Power, Identity, and Collective Action in Social Exchange*

BRENT SIMPSON, *University of South Carolina*
MICHAEL W. MACY, *Cornell University*

Abstract

Our research aims to bring collective action back into the study of structural determinants of power in social exchange. Previous research has focused primarily on the bargaining power of actors whose locations in exchange networks confer different risks of exclusion. We argue that structural position affects not only bargaining power but also the ability of low-power actors to organize against unequal bargaining power. We hypothesize that collective action among low-power actors is facilitated by identification with others who are structurally disadvantaged. We test two identity-theoretic expected utility models that specify how actors in a mixed-motive coalition game might take into account the payoffs to others in structurally equivalent positions. In the utilitarian model, actors maximize the greatest good to the greatest number. In the collectivist model, actors also seek to minimize in-group inequality. Results show some support for the utilitarian model among female participants and strong support for the collectivist model among both males and females. We speculate about causes of gender differences and identify directions for future exchange-theoretic research on social identity and socially embedded collective action.

Bringing Back Collective Action

Bertrand Russell (1938:12) regarded power as "the fundamental concept of social science," while Blau (1974:616) saw the study of social structure as the "distinctive task" of sociology. Social exchange theories link these ideas by

We wish to acknowledge the National Science Foundation (SES-9819249) for their generous support of this research. In addition, we thank Andreas Flache, Dave Willer, Robb Willer, and anonymous* Social Forces *reviewers for helpful suggestions on previous drafts. Direct correspondence to Brent Simpson, Department of Sociology, University of South Carolina, Columbia, SC 29208. E-mail bts@sc.edu.

 Social Forces, June 2004, 82(4):1373-1409

predicting interpersonal power from actors' locations in network structures (Skvoretz & Willer 1993).[1] For example, in a three-line network ($B_1 - A - B_2$), *A* has power over B_1 and B_2 because *A* has access to multiple exchange partners, each of whom has access only to *A*. But if we simply add a link between B_1 and B_2, *A* loses its structural advantage. In this triangle network, all three actors now have equal power because all are excluded with the same probability.

The predicted effects of network structure have been strongly supported in laboratory studies of bargaining behavior (Willer 1999). However, most applications are limited in scope by the assumption that low-power actors cannot act collectively to counter the structural advantage of high-power positions (Bonacich 2000). Emerson (1962) anticipated this point over four decades ago, when he argued that power inequality is "unstable for it encourages the use of power which in turn sets in motion . . . balancing operations" (34). Similarly, Thibaut and Kelley (1959:208) noted how "power provides the problem to be solved and coalition the means of solving it." These early exchange theorists (including Blau 1964) believed coalition formation among structurally disadvantaged actors to be a fundamental means through which structurally based power inequality is "balanced."

Despite the centrality of coalitions in the work of early theorists, network exchange researchers have only recently begun to address the interplay of power and coalition processes. This paper argues that it is time to bring collective action back in to the structural analysis of power.

We argue that structure affects power through two mechanisms related to collective action, one related to the *motivation* for coalition formation among low-power actors, the other to the *success* of those coalitions. The motivation for coalition formation results from the "bidding-war problem," which arises from differences in actors' exchange opportunities (e.g., the separation of the worker from the means of production). The success of coalitions, in turn, rests on solving the "free-rider problem" that arises because potential coalition members differ in their opportunities to benefit from coalition formation without shouldering part of the coalition costs. Free riding opportunities also depend on the structure of social relations, a problem overlooked in network exchange research in which collective action cannot occur.

Ironically, the free-rider problem is also overlooked in studies of social exchange in which collective action *can* occur. Cook and Gillmore (1984), for example, tested and found support for Emerson's argument that greater power imbalance would promote formation of coalitions to countervail inequality. Although important and suggestive, Cook and Gillmore's study uses an experimental design that precludes structural constraints on collective action. In the simple three-line network they studied, both low-power actors must agree to join the coalition in order to avoid a bidding war. But other networks can make collective action more problematic. Consider, for example, the M-Branch network of Figure 1, which contains low-power bidders (the Bs) competing to

FIGURE 1: The M-Branch Network

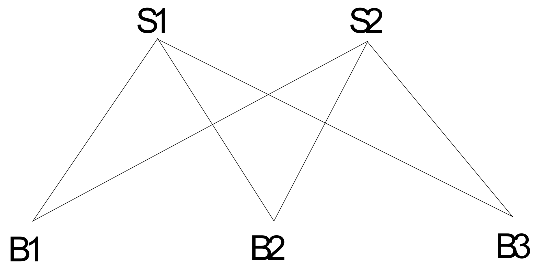

participate in exchange with two high-power sellers (Ss). If coalitions are precluded, the bidding wars among the Bs will produce exchanges that greatly favor S_1 and S_2. A coalition of Bs whose members agree to bargain as a unit will solve the bidding-war problem by transforming the network into an equal-power structure in which no actor risks exclusion. But now a new difficulty arises: the free-rider problem. With three low-power actors competing for two opportunities for exchange, only two coalition members are needed to avoid bidding wars (Willer & Skvoretz 1997). The nonmember then enjoys the same bargaining power as the coalition without having to share the "profit."[2]

More generally, a free-rider problem arises in networks in which a coalition improves the bargaining power not only of the members but also of other low-power actors who elect not to join (Bonacich 1995; Willer & Skvoretz 1997). In some networks, coalitions may actually benefit nonmembers much more than members, creating an incentive to "let George do it." But if too many low-power actors choose to free-ride, the coalition will fail, as happens sometimes in unions and cartels. In short, most previous research on power in exchange networks may have overstated power inequality by assuming that collective action cannot occur. Cook and Gillmore allowed for collective action but precluded the opportunity to free-ride in order to show that coalitions can overcome the bidding-war problem. This study may have understated power inequality by assuming that collective action cannot fail.

Recent game-theoretic studies of collective action in exchange networks show that coalition formation among the structurally disadvantaged depends decisively on the ability to overcome temptations to free-ride (Bonacich 1995; Willer & Skvoretz 1997). Our previous work (Simpson & Macy 2001) tested the ability of low-power actors to overcome the free-rider problem and form an effective coalition. Although only two members were needed for the coalition to succeed, three-member coalitions were as prevalent as two-member coalitions and, once established, were more stable than two-member coalitions. This anomaly is difficult to explain with standard game theory, because dropping out of a two-person coalition causes the coalition to fail, while dropping out of a three-person coalition does not. Why, then, were participants more likely to abandon two-person coalitions?

Research by Lawler and Yoon (1998) suggests one possible explanation. They studied how shared social identity between exchange partners affects exchange ratios in unequal-power networks. In half the conditions, participants were given information intended to activate identification with the exchange partner. As predicted by social identity theory (Tajfel 1981), participants in these conditions agreed to less unequal exchange ratios than did participants who were not led to identify with each other.

Suppose instead low-power actors identified with one another rather than with a high-power exchange partner. Bonacich and Bienenstock (1997) describe how certain exchange networks give rise to latent classes of structurally similar positions, linked not through exchange relations, as in the Lawler-Yoon study, but through a common fate in exchange relations with members of other classes. Shared fate, in turn, has been shown to be a strong basis of identity (Kramer & Brewer 1986). If identity inhibits inequality and exploitation between low- and high-power actors, then we might also expect a similar effect of identity on solidarity among low-power actors attempting to countervail unequal exchange with high-power actors.

We tested a formal specification of this identity-theoretic explanation. Like most theories of network exchange, our formal model assumes that choices are determined solely by the associated payoffs. However, based on our previous findings on coalition size, and on work by Lawler and Yoon, we relax the simplifying assumption that actors' utility functions assign zero weight to the payoffs received by others. We assume instead that utility is a weighted sum of the payoffs to self and others and explore alternative specifications of the weights suggested by social identity theory.[3] Specifically, we test two formal payoff transformation models of how identity leads actors to take into account the payoffs to others in structurally equivalent low-power positions. In the utilitarian model, actors maximize the greatest good to the greatest number. In the collectivist model, actors also seek to minimize in-group inequality. Both identity-theoretic models are tested against the conventional exchange-theoretic assumption that actors are motivated purely by self-interest.

In sum, while our previous research studied the effects of coalitions on bidding wars and power inequality, the current study focuses on the effects of social identity on the formation of coalitions despite the temptation to free-ride. To preview the study, the next section briefly summarizes research on social identity and collective action. From this summary, we elaborate alternative models of the effects of social identity on the willingness to join a coalition of other low-power actors. We then introduce the experiment designed to test the models. Results show some support for the utilitarian model and strong support for the collectivist model, but the results are conditioned by gender. The article concludes with a discussion of these gender effects based on a comparison with findings from related work. The conclusion also outlines avenues for future research suggested by these results.

Social Identity and Collective Action

This section begins with a brief review of social identity theory, focusing on Tajfel's (1981) fundamental distinction between personal and social identity, and how the theory applies to collective action. Social identity is defined as "that part of an individual's self-concept which derives from his knowledge of his membership in a social group (or groups) together with the value and emotional significance attached to that membership" (Tajfel 1981:225). In social identity theory, the distinction between "we" and "others" comes from cognitive self/other categorization that induces "a perceptual accentuation of intragroup similarities and intergroup differences" (Turner & Onorato 1999:21).

Personal identity, on the other hand, "is highlighted by thinking of the self in terms of unique attributes" (Deaux 1996:780). Personal identity therefore accentuates the distinction between individual and collective interests and thus activates a self-interested desire to maximize individual payoffs, even when this necessitates a reduction in the payoffs to others. In contrast, social identity blurs the distinction and thus mitigates the free-rider problem in collective action. Citing parental bonding and patriotism as examples, Coleman (1990:158) characterized identification as the process through which "one actor has adopted, or taken up, the other's interest." In-group cooperation then results "from the shared and mutual perception by in-group members of their interests as interchangeable" (Turner et al. 1987:65). For this reason, Brewer and Silver (2000:154) argue, identity is a "group resource that is critical to the ability of the group to mobilize collective action among its members or to recruit members into a social movement." A number of empirical studies support this general point (Brewer & Kramer 1986; Dawes et al. 1988; de Weerd & Klandermans 1999; Kollock 1998; Kramer & Brewer 1984, 1986).

