting associational memberships. Our claim in this regard is, of course, quite consistent with a huge social psychological literature on communication networks (for a summary, see Michener, DeLamater, and Schwartz 1986, pp. 418–21), and Granovetter would doubtless agree that the same number of ties will be more effectively employed if they are centralized than if they are decentralized.

Where we might part company is that our results suggest that the "bridging" function of weak ties is relatively unimportant for collective action. Collective action happens when a critical mass of interested and resourceful individuals can coordinate their efforts. There is generally no need for *all* the aggrieved population to be mobilized and no need for all the members of a population to be mutually reachable.

The fact that association membership promotes collective action has been said to demonstrate the importance of weak ties for providing bridges among friendship cliques. Our analysis indicates that the important structural fact about associations is that they are centralized. Since both theoretical approaches predict that organizations are important, it would be difficult to construct critical tests at an aggregate level. Case studies might indicate whether collective actions tend to be coordinated through the leaders of organizations (favoring the centralization analysis) or through nonleaders who use the organization as a recruiting ground for events outside the aegis of the organization (favoring the weak-ties analysis).

At least in principle, empirical tests could also be constructed for our predictions that organizing costs are less important constraints as group heterogeneity increases and that successful organizing is more a matter of *whom* you can mobilize than of *how many* you can mobilize. Centralization, and its effect on creating some very large organizer network, should remain important even when only a few members are mobilized because an organizer with a larger personal network is more likely to know the right few people.

REFERENCES

Aitchison, J., and J. A. C. Brown. 1963. *The Lognormal Distribution*. Cambridge: Cambridge University Press.

Alba, Richard D. 1981. "From Small Groups to Social Networks." *American Behavioral Scientist* 24 (5): 681–94.

Arabie, P. 1977. "Clustering Representations of Group Overlap." *Journal of Mathematical Sociology* 5:113–28.

Bator, Francis M. 1958. "The Anatomy of Market Failure." *Quarterly Journal of Economics* 72 (August): 351–79.

Bonacich, Philip, Gerald H. Shure, James P. Kahan, and Robert J. Meeker. 1976. "Cooperation and Group Size in the *N*-Person Prisoners' Dilemma." *Journal of Conflict Resolution* 20 (December): 687–706.

Brewer, M. B. 1985. "Experimental Research and Social Policy: Must It Be Rigor versus Relevance?" *Journal of Social Issues* 41:159–76.

Brubaker, E. R. 1982. "Sixty-eight Percent Free Revelation and Thirty-two Percent Free Ride? Demand Disclosures under Varying Conditions on Exclusion." Pp.

151–66 in *Research in Experimental Economics,* vol. 2. Edited by V. L. Smith. Greenwich, Conn.: JAI.

Burt, Ronald S. 1978. "Cohesion versus Structural Equivalence as a Basis for Network Subgroups." *Sociological Methods and Research* 7:189–212.

———. 1980. "Models of Network Structure." *Annual Review of Sociology* 6:79–141.

Chamberlain, John. 1974. "Provision of Collective Goods as a Function of Group Size." *American Political Science Review* 68 (March): 707–16.

Cohn, Samuel R. 1985. *The Process of Occupational Sex-Typing: The Feminization of Clerical Labor in Great Britain.* Philadelphia: Temple University Press.

Dawes, R. M. 1980. "Social Dilemmas." *Annual Review of Psychology* 31:169–93.

Dawes, R. M., J. McTavish, and H. Shaklee. 1977. "Behavior, Communication, and Assumptions about Other People's Behavior in a Commons Dilemma Situation." *Journal of Personality and Social Psychology* 35:1–11.

Dawes, R. M., and J. M. Orbell. 1982. "Cooperation in Social Dilemma Situations: Thinking about It Doesn't Help." Pp. 167–73 in *Research in Experimental Economics,* vol. 2. Edited by V. L. Smith. Greenwich, Conn.: JAI.

Doreian, P. 1974. "On the Connectivity of Social Networks." *Journal of Mathematical Sociology* 3:245–58.

Fireman, Bruce, and William A. Gamson. 1979. "Utilitarian Logic in the Resource Mobilization Perspective." Pp. 8–44 in *The Dynamics of Social Movements,* edited by Mayer Zald and John D. McCarthy. Cambridge, Mass.: Winthrop.

Freeman, L. C. 1979. "Centrality in Social Networks: Conceptual Clarification." *Social Networks* 1:215–39.

Frohlich, Norman, Thomas Hunt, Joe Oppenheimer, and R. Harrison Wagner. 1975. "Individual Contributions for Collective Goods: Alternative Models." *Journal of Conflict Resolution* 19 (June): 310–29.

Frohlich, Norman, and Joe A. Oppenheimer. 1970. "I Get By with a Little Help from My Friends." *World Politics* 23 (October): 104–20.

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

Hardin, Russell. 1971. "Collective Action as an Agreeable *N*-Prisoners' Dilemma." *Behavioral Science* 16:472–81.

