Title
The Incidence of Exchange Networks

Journal
Social Psychology Quarterly, 58(3)

ISSN
0190-2725

Author
Friedkin, Noah E

Publication Date
1995-09-01

DOI
10.2307/2787043

Peer reviewed

Social Psychology Quarterly
1995, Vol. 58, No. 3, 213–221

The Incidence of Exchange Networks*

NOAH E. FRIEDKIN
University of California

This paper advances a line of work on an expected-value model of social exchange, in which a power structure indicates opportunities for exchange and a sample space of exchange networks. When the probability distribution of the networks in this sample space is known, the expected-value model provides an excellent account of the distribution of benefits among actors in a variety of power structures. The model would be more elegant if it also predicted the probability distribution of exchange networks. Furthermore, in its current form, the model fails to account for the observed exchange payoffs in the Kite, a structure that has attracted considerable attention among exchange theorists. Here I advance the expected-value model by deriving the probability distribution of exchange networks from a simple process in which the probability of an exchange at time t *depends on the value of an exchange at time* t*-1. I show that this approach addresses the anomalies posed by the Kite structure.*

The interplay between social structure and process has been deemphasized in experimental research programs on social exchange networks, although such an interplay often is highlighted in discussions of social exchange phenomena (Molm and Cook 1993). Experimental work on exchange networks has focused mainly on accounts of the distribution of benefits that actors obtain from exchanges. Much less attention has been given to explaining how exchange processes affect the development of exchange networks. In this paper I show how the two accounts may be intertwined—that is, how the outcomes of exchange affect the likelihood of subsequent exchange and, in doing so, shape an emergent social structure.

In recent work on network exchange phenomena, the prevailing theoretical agenda has been to construct a measure of point-centrality which, when applied to a network of potential exchange transactions, correctly predicts the resources that actors acquire through negotiated agreements. Substantial controversy has developed in the pursuit of this agenda. Skvoretz and Willer's (1993) recent evaluation of alternative theoretical approaches showed that an expected-value approach (Friedkin 1992, 1993) provides the closest empirical fits when *observed* relative frequencies of exchange networks are employed in determining actors' dependencies and, in turn, their exchange ratios.

In the expected-value model, a power structure indicates opportunities for exchange relations and a sample space of transaction networks; the approach incorporates a micro-model of bargaining that predicts exchange outcomes in each of these possible networks. An expected-value model of social exchange outcomes has been latent in the literature on social exchange; its development does not represent a dramatic new direction for the field. The approach was presaged in Emerson's (1962:41) reference to French's (1956) formal theory of social power as a treatment of power toward which his own work might move. Emerson, like French, viewed power as a relation that defines opportunities for interpersonal events, such as exchange transactions (Emerson) or interpersonal influences (French). Although neither Emerson nor French carried forward the logic of this initial formulation, Emerson (1972:56) verged on the point of departure of an expected-value approach to social power when he described an exchange relation as giving rise to opportunities that "result in transactions with probability p_{yk}." The crucial element is the idea that a power structure implies a sample space of events or transaction networks; with this cornerstone in place, the outlines of the remaining theoretical development are relatively straightforward.

Although the findings of Skvöretz and Willer (1993:814) support an expected-value approach to social exchange, they also point out an anomaly in the findings on one

* This paper was presented at the Conference on Theory and Research on Group Processes, held in Los Angeles on August 4, 1994. I thank the editor and anonymous reviewers of *SPQ* for their helpful comments on an earlier draft of this paper.

structure—the Kite (shown in Figure 1). In this structure, there is an observed advantage for one position (B) when no such advantage was predicted. Furthermore, Skvoretz and Willer observe,

> In one exceptional case—the Kite network—*the structural potential of exclusion overrides the effect of differences in experienced exclusion.* Even though the central actor B is excluded from exchange more than twice as often as the peripheral A actors (41 percent versus 16 percent), B nevertheless earns moderately more points in exchange with actors in the A positions (14.05 versus 9.95 points) (1993:814).

With dependencies linked to rates of exclusion, as they are in the expected-value model, it appears that B should be disadvantaged rather than advantaged.

Another unsatisfactory feature of the expected-value approach is its reliance, for accurate predictions, on the observed relative frequencies of exchange networks (the fractions of trials of an experiment that eventuate in particular patterns of exchange transactions). A baseline assumption of equally likely exchange networks is seriously flawed (Friedkin 1993:186); no formal model has been developed, however, to account for the incidence of particular exchange networks.

In the present paper, I advance the expected-value model by showing how the relative frequencies of exchange networks may arise from a simple exchange process in which the probability of an exchange at time *t* depends on the value of an exchange at time *t*-1. Furthermore, I show that this approach explains the anomalies in the Kite structure.