At issue is the process through which social identity influences participation in collective action. Following Brewer and Silver (2000:160), we assume that, through identification with significant others, "the meaning of 'self-interest' is transformed to the group level" such that "group welfare becomes a part of the rational calculus by which an individual evaluates the costs and benefits of intended actions and potential outcomes." In the sections to follow, we consider two models of payoff transformation suggested by social identity theory. The utilitarian model maximizes the mean of in-group payoffs ("the greatest good to the greatest number"). The collectivist model strikes a balance between maximization of the mean and minimization of in-group inequality.[4]

Social Identity As Utilitarianism

Brewer and Silver (2000) posit a utilitarian effect of identity on cooperation. The utilitarian model assumes that social identity leads to a "transformation of social motives whereby an individualistic self-interested orientation is matched or superseded by the motivation to maximize joint or collective interests" (161). More formally, let *i* denote an in-group member. Then, according to the utilitarian model, the subjective utility *i* derives from a given payoff *P* is simply the mean payoff for the in-group, or

$$U_i = \frac{\sum_{j=1}^{N} P_j}{N} = \overline{P}, i \in j \tag{1}$$

where U_i is *i*'s utility, P_j is the objective payoff to any group member $j(i \in j)$, and *N* is the number with whom ego identifies (where $N = 1$ when identity is personal and *N* is the number of in-group members when identity is social). Note that when personal identity (self-interest) is salient, $U_i = P_i$ and the actor's subjective and objective payoffs are identical — the payoffs are not transformed.

Social Identity As Collectivism

Most applications of social identity theory to cooperation in social dilemmas assume some version of the utilitarian model, in which actors are motivated by a weighted sum of the payoffs to self and significant others, without regard to in-group differences in payoffs (Brewer & Silver 2000; Kollock 1998). However, this indifference to in-group inequality is inconsistent with a central tenet of social identity theory — that in-group members seek *maximal distinctiveness*. According to the maximal distinctiveness hypothesis, social identity leads actors to make choices that maximize the differences between the in-group and the out-group while minimizing within-group differences (Hogg 1996; Turner 1985). That the motivation to maximize group distinctiveness applies to resources, as well as attributes, has been shown repeatedly in minimal-group experiments (Tajfel 1982).[5]

Despite the centrality of the maximal distinctiveness hypothesis to social identity theory, it has remained curiously absent from research on identity and cooperation in social dilemmas, where the focus has remained almost exclusively on the effort to maximize in-group interests. Applied to the problem of coalition formation among low-power actors, within-group homogenization implies an effort to minimize the inequality of payoffs to coalition members (the in-group). When coupled with a motivation to increase within-group welfare (as assumed by the utilitarian model), between-group differentiation implies an effort to maximize the mean of the payoffs to coalition members, which in turn entails minimization of payoffs to high-power actors (the out-group). The effort to increase in-group payoffs while minimizing in-group inequality corresponds to a collectivist image of a unitary actor, in contrast to the utilitarian idea of a mere aggregation of individuals (Leung & Bond 1984; Triandis et al. 1985). As Deaux and Reid (2000:186) write, "collectivism implies an emphasis on group cohesion, common fate, distinction from out-groups, and shared norms and standards." Recent research (Brown et al. 1992) supports this collectivist effect of identity.

We implemented a collectivist model using the mean payoff to capture maximization of in-group payoffs (as well differentiation with the out-group) and the mean absolute difference in payoffs to operationalize inequality.[6] This gives the subjective utility to any in-group member i as

$$U_i = \overline{P} - \frac{\sum_{i=1}^{N} \sum_{j=1}^{N} \left| P_j - P_i \right|}{N(N-1)}, N > 1 \tag{2}$$

where U_i is i's utility, P_j is the objective payoff to any group member $j(i \in j)$, and N is the number with whom ego identifies. When personal rather than social identity is salient, $N = 1$ and in-group differences are not possible; the collectivist model then implies $U_i = P_i$.

Application to Coalitions in Network Exchange

These models of social identity make competing predictions about participation in collective action among low-power actors in exchange networks. As detailed below, the predictions of both models also differ from the baseline assumption, normally employed in exchange network research, namely, that actors are motivated by self-interest. We tested these predictions using the M-branch network (Figure 1) in which three low-power bidders (Bs) compete to be included in an exchange with one of two high-power ellers (S). The Bs could benefit from forming a coalition but they could benefit even more by free riding on the bargaining power they gain if enough others form an effective coalition.[7]

TABLE 1: Payoff Matrix for Low-Power Actors in the M-Branch Network When Coalitions Are Possible, with Probability of Cooperation As Smoothed Best-Reply

Number of Cooperative Alters (B_1 and B_2)	B_3's Strategy C	B_3's Strategy D	Probability of Cooperation
0	**4** 4, 4	**4** 4, 4	.50
1	**6** 11, 6	**4** 4, 4	.73
2	**8** 8, 8	**11** 6, 6	.31

Note: Values given in bold are payoffs to B_3, the focal actor. B_1, B_2, and B_3 are as defined in Figure 1.

To illustrate, assume that each S-B exchange relation negotiates the division of 24 points and that each position is allowed only one exchange during any given round of negotiation. In the absence of an effective coalition, the three Bs must bid against each other for inclusion in two exchanges, with one actor excluded and earning 0. Although competition to avoid exclusion can drive bids up to the maximum of 23 points, actual exchange ratios rarely reach such extremes (Bonacich & Friedkin 1998; Cook & Emerson 1978; Cook et al. 1983). Two decades of research on strong-power networks like the M-branch show that exchange ratios average 18:6 in favor of the high-power actor (Bonacich 2000; Markovsky, Willer & Patton 1988; see Simpson & Macy 2001 for a formal derivation of these bidding-war payoffs). That leaves the low-power actor with 6 points. With three low-power actors competing for inclusion in two exchanges, each has a two-thirds chance that its 18-point offer will be accepted. Thus, in the absence of an effective coalition, the expected payoff is two-thirds of 6 points, or 4 points.

If, on the other hand, all three low-power actors join the coalition, they will have no risk of exclusion and will thus enjoy equal bargaining power with the high-power actors, leading to equal exchange ratios (Cook & Gillmore 1984; Willer & Skvoretz 1997). Thus the three coalition members earn 12 points in each of two exchanges, for a total of 24 points. Assuming coalition points are divided equally among the three members, each receives a payoff of 8 points.

Now suppose one of the three low-power actors decides to defect from the coalition. The two coalition members make identical 12-point offers. However, the nonmember must offer 13 (earning 11) in order to avoid exclusion. The coalition members have no incentive to improve one of their 12-point offers

TABLE 2: Utilitarian Payoff Transformation of Table 1 Payoffs

Number of Cooperative Alters (B_1 and B_2)	B_3's Strategy C	D	Probability of Cooperation
0	4	4	.50
1	7.7	4	.83
2	8	7.7	.53

Note: Payoffs are to B_3, the focal actor. B_1, B_2, and B_3 are as defined in Figure 1.

because they will then be effectively undermining their other 12-point offer. They are better off to participate in a single 12–12 exchange and allow the defector to earn 11 (see Simpson & Macy 2001 for empirical support).[8] The coalition thus earns a total of 12 points, which is divided equally between the two members, giving each a payoff of 6 points. Although the payoff of 6 to each of two coalition members is not as good as the 8 points they earn in a three-person coalition, it is still much better than the 4-point payoff they would receive if the coalition were to fail.

Table 1 summarizes these payoffs in matrix form, where the columns display payoffs depending on whether the focal actor decides to join or not join the coalition and the rows give the total number of alters (other low-power actors) who join. These payoffs show that, for the M-branch network, a coalition of two members is a strict Nash equilibrium (i.e., a set of strategies such that any unilateral deviation leads to a lower payoff). For the two-person coalition, the nonmember is worse off by joining the coalition (the payoff will decline from 11 to 8) and any member is worse off by unilaterally abandoning the coalition (the payoff will decline from 6 to 4).[9] Assuming the actors are self-interested and able to coordinate a pure-strategy equilibrium (one in which each actor always plays the same strategy), we can expect to observe stable coalitions of two members. Three-member coalitions should be unstable.

In standard game theory, the Nash equilibrium assumes that ego's utility is solely a function of ego's payoffs. This is also the assumption in almost all network exchange research. In deciding whether to accept an offer, ego takes into account only the payoff to ego and ignores the payoff to alter. We continue to make that assumption in the model of the bargaining relation between high- and low-power actors. However, we relax that assumption in the model of the decision whether to join a coalition with other actors in a structurally equivalent position.

The two identity-theoretic models in equations 1 and 2 assume instead that ego takes into account the payoffs to alters with whom ego identifies. Applying the utilitarian transformation to the objective payoffs of Table 1 produces the subjective utilities given in Table 2. For simplicity, Table 2 displays only the

TABLE 3: Collectivist Payoff Transformation of Table 1 Payoffs

Number of Cooperative Alters (B_1 and B_2)	B_3's Strategy C	D	Probability of Cooperation
0	4	4	.50
1	4.4	4	.56
2	8	4.4	.81

Note: Payoffs are to B_3, the focal actor. B_1, B_2, and B_3 are as defined in Figure 1.

utilities for focal actor B_3. Because the utilitarian transformation takes the average of in-group members' objective payoffs, the values of Tables 1 and 2 differ only for coalitions that produce unequal payoffs to low-power actors (i.e., two-person coalitions). When any two Bs cooperate and the third defects, the defector receives 11 and each cooperator receives 6. Applying the utilitarian transformation function to these objective payoffs gives $U_{B3} = 7.7$ for any two-person coalition, whether or not B_3 is a member.

The collectivist model in equation 2 transforms the objective payoffs of Table 1 into the subjective utilities given in Table 3. The values of Tables 1 and 3 differ only for two-person coalitions, which produce unequal objective payoffs to in-group members. When any two Bs cooperate and the third defects, the defector receives 11, and each cooperator receives 6. Applying the collectivist function for $i = B_3$ gives $U_{B3} = 4.4$ for any two-person coalition, whether or not B_3 is a member. For all other coalition sizes in the M-branch, the subjective utilities of Table 3 are equivalent to the objective payoffs of Table 1, as well as the utilitarian subjective utilities given in Table 2.

From Utilities to Cooperation

Game theory typically assumes deterministic choices based on strict preference, that is, ego will always choose the better alternative, even if both choices have nearly identical utility. Harsanyi (1973) has proposed a more plausible theory, based on smoothed best response, in which the larger the difference in utility across alternatives, the greater the probability that a player will choose the strategy with the higher utility. Following Harsanyi, Appendix A specifies functions (see equations A1 and A2) for translating the utilities of Tables 1–3 into expected rates of cooperation. Table 1 gives the utilities and corresponding choice probabilities for the self-interest (personal identity) model in which utilities equal the untransformed payoffs, with $p = 0.5$ for zero cooperative alters, rising to $p = 0.73$ for one, and then falling to $p = 0.31$ for two. We do not expect the self-interest model to hit these marks, but we should at least expect to see a sharp decline in cooperation as the expected number of cooperative alters increases from one to two.