———. 1982. *Collective Action.* Baltimore: Johns Hopkins University Press for Resources for the Future.

Head, John G. 1974. *Public Goods and Public Welfare.* Durham, N.C.: Duke University Press.

Kelley, H. H., and J. Grzelak. 1972. "Conflict between Individual and Common Interest in an *N*-Person Relationship." *Journal of Personality and Social Psychology* 21:190–97.

Marsden, Peter V. 1981. "Introducing Influence Processes into a System of Collective Decisions." *American Journal of Sociology* 86:1203–35.

———. 1983. "Restricted Access in Networks and Models of Power." *American Journal of Sociology* 88:686–717.

Marwell, Gerald, and Ruth E. Ames. 1979. "Experiments on the Provision of Public Goods. I. Resources, Interest, Group Size and the Free-Rider Problem." *American Journal of Sociology* 84 (May): 1335–60.

Marwell, Gerald, and Pamela Oliver. 1984. "Collective Action Theory and Social Movements Research." *Research in Social Movements, Conflict and Change* 7:1–28.

Messick, D. M., and C. L. McClelland. 1983. "Social Traps and Temporal Traps." *Personality and Social Psychology Bulletin* 9:105–10.

Messick, D. M., H. Wilke, M. B. Brewer, R. M. Kramer, P. E. Zemke, and L. Lui. 1983. "Individual Adaptations and Structural Changes as Solutions to Social Dilemmas." *Journal of Personality and Social Psychology* 44:294–309.

Michener, H. A., J. D. DeLamater, and S. A. Schwartz. 1986. *Social Psychology*. San Diego: Harcourt Brace Jovanovich.

Oberschall, Anthony. 1973. *Social Conflict and Social Movements*. Englewood Cliffs, N.J.: Prentice-Hall.

———. 1980. "Loosely Structured Collective Conflict: A Theory and an Application." *Research in Social Movements, Conflict and Change* 3:45–68.

Oliver, Pamela E. 1980. "Rewards and Punishments as Selective Incentives for Collective Action: Theoretical Investigations." *American Journal of Sociology* 85 (May): 1356–75.

———. 1984. " 'If You Don't Do It, Nobody Else Will': Active and Token Contributors to Local Collective Action." *American Sociological Review* 49:601–10.

Oliver, Pamela E., and Gerald Marwell. 1987. "Towards a Theory of the Critical Mass. IV. Selectivity, Group Size, Organizing Costs, and Heterogeneity in Mobilizing for Collective Action." Unpublished manuscript. Department of Sociology, University of Wisconsin—Madison.

———. 1988. "The Paradox of Group Size in Collective Action: A Theory of the Critical Mass. II." *American Sociological Review* 53:1–8.

Oliver, Pamela E., Gerald Marwell, and Ralph Prahl. 1985. "Organizer and Network Characteristics as Predictors of Collective Action through All-or-None Agreements." Paper presented at the annual meetings of the American Sociological Association, Washington, D.C.

Oliver, Pamela E., Gerald Marwell, and Ruy Teixeira. 1985. "A Theory of the Critical Mass. I. Interdependence, Group Heterogeneity, and the Production of Collective Action." *American Journal of Sociology* 91:522–56.

Olson, Mancur, Jr. 1965. *The Logic of Collective Action*. Cambridge: Harvard University Press.

Samuelson, Paul A. 1954. "The Pure Theory of Public Expenditure." *Review of Economics and Statistics* 36 (November): 387–89.

Schelling, Thomas C. 1973. "Hockey Helmets, Concealed Weapons, and Daylight Saving: A Study of Binary Choices with Externalities." *Journal of Conflict Resolution* 17 (September): 381–428.

Schofield, Norman. 1975. "A Game Theoretic Analysis of Olson's Game of Collective Action." *Journal of Conflict Resolution* 19 (September): 441–61.

Smith, Jan. 1976. "Communities, Associations, and the Supply of Collective Goods." *American Journal of Sociology* 82 (September): 291–308.

Snow, David A., Louis A. Zurcher, Jr., and Sheldon Ekland-Olson. 1980. "Social Networks and Social Movements: A Micro-structural Approach to Differential Recruitment." *American Sociological Review* 45:787–801.

Tilly, Charles. 1978. *From Mobilization to Revolution*. Reading, Mass.: Addison-Wesley.

Van de Kragt, A. J. C., J. M. Orbell, and R. M. Dawes. 1983. "The Minimal Contributing Set as a Solution to Public Goods Problems." *American Political Science Review* 77:112–22.

Walsh, Edward, and Rex Warland. 1983. "Social Movement Involvement in the Wake of a Nuclear Accident: Activists and Free Riders in the TMI Area." *American Sociological Review* 48:764–80.

White, H. C., S. A. Boorman, and R. L. Breiger. 1976. "Social Structure from Multiple Networks. I. Blockmodels of Roles and Positions." *American Journal of Sociology* 81:730–80.