This self-interest prediction provides a benchmark against which to measure the effects of social identity on the willingness to join the coalition. The far right column of Table 2 reports the probabilities of cooperation estimated from the utilitarian's subjective utilities using equations 1, A-1, and A-2. When no alters join, the subjective payoff for cooperation and defection are equal (4 points). Thus B_3 is indifferent between cooperation and defection (p_{B3} = .5). When one alter cooperates p_{B3} = .83, and when two alters cooperate p_{B3} = .53.

The utilitarian model generates a larger probability of cooperation for ego when one alter joins the coalition than when two others join, for the following reason: B_3 can greatly improve the payoffs to *all* in-group members by joining with one alter to form a two-person coalition. On the other hand, when two alters join, the subjective utility for cooperation (U_i = 8) is only slightly larger than the subjective payoff for defection (U_i = 7.7). That is, a third cooperator has only a small impact on the average earnings of in-group members. Thus, the utilitarian model predicts a lower probability of cooperation when two alters join than when one alter joins.

Comparison of Tables 1 and 2 shows no difference in the prediction of a nonmonotonic effect of the number of cooperative alters. The key difference is that the utilitarian model predicts a main effect of social identity, such that the rate of cooperation is higher when identity is social than when it is personal.

The estimated probabilities of cooperation for the collectivist model, formalized in equations 2, A-1, and A-2 are given in the rightmost column of Table 3. When one other joins the coalition, p_{B3} = .56 and p_{B3} = .81 when two others join. Thus, in striking contrast to both the self-interest and the utilitarian models, the collectivist model yields a much greater likelihood of cooperation when two alters cooperate than when one alter cooperates. Although the *average* payoff to in-group members is relatively high for two-person coalitions, the degree of in-group inequality is also high. Compared to the utilitarian model, the collectivist model predicts a reluctance to join a coalition of one other, even though joining one other would increase both coalition success and ego's personal (objective) payoff. This reluctance stems from the relatively high degree of in-group-inequality that results when any two actors form a coalition and the third free-rides. Because the collectivist values in-group equality, outcomes that yield unequal payoffs (e.g., two-person coalitions) are discounted. On the other hand, when all three Bs join the coalition, all receive relatively high and equal payoffs. Therefore, when two others join the coalition, the effect of social identity on cooperation is greater for collectivists (who value in-group equality) than for utilitarians (who do not).

FIGURE 2: Predicted Rates of Cooperation for the Utilitarian Model

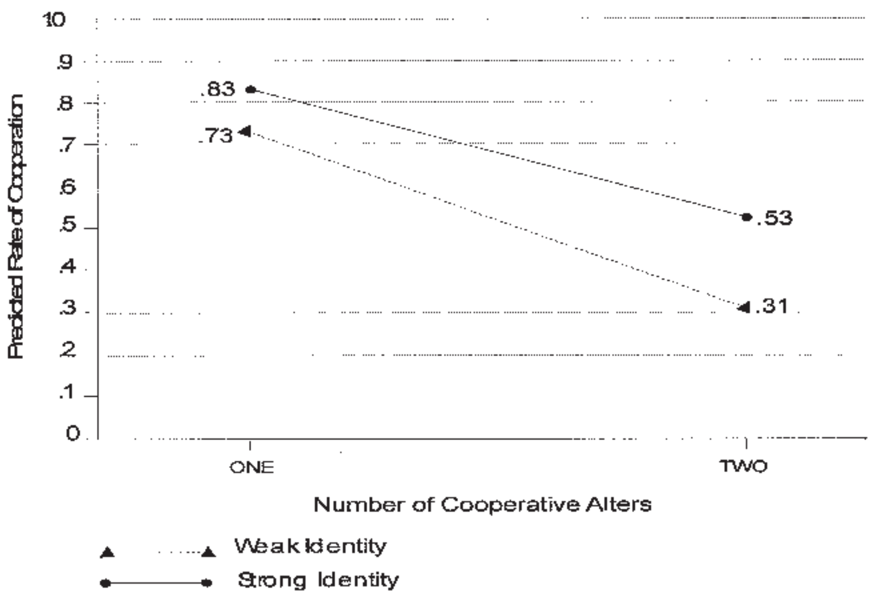

FIGURE 3: Predicted Rates of Cooperation for the Collectivist Model

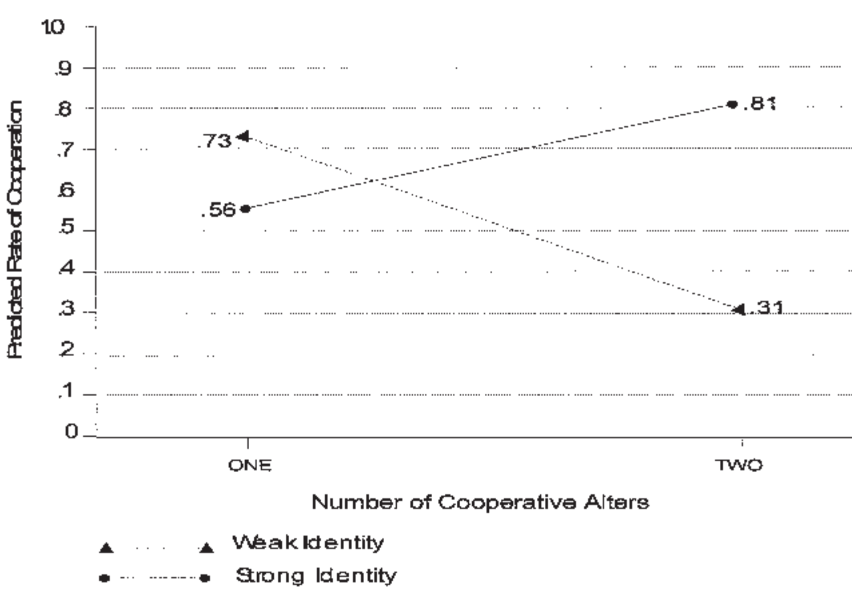

TABLE 4: Predictions of Three Models about the Effects of Identity and Alters' Cooperation on the Probability of Ego's Cooperation

	Self-Interest	Utilitarian	Collectivist
Alters' cooperation	(–)	(–)	0
Social identity	0	(+)	0
Interaction	0	0	(+)

Hypotheses

Figures 2 and 3 display the predicted rates of cooperation for the utilitarian and collectivist models, respectively. Each figure graphs the probability of cooperation (Y-axis) against the number of cooperative alters (X-axis). The dashed lines give predicted cooperation rates when personal identity is salient (i.e., when objective payoffs are not transformed). The plotted values, which are identical for that case in Figures 2 and 3, are taken directly from the rightmost column of the objective payoff matrix given in Table 1. The solid lines give predicted probabilities of cooperation when social identity is salient, as transformed by the utilitarian (Figure 2) and the collectivist (Figure 3) utility functions. These values are taken from the rightmost columns of Tables 2 and 3, respectively.

For ease of presentation, Table 4 organizes these derivations of the self-interest, utilitarian, and collectivist models as a series of three hypotheses, based on the main effects of (1) alters' cooperation, (2) social identity, and (3) their interaction.

Hypothesis 1: Ego's cooperation rate decreases with the number of cooperative alters.

Hypothesis 1 follows from the self-interest and utilitarian models, as illustrated by the negative slopes of both lines of Figure 2. In contrast, the collectivist model predicts no main effect for alters' cooperation, as shown by the opposite sign of the slopes in Figure 3. The utilitarian model also predicts a main effect for social identity, as given by hypothesis 2:

Hypothesis 2: Ego's cooperation rate increases with the strength of social identity.

Hypothesis 2, which corresponds to the utilitarian model, is indicated by the vertical distance between the two lines in Figure 2. In contrast, the self-interest model predicts no main effect. If coalition behavior is motivated strictly by self-interest, social identity will not result in a transformation of outcomes and thus will have no effect on cooperation.

While the self-interest and utilitarian models predict identical effects for alters' cooperation, the self-interest and collectivist models give convergent predictions about the main effects of social identity, but for very different reasons. As shown in Figure 3, the collectivist model predicts a positive effect of social identity when two alters cooperate but a *negative* effect when one cooperates, due to the increase in inequality. Simply put, both main effects are suppressed by their interaction, as formalized in hypothesis 3:

> *Hypothesis 3:* The effect of number of cooperative alters depends on the strength of social identity.

Because the self-interest model assumes no transformation of objective payoffs, it predicts changes in ego's cooperation on the basis of alters' behavior, an effect that does not depend on whether ego identifies with other low-power actors. Similarly, as shown in Figure 2, the utilitarian model predicts that alters' cooperation and social identity have only additive effects on ego's cooperation.

In contrast, Figure 3 shows the strong interaction effect predicted by the collectivist model. The predicted interaction reflects the value the collectivist places on in-group equality, which the utilitarian and self-interest models do not share. In the M-branch, when two actors cooperate and a third defects, the level of in-group inequality is greatly increased. On the other hand, when all three actors cooperate, all receive relatively high *and* equal payoffs.

By testing the divergent predictions outlined in Table 4, we can determine whether identity affects cooperation, and, if so, whether the effects reflect a concern for group *welfare* (hypotheses 1 and 2), as suggested by most previous social dilemma research, or group *distinctiveness* (hypothesis 3).

Gender, Identity, and Cooperation

Gender differences are relatively uncommon in network exchange experiments (see Molm 1988). However, our design differs from almost all these studies in that it poses a social dilemma. Previous social dilemma research suggests the need to block participants on gender. Kramer and Brewer's (1984) study of social identity and cooperation in social dilemmas found that males were more responsive than females to a manipulation of social identity. Another social dilemma study (Simpson 2003) tested gender differences in cooperation when the motivation for defection is decomposed into fear (of being suckered by a partner who is expected to defect) and greed (the temptation to sucker a partner who is expected to cooperate). Results confirmed that males respond more noncooperatively than females to the greed component. The social dilemma posed by coalition formation in networks like the M-branch entails greed but not fear. If no one else joins the coalition, ego's payoff is the same whether or not ego cooperates. Hence, there is no fear of being suckered. However, a low-power actor who expects both alters to join the coalition is better off free riding. The actor who does so is responding to the greed component.

Thus, while our main theoretical interest is the effect of social identity on collective action among structurally disadvantaged actors, previous work on social identity and social dilemmas suggests the need to block participants on gender. We therefore consider coalition behavior separately for males and females.

Experimental Design

Participants were undergraduates at a large private university, recruited via flyers advertising the opportunity to earn $10 to $20 for participation in a one-hour experiment. We emphasized participants' self-interest in earning money in order to satisfy the scope conditions of the game-theoretic predictions and to provide a strong test of the identity-theoretic alternatives. As detailed below, the instructions repeatedly stressed that each participant's cash reward would be determined by the total number of points that participant accumulated during the experiment. A total of 114 students (50 males and 64 females) participated.

Settings and Procedures

Participants were scheduled in groups of four.[10] Upon entering the laboratory, each participant was escorted to a private station equipped with a computer. Participants were aware that they were in a group but never actually met one another. Online instructions informed participants that they would not see other participants at any point during or after the experiment. The instructions then described the experiment in detail, followed by several practice sessions and an on-line quiz that assessed participants' understanding of the procedure. Each incorrect response was followed by a detailed explanation of the correct answer. As explained below, very few participants had problems understanding the procedures and most reported being highly involved in the experiment.

The network used was the M-branch of Figure 1.[11] Although participants were led to believe that the other network positions were occupied by students, in reality all but the participant's position were occupied by computer-simulated actors. Computer simulants were needed in order to manipulate the number of other low-power actors who join the coalition. The four students in each scheduled group actually participated in independent, randomly assigned, conditions.

For clarity, the instructions referred to the network as a market, and to network positions as sellers and bidders (cf. Bonacich 2000; Kollock 1994; Lawler & Yoon 1998). Positions in the M-branch labeled S and B were sellers and bidders, respectively. Each participant was led to believe that she had been randomly assigned to the role of B_3. The instructions introduced the task as follows:

> Bidders and sellers will earn points by bargaining with others in the game. Your objective as a bidder is to "exchange" with sellers. Exchange occurs when a seller accepts a bidder's offer. When a bidder and seller exchange, they both earn points. Because you are paid according to the number of points you earn, the more exchanges you make (and the more points you get in each exchange), the more money you earn.

The instructions then explained how to make offers, complete exchanges, and observe bargaining among others. Participants had full knowledge of the value of others' offers and exchanges.

To be consistent with related research (Simpson & Macy 2001; Willer & Skvoretz 1997), only bidders could make offers.[12] Each bidder was allowed to make as many offers as desired until one of the following occurred: the bidder's offer was accepted, both sellers accepted offers from the other two bidders, or time expired. If the participant's offer was not accepted, she was excluded from exchange in that round.

The (actual or simulated) occupant of each position was limited to one exchange on any given round. Since there were three bidders and two sellers, at least one bidder was excluded (and earned 0) on any given round of negotiations. This structure was designed to create bidding wars, such that exchange ratios would strongly favor sellers. We used these power inequalities to motivate low-power actors' option to form a coalition in the second stage.

A Two-Stage Experiment

The experiment was divided into two games, each of which lasted 15 rounds. In game 1, bidders were not allowed to form coalitions. After game 1, the computer informed participants that bidders would now have the option to form a bargaining coalition. The instructions explained the coalition option as follows:

> At the beginning of each round, each bidder will be asked if he or she wishes to join the coalition for that round. If a coalition forms, each member will vote on the bid he or she wants the coalition to send to sellers. Each member gets one vote, the average of which is the coalition bid. The coalition bid will be sent by all members and cannot be changed during the round. All points earned by the coalition will be divided equally among coalition members.
>
> Bidders who do not join the coalition will not vote on the coalition bid, nor will they be required to send the coalition bid. (Bidders who do not join will negotiate just as in game 1.) Points earned by bidders who are not members of the coalition will not be shared with others. It takes at least two bidders to

form a coalition. If less than two bidders vote to join, no coalition will exist for that round.

The instructions explained in detail the rules for allocating points earned by a coalition, illustrated with hypothetical examples. However, at no point did the instructions suggest that coalition outcomes were better or worse than outcomes in which coalitions did not form.

At the beginning of each round of game 2, the participant was asked whether she wished to join the coalition for that round. The participant and other potential coalition members ostensibly decided whether to join the coalition simultaneously. If the participant joined, she was told which others joined (if any), and prompted to submit a vote for the coalition bid. (Bidders who did not join the coalition did not vote, nor did they observe the voting process.) The average of coalition votes was the coalition bid. The coalition bid was sent automatically from each coalition member to each seller and could not be changed during a round. In other words, coalitions did not bargain as a single unit; rather, the coalition prevented bidding wars by simply restricting the number of offers, a much simpler way to accomplish the same end.

During the bidding, the computer screen displayed a list of the coalition members and the coalition bid, whether or not the participant was a member. When only one bidder voted to join the coalition, the bidder was not bound by the coalition rules. However, the participant's screen identified the bidder as a coalition member. This allowed the participant to see which others were willing to join a coalition for that round and to signal to other bidders a willingness to join.

Simulated Actor Strategies

Appendix B gives a detailed account of the simulated actors' strategies employed in games 1 and 2. The strategies were designed to appear highly realistic and to achieve two necessary objectives: (1) to generate bidding wars (which always occur in game 1, and also in game 2 if fewer than two low-power actors join the coalition) and (2) to give experimental control over the number of others who joined the coalition, thereby allowing a more parsimonious test of the hypotheses than would be possible if other positions were occupied by human participants. Manipulation checks (reported below) indicate that both objectives were fully realized.

Experimental Manipulations and Dependent Measures

The experiment used a two by three factorial design (one or two cooperative alters and three levels of identity: in-group, out-group, and none). Both factors were between-subject and were fully crossed. In addition, the multiple rounds in each stage allowed repeated within-subject measures of participant behavior.

TABLE 5: Schedule of Alters' Coalition Behavior in Game Two, by Treatment Condition

Round	One Alter Joins B_1	One Alter Joins B_2	Two Alters Join B_1	Two Alters Join B_2
1–2	C	D	C	C
3	D	C	C	D
4	C	D	D	C
5	D	D	C	D
6–15	C	D	C	C

Factor 1: Cooperative Alters

The first factor was the number of simulated bidders that typically joined the coalition. In the one-cooperator condition, B_1 almost always joined and B_2 almost always defected, with the exceptions noted in Table 5. In the two-cooperator condition, B_1 and B_2 always joined the coalition, with exceptions noted in Table 5. In both conditions, the exceptions occurred early in the game, in rounds 3 and 5, to suggest simulant exploration prior to fixing on a preferred strategy (see Appendix B). We analyzed participant behavior only in rounds 8–15, after participants were fully aware that the other players were locked into permanent roles as joiners or free-riders.[13]

We excluded the case of both simulants settling into a pattern of free riding because participant responses to this condition would provide no information that could be used to test our hypotheses. When no alters join, an actor should be indifferent between joining and not joining the coalition, given the inability of one-person coalitions to control bidding wars, as supported by our earlier study (Simpson & Macy 2001). Accordingly, all three models (self-interest, utilitarian, and collectivist) give identical predictions for cases in which no alters join the coalition. The predictions differ only for cases in which at least one alter joins the coalition.

Factor 2: Social Identity

We tested competing explanations for cooperation by manipulating participants' identification with incumbents of structurally equivalent positions in game 2. Shared fate (Sherif et al. 1961) has been shown to be a highly salient basis of identity, but it is very difficult to design an experiment that can prevent this effect from being confounded by self-interested responses to the incentive structure of the game. Therefore, in order to tease apart the effects of identity and incentives, we manipulated an attribute orthogonal to incentives to strengthen or dampen identification with occupants of structurally equivalent positions.

Following a standard procedure (Kollock 1998), we varied the information given to participants about the school affiliation of the four other actors in the game. At the beginning of the experiment, participants in all conditions read "Students from [Large Private] University and [Small Private] College are participating in the study," but were given no other information until the conclusion of Game 1. (Actual school names were given to the participants but are deleted here.) Mentioning school affiliation early reduced the risk of later statements surprising participants. In addition, this statement made others' identity ambiguous in the no-identity conditions. After game 1, all participants read a statement that identified the roles they would play in game 2. This statement was modified slightly, in accord with the identity condition. The following passage identifies the modifications, with <u>in-group information</u> underlined, and the [out-group information] bracketed:

> We are now ready to begin the second game, which will differ slightly from the first. Before the first game, you and two <u>LPU</u> [SPC] students were assigned to bidder roles, and two <u>SPC</u> [LPU] students were assigned to seller roles. In game 2, each participant will continue to occupy the same role s/he occupied in the first game. But in each round of game 2, bidders (You, B_1 and B_2) will be able to form a "bargaining coalition."

Alters' school affiliation was not referred to at any other point during the experiment.

In the in-group condition, participants decided whether to join a coalition with other bidders ostensibly from their own school (or in-group) against sellers from another local school (out-group). In this case, school-based and network-based identities were aligned. In the out-group condition, participants decided whether to join a coalition with bidders ostensibly from another local school (out-group) against sellers from their own school (in-group). In these conditions, school-based and network-based identities were in conflict. As a manipulation check on the effects of school affiliation (discussed below), we also included a no-information condition that eliminated reference to others' school affiliation before game 2.

The two fully crossed factors yielded six treatment groups. In each group, the primary dependent measure was whether the participant voted to join the coalition in each of 15 rounds of game 2. We also measured exchange ratios for games 1 and 2, but these were used primarily as manipulation checks (reported below).

After the 15 rounds of game 2, each participant was paid according to the number of points he/she accumulated during the study and then asked to complete a post=treatment survey. Survey responses identified six participants who failed to understand the instructions or who expressed suspicion about the use of computer simulants. Data for these six participants were therefore isolated, leaving 108 participants (59 females and 49 males).[14] A few weeks after the experiment, we mailed all participants a detailed explanation of the study, including all deceptive aspects.

Results

We report results in two parts, beginning with a discussion of the manipulation checks. We then report results from game 2 that test the hypotheses.

Manipulation Checks

A fair test of the hypotheses requires that two conditions be realized. First, we know from Cook and Gillmore's (1984) experiment that coalitions do not form under equal power conditions. Thus, in the absence of coalitions, bidding wars must produce sufficient power inequality to motivate coalitions. We used exchange ratios for game 1 to determine whether bidding wars produced the power inequalities needed to motivate coalition formation in game 2. Results show this condition was realized. The average exchange ratio for participants in game 1 was 17.9:6.1, favoring high-power actors, virtually identical to the ratio of 18:6 observed in previous studies of strong-power networks. It is important to note that these payoffs did not differ across low-power positions (i.e., between participants and simulants); nor did rates of exclusion.

Second, two-person and three-person coalitions must produce the expected exchange ratios with high-power actors, such that someone who free rides on a two-person coalition earns a higher payoff than does a member of a three-person coalition. The exchange ratios reported in Appendix C show that the observed payoff structure corresponded to the expected payoff matrix in Table 1. Thus, if we find that three-person coalitions are preferred over two-person coalitions, it cannot be because participants earned more by joining than by free riding.

As an additional check on the payoff structure required for the coalition game, we tested to make sure that shared school affiliation did not affect bargaining in the out-group condition (when participants shared school affiliation with high-power actors but not with other low-power actors). Lawler and Yoon (1998) found that exchange partners who were led to identify with each other agreed to less unequal exchange ratios than did partners who did not identifty with each other. It was therefore important for our study that participants not identity with high-power actors in the out-group condition when they shared the same school affiliation, because this might alter bargaining behavior and distort the payoff schedule assumed in the predictions about coalition formation.

TABLE 6: Analysis of Variance in Cooperation, by Alters' Cooperation, Social Identity, and Gender

	In-group and Out-group		In-group Only	
	$F(1, 108)$	p	$F(1, 108)$	p
Gender (G)	.544	.46	.501	.48
Alters' Cooperation (AC)	4.972	.03	6.660	.01
In-group (IN)	5.306	.02	5.562	.02
Out-group (OUT)	.256	.61		
G × AC	.687	.41	.180	.67
G × IN	1.497	.22	.645	.42
AC × IN	1.708	.20	3.107	.08
G × OUT	1.021	.32		
G × AC × IN	3.655	.06	4.060	.05
G × AC × OUT	.093	.76		

To check this, we compared exchange outcomes in the out-group and no-information conditions. According to social identity theory, when social identity is salient, actors suppress perceptions of in-group differences. If (as we suspect) shared fate as low-power actors is more salient than school affiliation, participants can be expected to tune out information about school affiliation and continue to identifty with other low-power actors even when they are from another school. In this case we should observe similar exchange outcomes in the out-group and no-information conditions. However, if school affiliation is more salient than shared fate, participants should ignore in-group differences in structural location and, in concord with Lawler and Yoon (1998), we should observe differences between exchange ratios in the no-information conditions and out-group conditions (where participants shared school affiliation with high-power actors).

Results strongly support the validity of the design. Game 2 exchange outcomes were nearly identical in the out-group and no-information conditions, regardless of coalition size. If participants in the out-group conditions had identified with high-power actors, payoffs from these conditions should have differed from the payoffs in the no-information conditions. That no comparison resulted in significant differences suggests that school affiliation was not a salient basis of identity between incumbents of low-and high-power positions in the out-group conditions.

As a final check on the design, we tested whether the out-group and no-information conditions differed in the rates of cooperation by participants. If we observe lower cooperation rates in the out-group condition (when affiliations of low-power actors differ) than in the no-information condition, we would be led to suspect that participants identified with high-power actors in the out-group condition, and the effects on bargaining behavior might then

TABLE 7: Analysis of Variance in Cooperation by Alters' Cooperation and Social Identity, for Females and Males

	Females				Males			
	In-group and Out-group		In-group Only		In-group and Out-group		In-group Only	
	$F(1, 59)$	p	$F(1, 59)$	p	$F(1, 49)$	p	$F(1, 49)$	p
Alters' cooperation (AC)	5.465	.02	5.123	.03	.845	.36	2.061	.16
In-group (IN)	6.898	.01	5.670	.02	.524	.47	1.073	.31
Out-group (OUT)	1.299	.26			.113	.74		
AC × IN	.203	.65	.036	.85	4.659	.04	6.326	.02
AC × OUT	.214	.65			.000	1.00		

violate the assumptions on which the predicted game payoffs are based. Again, results are highly reassuring. We found almost identical rates of cooperation in the out-group and no-information conditions, regardless of coalition size, a finding also supported in the analyses below.

Gender, Utilitarianism, and Collectivism

We analyzed cooperation rates for the final 8 of the 15 rounds of game 2. The analyses do not include the initial 7 rounds because, as detailed in Appendix B, the simulated buyers were programmed with a mixed strategy during rounds 1 to 5 of game 2. We therefore needed to give participants sufficient time to form strong expectations of alters' cooperation after the simulants had settled on a fixed strategy. We also wanted to give participants sufficient time to experience the effect of coalitions on bargaining. Extensive robustness checks (using still fewer rounds) give almost identical results.

Table 6 reports overall differences in coalition behavior in a 2 (gender) × 2 (cooperative alters: one vs. two) × 2 (a dummy variable, with the in-group coded as 1) × 2 (a dummy variable, with out-group coded as 1) ANOVA. As shown in the first column, the initial analysis revealed significant main effects for alters' cooperation ($p = .03$), and in-group ($p = .02$). However, both main effects were qualified by a three-way interaction between gender, in-group, and alters' cooperation ($p = .06$).

It is important to note that the out-group dummy variable revealed no effect ($p = .61$), nor did the term interact with any other variables ($p > .30$ for all interaction terms). On the basis of these results and the manipulation checks reported in the previous section, we therefore aggregated the out-group and no-information conditions into a single weak-identity category (in which shared fate is not reinforced by school affiliation). We compare this with a

FIGURE 4: Observed Cooperation Rates among Females

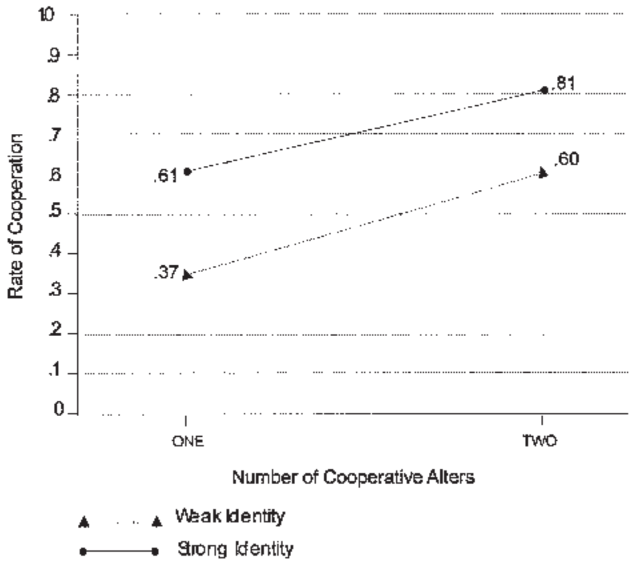

FIGURE 5: Observed Cooperation Rates among Males

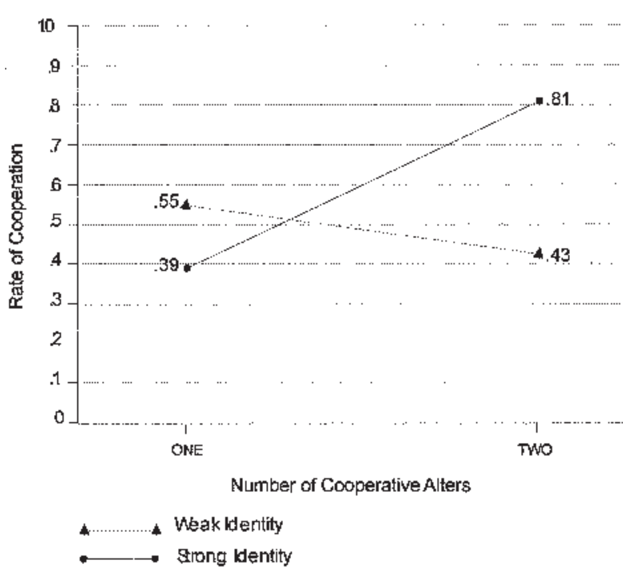

strong-identity condition in which shared fate is reinforced by shared school affiliation (the in-group condition). This facilitates a more straightforward interpretation of interaction effects. (As a robustness check, we also report results with the two conditions disaggregated).

As shown in the rightmost columns of Table 6, the aggregated data yield very similar results, with significant main effects for alters' cooperation and in-group membership, and a significant three-way interaction between gender, identity, and alters' cooperation (p = .05). We therefore conducted separate analyses for males and females.

Table 7 reports the effects of alters' cooperation and social identity separately for males and females. The results are graphically depicted in Figures 4 and 5 for females and males, respectively. As in Table 6, the results are unaffected by aggregation of the out-group and no-identity treatments as a single weak-identity condition.

Hypothesis 1, based on the self-interest and utilitarian models, predicts a decrease in the rate of cooperation by participants as the number of cooperative alters increases from one to two. Table 7 reports a significant main effect of alters' cooperation among females (p = .03). However, Figure 4 shows that the effect is in the opposite direction from that predicted by hypothesis 1. Females' cooperation rates were higher when two alters cooperated, independent of social identity. Hypothesis 1 is also not supported among males. Data from males revealed no main effect for alters' cooperation (p = .16).

Hypothesis 2, based on the utilitarian model, predicts higher cooperation when social identity is reinforced by common school affiliation than when it is not. The self-interest and collectivist models predict no main effect. As shown in Table 7 (second column) and Figure 4, results from females yield strong support for this prediction. Cooperation rates were higher in the strong-identity (in-group) condition than in the two aggregated weak-identity conditions (p = .02).

As with hypothesis 1, results from males are inconsistent with the predicted main effect. As shown in the rightmost column of Table 7, there was no main effect of identity on cooperation (p = .31). Thus, data from females provide support for hypothesis 2 and the utilitarian model, while data from males support the collectivist model, which predicts no main effect.

Hypothesis 3, based on the collectivist model, predicts an interaction between identity and the number of cooperative alters. We found no interaction effect among females (p = .85), confirmed by the parallel lines in Figure 4. The slopes do not change with the strength of identity.

In striking contrast, Table 7 (rightmost column) and Figure 5 show strong support for the predicted interaction among males (p = .02). Here the effects of identity depend decisively on the number of cooperative alters, as predicted by the collectivist model. Social identity increased cooperation when two alters cooperated, but *decreased* cooperation when one alter cooperated.

In sum, data from females fit two of the predictions of the utilitarian model, the positive effect of social identity, and the absence of an interaction between identity and alters' cooperation. However, the main effect of alters' cooperation, while significant, was in the wrong direction, with female cooperation rates *higher* when two alters cooperated. This positive effect among females is consistent with the collectivist model. Even more striking was the fit between data from males and the collectivist model, clearly evident in a comparison of Figures 3 and 5. For males, social identity increased cooperation when two alters cooperated, but decreased cooperation when one alter cooperated.

The most plausible conclusion from these data is that identity matters, and it matters because it motivates an effort not only to increase in-group payoffs but also to decrease in-group inequality. The latter possibility is strongly suggested by social identity theory and research on the maximal distinctiveness hypothesis, but it has been curiously overlooked in previous research on identity and cooperation in social dilemmas.

Discussion

We argue that structure affects power through two mechanisms, one related to the *motivation* for collective action, the other to the *success* of collective action. The motivation for collective action results from bidding wars, based on actors' locations in social structures. The success of collective action, in turn, rests on solving the free-rider problem. In short, social structures determine power not only by shaping opportunities to exchange, but also by shaping opportunities to act collectively to counteract structural disadvantages.

This argument is not new. Collins (1985), for example, argues that unequal exchange between capital and labor reflects not only a structural asymmetry in the ability to exclude workers from access to the means of production, but also the structural constraints on the ability of workers to act collectively to protect themselves from exclusion. Similarly, Marx (1844) points to both the asymmetry between labor and capital in their dependence on exchange with each other, and the greater difficulty entailed in the organization of workers: "The capitalist can live longer without the worker than the worker can live without him. Combination among the capitalists is habitual and effective, while combination among the workers is forbidden and has painful consequences for them" (19).

Our work focuses on structures deliberately chosen to make collective action among low-power actors possible but uncertain. Consistent with the argument that networks constrain free riding on coalitions, we found less cooperation in the M-branch exchange network than was observed in the Cook and Gillmore (1984) study. However, we found that low-power actors were more likely to join than free-ride when two others joined a coalition, even though the coalition needed only two members to obtain an equal exchange with a high-power actor.

The reluctance to free-ride does not, however, lead us to embrace the assumption that collective action to balance structural inequalities is unproblematic. On the contrary, we suspect that the explanation for the surprising frequency of three-person coalitions lies in another structural mechanism that researchers have largely overlooked. Social relations may not only structure *incentives* to bargain and organize, they may also structure social *identities*, based on perceptions of shared fate. Identification with occupants of structurally equivalent low-power positions leads actors to take into account the payoffs to others like themselves.

Although our theoretical interest centers on structural determinants of identity, it is difficult to tease apart the incentive and identity effects of structural disadvantage by manipulating network configurations. Any structure that assigns a shared fate to a set of positions also assigns an isomorphic common interest (Bonacich Bienenstock 1997). Therefore, we induced identification with (or against) other low-power positions using a cultural attribute — school affiliation — that is orthogonal to incentives and can be easily crossed with structural similarity.

We tested two specifications of the effects of social identity on payoff transformations. The utilitarian model assumes that utility functions take into account the payoffs to all in-group members. This is similar to the approach taken in most previous studies of the effect of identity on cooperation in social dilemmas. However, social identity theory suggests a somewhat more complicated mechanism: the maximization of group distinctiveness. Accordingly, we also tested a collectivist utility function that maximizes the mean payoff while giving equal weight to the similarity of payoffs among in-group members. We found support for both models, but not among the same participants. Females seem to respond to identity as a goal amplification effect (Kramer & Goldman 1995), while males respond as a goal transformation. Figure 4 shows that, for females, social identity simply increases baseline levels of cooperation. As shown in Figure 5, among males, the effect of identity on cooperation changes directions and depends on the number of cooperative alters.

Most important, when social identity was salient, both males and females were much more likely to cooperate when two alters cooperated than when one alter cooperated. This is clearly contrary to the dictates of pure self-interest. It also contradicts the utilitarian model based on the maximization of group interests. Instead, the results support the collectivist model in which in-group members also seek to maximize similarity within. Although this possibility has strong theoretical and empirical support in other avenues of social identity research, it has been overlooked in previous studies of identity and cooperation in social dilemmas. Our findings suggest that future work in this area needs to consider not only the effects of identity on the level of in-group payoffs but also on the variance.

Gender and Coalitions

Although gender differences are not the focus of this study, some related research led us to anticipate this possibility, clearly confirmed in Figures 4 and 5. It is important to note that the gender differences are evident primarily in the personal identity (self-interest) conditions, contrary to the finding by Kramer and Brewer (1984) that males and females responded differently to manipulation of social identity. Our results are consistent, however, with recent research on males' greater responsiveness to the greed component in social dilemmas (Simpson 2003). Our coalition design tempted participants with an opportunity to free-ride (greed) but did not pose a risk of being suckered (fear). Comparison of Figures 4 and 5 shows that in the weak identity conditions, when payoffs to self are relatively more important, males are more likely than females to take the bait. In the social identity conditions, when payoffs to alter become more salient, the temptation diminishes and so too does the difference between male and female behavior. Thus, the gender differences from this study are consistent with those reported by Simpson (2003) for a different type of collective action setting.[15]

Conclusion

This study breaks new ground by bringing collective action back into the study of structural disadvantage. As such, we invite readers to regard these results as tentative. The important lesson to draw from this research is the need for much more attention to a problem that has been badly neglected. We therefore conclude by identifying the directions for future work we see as most pressing.

Foremost, we need to find ways to directly manipulate the structural basis of identity. While the use of school affiliation to induce in-group bias is a useful starting point, it is clearly not an adequate test of a structural theory of identity. We assume that school affiliation reinforces identification with other low-power actors, much the way union affiliation reinforces class identification with other structurally disadvantaged workers. However, we do not have a direct test of the theory of structural sources of identity. Having found strong indirect support for the theory, we plan a follow-up study to test whether structural equivalence can produce similar effects when actors share no other basis of identity.

We also see the need for research that more fully explores the effects of social identity on coalition joining. Our study is limited to utility functions that take into account the payoffs to in-group members. This is a convenient starting point because it integrates the effects of incentives and identities within a single formal model. However, identity may have other important effects on participation in coalitions, such as increasing the expectation that others will cooperate (Yamagishi & Kiyonari 2000), or increasing feelings of guilt for free riding on cooperative members of the in-group (as suggested by one reviewer).

The effects of social identity reinforce Bonacich and Friedkin's (1998) call for the explicit incorporation of extraegoistic motives into models of exchange networks. These motives "have been viewed as confounding the study of structural effects, but they may be worth including as part of the specification of a general social exchange model" (170). More generally, as noted by Emerson (1987), sociological models of exchange differ decisively from economic models of a fully connected market by theorizing interactions embedded in a web of social relationships. Yet social exchange theory has been reluctant to abandon the economic model of the self-interested actor. Recent exceptions (Cook & Hegtvedt 1986; Lawler & Yoon 1998) demonstrate the need for a relational conception of the actor, in which behavior is influenced by social ties to significant others.

We also need to know more about the return path, in which behavior alters the structure of social ties. Actors often can and do form coalitions to successfully transform power structures and countervail existing power inequalities. Exchange network theories must move away from the traditional conception of social structure as fixed and unyielding. Our study is only a first step in this direction; much remains to be done. For example, we did not allow countermobilization by high-power actors, either via the formation of countercoalitions (see Borch & Willer 2001) or by divide-and-rule tactics (such as collaboration with free riders). These are important directions for future research, but we caution that Emerson's balancing processes need to be better understood before we complicate the picture with counterbalancing.

A related point is that future research needs to address questions from the social psychological literature on coalitions, such as what rules coalition members use to divide resources (Gamson 1961; Komorita & Chertkoff 1973). In our coalition design, profits from coalition members' exchanges were redistributed automatically and equally to all members. Given that low-power positions in our study were structurally equivalent, no actor had a *structural* basis for making claims to a larger share of coalition earnings. Thus, the equal redistribution rule is a reasonable simplification in the current context. The assumption is less tenable, however, for networks in which potential coalition members are not structurally equivalent and, as a result, may have different effects on coalition success. We expect rules governing intra=coalition resource distribution will be especially important in these networks. More generally, future research might fruitfully employ a network exchange approach to illuminate the *structural* determinants of the processes that have long concerned social psychologists interested in coalitions.

Our results also suggest new directions for the growing body of literature on socially embedded collective action (Marwell, Oliver & Prahl 1988; Snow, Zurcher & Ekland-Olson 1980). These studies show how networks help *solve* collective action problems. The possibility that networks might *create* these problems remains largely unexplored. Moreover, we need to explore how these effects vary with network structure. The motivation for staying outside the coalition depends decisively on the shape of the network structure (Bonacich 1995; Willer & Skvoretz 1997). For example, Willer and Skvoretz (1997) show that coalition formation in some networks poses an assurance game, characterized by fear but not greed. Recent research (Simpson 2003) found no gender differences in games like Assurance, while we found that males and females respond differently in a chicken-like game characterized by greed but not fear. Future research should further explore the possibility that the effects of gender and social identity depend on the incentive structure of network configurations.

In closing, we hope our research has shed new light on the dynamic and reciprocal relation between network structure and the ability to organize as well as bargain. We also hope these early findings will encourage others to bring collective action back into sociological research on power in social exchange.

APPENDIX A: Translating Payoffs into Cooperation Rates

We assume that, for any given rate at which players discount future payoffs, the larger the difference in utility across immediate choices, the greater the probability that players will choose the strategy with the higher utility. We formalize this using Harsanyi's (1973) purification theorem, which gives probabilities for cooperation based on smoothed best responses, given random payoff perturbations (Fudenberg & Levine 1999). In other words, we assume that a choice between two strategies with almost identical utilities will be much closer to 0.5 probability than to 1.0, as would be the case if we assumed strict preference.

More precisely, let $U_i(S)$ represent the utility of strategy S for player i, where S = [C,D]. (We assume for the moment that utility is the same as the payoff for a given strategy.) We then norm the utility of cooperation as

$$\upsilon_i = \frac{U_i(C)}{U_i(C) + U_i(D)} \quad \text{(A1)}$$

and then derive the probability p that i cooperates as

$$p_i = \frac{1}{1 + e^{(1-2\upsilon_i)M}} \quad \text{(A2)}$$

where M is a slope parameter for the cumulative logistic function and $M > 1$. If $M = 1$, the function approaches linearity such that an increase in the relative payoff yields a proportional increase in the likelihood of cooperation (but within a truncated range). As M increases, the function approaches the discrete step function for deterministic strict preference, in which values around the inflection point produce large shifts in the probability of cooperation. We chose an intermediate value, $M = 5$, that preserves the strict-preference tipping point but relaxes the deterministic prediction. At $M = 5$ the stochastic function is sigmoid and approximates a cumulative normal logistic function. Application of this function to the utilities given in Tables 1–3 results in the predicted cooperation rates given in the rightmost column of each table.

APPENDIX B: Simulated Actor Strategies

The strategies of simulants were based on previous studies using programmed actors (Cook & Gillmore 1984; Willer & Skvoretz 1997). Programmed sellers accepted their highest offers; and chose randomly among tied offers. These strategies were then evaluated and refined in pretests to appear highly realistic. A seller monitored changes in its best offer as bidding continued. Pretests suggested that allowing sellers to accept a few offers prematurely increased realism. Otherwise, sellers waited until the value of the best offer stopped increasing, indicating that bidding had stopped or slowed. Sellers followed this bargaining algorithm throughout the experiment with two exceptions. To prevent mechanistic behavior, we programmed one of the two sellers to deviate from the algorithm in rounds 4 and 8. In these two rounds, this seller refused all offers if the best offer was less than 20 points.

During game 1, simulated bidders used a Monte Carlo algorithm to choose an initial offer, with a normal probability distribution centered on the value of its exchange in the previous round, with a range of ±3. Specifically, imagine that B1 and S1 exchanged at 13:11 (favoring S1) at round t. Then at $t + 1$, B1 offered 13 (with $p = .25$); 12 or 14 (each with $p = .1875$), 11 or 15 ($p = .125$) or 10 or 16 ($p = .0625$). If B1 did not exchange the previous round, its initial offers were based on B2's exchange outcome. If B2 did not exchange, B1's initial offers were based on the subject's exchange. For the first round, a previous exchange ratio of 12:12 was assumed.

After initial offers, each simulated bidder responded to competing offers from other bidders using a similar Monte Carlo algorithm, but with the probability distribution now centered on an expected offer that was one point more than the competing offer, with a range of –1 to +3. If the subject exchanged before both simulated bidders, the simulants responded to each other's offers until a chance refusal to outbid triggered acceptance by one of the simulated sellers. Thus, the participant observed the bidding process continuinge if she exchanged first.

The Monte Carlo methods prevented the appearance of mechanistic behavior by the simulated bidders. To add to realism, simulants did not immediately respond to competing offers but instead adjusted their offers at randomly selected intervals. These delays prevented simulants from appearing overly responsive to the participant's actions and ensured that the participant and each simulant were included in about the same number of exchanges, which would be expected if all bidder positions were occupied by human actors.

In game 2, the simulated bidders voted whether to join the coalition. Simulant voting was predetermined, based on the treatment condition (see text) and the round, as shown in Table 5. These decisions were one of the experimental manipulations, and the rationale is elaborated below in the discussion of the treatment conditions. Note that simulants were programmed with a mixed strategy early in game 2 (with deviations in rounds 3 and 5), to give the impression that they were exploring their options before settling on a pure strategy.

If fewer than two bidders joined a coalition, the simulants used the same bidding strategies as in game 1. However, if two or more bidders formed a coalition, the simulated coalition members voted on the coalition bid using a Monte Carlo algorithm, with a normal probability distribution centered on the last group vote and with a range of ±3. Finally, if the participant and one simulated bidder formed a coalition and the other simulant did not, the simulated free rider used the same Monte Carlo bidding algorithm as in game 1, with an expectation of beating the coalition offer by 1 point.

TABLE A-1: Observed Payoffs (Not Transformed) and Predicted Probability of Cooperation As Smoothed Best Reply

Number of Cooperative Alters (B_1 and B_2)	B_3's Strategy C	B_3's Strategy D	Probability of Cooperation
0	**6.6** 6.6, 6.6	**6.9** 6.9, 6.6	.47
1	**8.2** 15.15, 8.2	**6.4** 6.4, 6.4	.65
2	**11.11** 11.11, 11.11	**15.15** 8.2, 8.2	.32

Notes: Values given in bold are payoffs to B_3, the focal actor. B_1, B_2, and B_3 are as defined in Figure 1.

APPENDIX C: Coalition Payoffs in Game 2

We used exchange ratios for game 2, broken down by coalition size, to determine whether the observed incentive structure corresponded to the incentives derived from the theory of countervailing power (Table 1). Manipulation checks confirm the similarity of the incentive structures.

The observed payoffs from game 2 are reported in Table A1. These payoffs show that free riding on a two-person coalition yields a better payoff than cooperating with two others. Similarly, two-person coalitions yield a better payoff to members than defecting when one other cooperates. As can be seen by comparing Table A1 with Table 1, for each coalition size, the observed payoffs are about 3 points greater than the corresponding prediction, but the ordinal relationship is preserved, leaving the observed marginal returns for cooperation very close to the predicted values. Indeed, the marginal return for cooperation when two alters join the coalition is actually even more strongly negative (–4.04) than predicted (–3).

Notes

1. For an overview of social exchange theories, see Molm and Cook (1995). For representative studies, see Cook and colleagues (1983), Bienenstock and Bonacich (1992), Friedkin (1992), Markovsky, Willer, and Patton (1988), and Whitmeyer (1999).

2. The nonmember is a free rider in that she enjoys the public benefits of coalition formation while avoiding the private costs associated with joining the coalition. A familiar example in natural settings is the worker who enjoys the benefits of a union while refusing to join and pay dues.

3. Following Lawler and Yoon (1998), our research is based on social identity theory, rather than role-identity or collective-identity theory. In very broad strokes, the three

approaches focus on three sources of identity: groups, roles, and shared beliefs/ commitments. Differences between these three approaches are discussed by Stryker (2000). Klandermans and de Weerd (2000) compare collective-and social-identity theories in application to collective action. We leave exploration of alternative theories (and sources) of identity in coalition formation to future research.

4. Utilitarianism and collectivism are not the only possible payoff transformations. Other possibilities are rational altruism ("maximize the payoffs to other in-group members without regard to your own") egalitarianism ("minimize the differences in payoffs to all in-group members, regardless of the mean"), and competitivism ("maximize differences in payoffs to in-group and out-group, without regard to individual members' outcomes"). Our focus on the utilitarian and collectivist models is based on previous research by Brewer, Kollock, Yamagishi, and others; we leave other possible specifications for future research.

5. The maximal distinctiveness hypothesis complements Brewer's theory of optimal distinctiveness, which "postulates that social identification is derived from the opposing forces of two universal human motives — the need for inclusion and assimilation on the one hand and the need for differentiation from others on the other . . . . The need for inclusion/belonging is met within the in-group; the need for differentiation, by distinctions between in-groups and out-groups" (Brewer & Silver 2000:155). Thus, while optimal distinctiveness determines the *level* or *basis* of social identification, maximal distinctiveness processes are consequences of these "optimally distinct" in-group/out-group distinctions.

6. See Van Lange (1999) for a similar formal specification of how prosocial actors strike a balance between joint outcomes and equality of outcomes.

7. For a formal derivation of this coalition game, see Simpson and Macy (2001). For related discussion of the social dilemmas entailed in coalition formation in exchange networks, see Bonacich (1995) and Willer and Skvoretz (1997). For a theory of "who forms coalitions with whom," see Bonacich (2000).

8. Why does the coalition settle for one exchange rather than offer fourteen and try to outbid the free rider for the second exchange? The free rider only has to outbid the lower of the coalition's two offers to avoid exclusion. Thus, the coalition is better off earning 12 points in one exchange than excluding the free rider by making two offers that are so high the free rider cannot improve on the lower of the two offers.

9. In a weak Nash equilibrium, unilateral deviation does not change the payoff. In the M-branch coalition game, zero cooperation is a weak Nash equilibrium. At zero cooperation, players are indifferent between defection and unilaterally switching to cooperation. However, should one player cooperate out of indifference, the others will no longer be indifferent and will now prefer cooperation.

10. We used the M-branch network, which contains five positions, but only four participants were scheduled for each session. However, the arrangement of partitions ensured that participants could not tell how many others were in the laboratory.

11. We decided against using the simpler network used in our previous work, which contained one high power actor who could exchange with any two of three low-power

actors. Although that network is behaviorally equivalent to the M branch used in this study, we didn't use the simpler structure because, unlike the M branch, using it would have required us to explain to participants that some positions could exchange only once, while others could exchange more than once. More substantively, we wanted to use a network that made salient in-group/out-group distinctions and felt that a group of two high-power actors would be more salient than a group of one.

12. Most network exchange experiments allow all positions to make offers and counter-offers. While it is relatively straightforward to create realistic simulated bidders that respond convincingly to a participant's offers, it is much more difficult to create simulated sellers that react to participants' offers with credible counter-offers and related negotiation activity. Thus, as in our previous work and related work by Willer and Skvoretz (1997), only bidders were allowed to make offers. Sellers could either accept standing offers or refuse them. It is important to note that Willer and Skvoretz (1997) show that such decision strategies on the part of simulated high-power actors produce power inequalities similar to those that emerge when high-power positions are occupied by human actors who *can* make offers.

13. Alternatively we could have provided participants with knowledge of others' cooperation by using sequential play with rotating order. However, this would allow two-person coalitions with rotating opportunities to free-ride, yielding equal payoffs to all in-group members. Equal payoffs with two-person coalitions would preclude a test of the collectivist model.

14. The post-treatment survey detected suspicion by asking participants to describe any aspect of the study they found hard to believe. This item may overestimate suspicion by motivating participants to search post hoc for deceptive aspects of the study. Reanalysis with suspicious participants included gave virtually identical results. This suggests that suspicion may have been induced by the survey item or suspicious participants became immersed in role-playing. For example, one participant who suspected the other bidders were fake nevertheless responded to a subsequent item that asked about motives for joining the coalition by stating "I joined to try and establish a bond with the other players, so that we could work together."

15. One reviewer asked whether the gender differences observed in this study may have resulted from the use of competitive terms such as *market* and *bidders,* or by our manipulation of identity using school affiliation (which may be associated with athletic rivalries and hence may be more important to males). However, a gender difference in the salience of school affiliation would lead us to expect gender differences in coalition behavior in the in-group condition — a difference we did not observe. It is more difficult to rule out the possibility that gender differences were induced by the framing of the situation as a competitive market exchange. However, gender differences have not been observed in a number of exchange studies using similar market-like scenarios (Kollock 1994; Lawler & Yoon 1998).

References

Bienenstock, Elisa J., and Phillip Bonacich 1992. "The Core As a Solution to Exclusionary Networks." *Social Networks* 14:231–44.

Blau, Peter M. 1964. *Exchange and Power in Social Life*. Wiley.

———. 1974. "Presidential Address: Parameters of Social Structure." *American Sociological Review* 39:615–35.

Bonacich, Phillip. 1995. "Four Kinds of Social Dilemmas within Exchange Networks." *Current Research Issues in Social Psychology* 1:1–7. [http://www.uiowa.edu/~grpproc]

———. 2000. "The Duality of Exchange and Coalitions: Equalizing Power in Exchange Networks." *Rationality and Society* 12:353–76.

Bonacich, Phillip, and Elisa J. Bienenstock. 1997. "Latent Classes in Exchange Networks: Sets of Positions with Common Interests." *Journal of Mathematical Sociology* 22:1–28.

Bonacich, Phillip, and Noah E. Friedkin. 1998. "Unequally Valued Exchange Relations." *Social Psychology Quarterly* 61:160–71.

Borch, Casey, and David Willer. 2001. "High Power versus Low Power Coalitions: Testing the Effects of Collective Action." Paper presented at the annual meeting of the American Sociological Association, Anaheim, California.

Brewer, Marilynn B., and Roderick M. Kramer. 1986. "Choice Behavior in Social Dilemmas: Effects of Social Identity, Group Size, and Decision Framing." *Journal of Personality and Social Psychology* 3:543–49.

Brewer, Marilynn B., and Michael D. Silver. 2000. "Group Distinctiveness, Social Identification, and Collective Mobilization." Pp. 153–71 in *Self, Identity and Social Movements,* edited by Sheldon Stryker, Timothy J. Owens and Robert W. White. University of Minnesota Press.

Brown, Rupert, Steve Hinkle, Pamela G. Ely, Lee Fox-Cardamone, Pamela Maras, and L.A. Taylor. 1992. "Recognizing Group Diversity: Individualist-Collectivism and Autonomous-Relational Social Orientation and Their Implications for Intergroup Processes." *British Journal of Social Psychology* 31:327–42.

Coleman, James S. 1990. *Foundations of Social Theory.* Harvard University Press.

Collins, Randall. 1985. *Four Sociological Traditions.* Oxford University Press.

Cook, Karen S., and Richard M. Emerson. 1978. "Power, Equity, and Commitment in Exchange Networks." *American Sociological Review* 43:721–39.

Cook, Karen S., Richard M. Emerson, Mary R. Gillmore, and Toshio Yamagishi. 1983. "The Distribution of Power in Exchange Networks: Theory and Experimental Results." *American Journal of Sociology* 89:275–305.

Cook, Karen S., and Mary R. Gillmore. 1984. "Power, Dependence, and Coalitions." *Advances in Group Processes* 1:27–58.

Dawes, Robyn M., Alphons J.C. van de Kragt, and John M. Orbell. 1988. "Not Me or Thee but We: The Importance of Group Identity in Eliciting Cooperation in Dilemma Situations: Experimental Manipulations." *Acta Psychologica* 68:83–97.

Deaux, Kay. 1996. "Social Identification." Pp. 777–98 in *Social Psychology: Handbook of Basic Principles*, edited by E. Tory Higgins and Arie W. Kruglanski. Guilford Press.

Deaux, Kay, and Anne Reid. 2000. "Contemplating Collectivism." Pp. 172–90 in *Self, Identity, and Social Movements,* edited by S. Stryker, T.J. Owens, and R.W. White. University of Minnesota Press.

de Weerd, Marga, and Bert Klandermans. 1999. "Group Identification and Social Protest: Farmers' Protest in the Netherlands." *European Journal of Social Psychology* 29:1073–95.

Emerson, Richard M. 1962. "Power-Dependence Relations." *American Sociological Review* 27:31–41.

———. 1987. "Toward a Theory of Value in Social Exchange." Pp. 11–45 in *Social Exchange Theory*, edited by Karen S. Cook. Sage Publications.

Friedkin, Noah E. 1992. "An Expected Value Model of Social Power: Predictions for Selected Exchange Networks." *Social Networks* 14:213–30.

Fudenberg, Drew, and David K. Levine. 1999. *The Theory of Learning in Games.* MIT Press.

Gamson, William A. 1961. "An Experimental Test of a Theory of Coalition Formation." *American Sociological Review* 26:565–73.

Harsanyi, John. 1973. "Games with Randomly Disturbed Payoffs." *International Journal of Game Theory* 2:1–23.

Hogg, Michael A. 1996. "Intragroup Processes, Group Structure, and Social Identity." Pp. 65–93 in *Social Groups and Identities*, edited by W. Peter Robinson. Butterworth-Heinemann.

Klandermans, Bert, and Marga de Weerd. 2000. "Group Identification and Political Protest." Pp. 68–92 in *Self, Identity, and Social Movements,* edited by Sheldon Stryker, Timothy J. Owens and Robert W. White. University of Minnesota Press.

Kollock, Peter. 1994. "The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust." *American Journal of Sociology* 100:313–45.

———. 1998. "Transforming Social Dilemmas: Group Identity and Cooperation." Pp. 186–210 in *Modeling Rational and Moral Agents*, edited by P. Danielson. Oxford University Press.

Komorita, Samuel, and Jerome M. Chertkoff. 1973. "A Bargaining Theory of Coalition Formation." *Psychological Review* 40:149–62.

Kramer, Roderick M., and Marilynn B. Brewer. 1984. "Effects of Group Identity on Resource Use in a Simulated Commons Dilemma." *Journal of Personality and Social Psychology* 46:1044–57.

———. 1986. "Social Group Identity and the Emergence of Cooperation in Resource Conservation Dilemmas." Pp. 205–34 in *Experimental Studies of Social Dilemmas*, edited by H. Wilke, C. Rutte, and D.M. Messick. Peter Lang.

Kramer, Roderick M., and Lisa Goldman. 1995. "Helping the Group or Helping Yourself? Social Motives and Group Identity in Resource Dilemmas." Pp. 49–67 in *Social Dilemmas: Perspectives on Individuals and Groups*, edited by D.A. Schroeder. Praeger.

Lawler, Edward J., and Jeongkoo Yoon. 1998. "Network Structure and Emotion in Exchange Relations." *American Sociological Review* 63:871–94.

Leung, Kwok, and M.H. Bond. 1984. "The Impact of Cultural Collectivism on Reward Allocation." *Journal of Personality and Social Psychology* 47:793–804.

Markovsky, Barry, David Willer, and Travis Patton. 1988. "Power Relations in Exchange Networks." *American Sociological Review* 53:220–36.

Marwell, Gerald, Pamela Oliver, and Ralph Prahl. 1988. "Social Networks and Collective Action: A Theory of the Critical Mass. Part 3." *American Journal of Sociology* 94:502–34.

Molm, Linda D. 1988. "Status Generalization and Power-Imbalanced Dyads: The Effects of Gender on Power Use." Pp. 86–109 in *Status Generalization: New Theory and Research*, edited by M. Webster Jr. and M. Foschi. Stanford University Press.

Molm, Linda D., and Karen S. Cook. 1995. "Social Exchange and Exchange Networks." Pp. 209–35 in *Sociological Perspectives on Social Psychology*, edited by Karen S. Cook, Gary A. Fine, and J.S. House. Stanford University Press.

Russell, Bertrand. 1938. *Power: A New Social Analysis.* W.W. Norton.

Sherif, Muzafer, O.-J. Harvey, B. Jack White, William R. Hood, and Carolyn W. Sherif. 1961. *Intergroup Conflict and Cooperation: The Robbers Cave Experiment.* University of Oklahoma Press.

Simpson, Brent. 2003. "Sex, Fear, and Greed: A Social Dilemma Analysis of Gender and Cooperation." *Social Forces* 82:35–52.

Simpson, Brent, and Michael W. Macy. 2001. "Collective Action and Power Inequality: Coalitions in Exchange Networks." *Social Psychology Quarterly* 64:88–100.

Skvoretz, John, and David Willer. 1993. "Exclusion and Power: A Test of Four Theories of Power in Exchange Networks." *American Sociological Review* 58:801–18.

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

Stryker, Sheldon. 2000. "Social Psychology and Social Movements: Cloudy Past and Bright Future." Pp. 1–20 in *Self, Identity, and Social Movements,* edited by Sheldon Stryker, Timothy J. Owens, and Robert W. White. University of Minnesota Press.

Tajfel, Henri. 1981. *Human Groups and Social Categories.* Cambridge University Press.

———. 1982. "Social Psychology of Intergroup Relations." *Annual Review of Psychology* 33:1–39.

Thibaut, John, and Harold H. Kelley. 1959. *The Social Psychology of Groups*. Wiley.

Triandis, Harry C., Kwok Leung, Marcelo J. Villareal, and Felicia L. Clack. 1985. "Allocentric versus Idiocentric Tendencies: Convergent and Discriminant Validation." *Journal of Research in Personality* 19:395–415.

Turner, John C. 1985. "Social Categorization and the Self-Concept: A Social Cognitive Theory of Group Behavior." *Advances in Group Processes* 2:77–121.

Turner, John C., Michael A. Hogg, Penelope J. Oakes, Stephen D. Reicher, and Margeret S. Wetherell. 1987. *Rediscovering the Social Group: A Self-Categorization Theory*. Basil Blackwell.

Turner, John C., and Rina S. Onorato. 1999. "Social Identity, Personality, and the Self-Concept: A Self-Categorization Perspective." Pp. 11-46 in *The Psychology of the Social Self*, edited by Tom R. Tyler, Roderick M. Kramer, and Oliver P. John. Lawrence Erlbaum Associates.

Van Lange, Paul A.M. 1999. "The Pursuit of Joint Outcomes and Equality in Outcomes: An Integrative Model of Social Value Orientation." *Journal of Personality and Social Psychology* 77:337–49.

Whitmeyer, Joseph M. 1999. "Convex Preferences and Power Inequality in Exchange Networks." *Rationality and Society* 11:419–42.

Willer, David. 1999. *Network Exchange Theory*. Praeger.

Willer, David, and John Skvoretz. 1997. "Games and Structures." *Rationality and Society* 9:5–35.

Yamagishi, Toshio, and Toko Kiyonari. 2000. "The Group As the Container of Generalized Reciprocity." *Social Psychology Quarterly* 63:116–32.