When actors are limited to no more than one exchange per trial of an experiment, the social exchange process in a Kite structure *evolves* over successive trials *from* a state in which the central B position is advantaged in exchanges with the peripheral A positions *into* a state in which this advantage has disappeared. This evolution occurs because of changes in the probability of future transactions between actors who are taking into account the benefits received from past exchanges. Thus B's favorable exchange ratios are based on the experience of actors *before* those exchanges in which B's favorable exchange ratios are manifested. Once B's power is manifested in profitable exchanges with A, the probability distribution of exchange networks is altered by the actors in the A positions; they begin exchanging more frequently with one another and thereby begin to more frequently exclude the actor in the B position. As B's vulnerability to exclusion increases, B's favorable exchange ratios disappear.

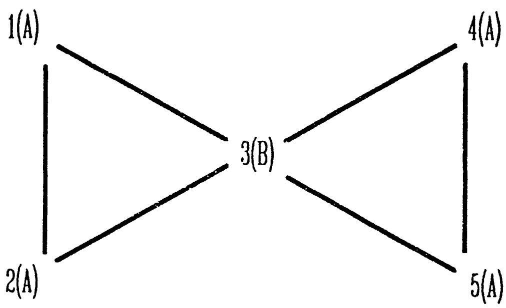

Figure 1. The Kite Structure

THE FORMAL MODEL

I describe the expected-value model in four steps. Apart from the determination of the relative frequency of exchange networks, these steps are the same as those I reported previously (Friedkin 1992, 1993). With an analysis of the Kite structure I illustrate the calculations that are entailed by the model. I focus on exchange networks composed of two-party transactions, in which each transaction provides one actor with a fraction of some amount of resources and the other actor with the remaining fraction. Such exchange networks have been the primary focus of experimental work on social exchange in recent decades.

Power Structures and Exchange Relations

The expected-value approach starts with the delineation of a power structure—that is, a network composed of 1) points indicating collective or individual actors and 2) undirected lines indicating potential transactions for each actor. The presence of a line between two actors indicates that they can exchange; the absence of a line between two actors indicates that they cannot. Thus, for example, the Kite structure contains six possible transactions or exchange relations:

$$\{1\text{–}2,\ 1\text{–}3,\ 2\text{–}3,\ 3\text{–}4,\ 3\text{–}5,\ 4\text{–}5\} = \text{Kite}.$$

In Figure 1 I have labeled the points with both numbers and letters; I use one or the other of these identifiers depending on which is more convenient in describing a procedure or finding. The numbers identify the nodes of the network; the letters signify the structurally distinct positions.

I make two simplifying assumptions about actors and potential transactions in a power structure. First, actors are assumed to be *rational* actors who seek to maximize their net receipt of resources over any set of transaction opportunities provided to them. (Later in the paper I replace this broad assumption of rational action with operational statements describing the type of action that is assumed to occur.) Second, the power structure is assumed to be *stable* with respect to the configuration of potential exchange transactions. Although the identities of the actors who occupy the various positions in the structure may change, the structure is stable as long as the configuration of possible exchanges remains the same.

Associated with a power structure is a *sample space* composed of the K different patterns of exchange transactions $\{R_1, R_2, \ldots, R_K\}$ that might occur. I refer to these patterns of transactions as R-nets. An R-net is a *subgraph* of the power structure that is feasible under a given set of empirical or theoretical conditions. Experiments on social exchange typically limit an actor to no more than a predetermined number of transactions with different actors on each trial of the experiment.[1] Each R-net in the sample space of a power structure is assumed to be *maximal* in that no other feasible transaction could occur. Rational actors do not absent themselves from exchange opportunities. Hence the relational networks should be maximal with respect to the number of feasible transactions.

The unrestricted sample space of the Kite structure contains 64 R-nets; however, under a one-exchange regime (see note 1) and the assumption of maximality, the number of feasible R-nets is reduced to five: $\{2\text{–}3, 4\text{–}5\} = R_1$, $\{1\text{–}3, 4\text{–}5\} = R_2$, $\{1\text{–}2, 4\text{–}5\} = R_3$, $\{1\text{–}2, 3\text{–}5\} = R_4$, and $\{1\text{–}2, 3\text{–}4\} = R_5$. One of these networks *must* occur on any given trial of an experiment in which rational actors are limited to no more than one exchange per trial.

Relative Frequency of Exchange Networks

In the elaborated model, R-nets are formed via sequences of random exchanges in which the probability of an exchange is determined by the weight of benefits provided by the exchange to the parties involved.[2] This weight is formulated as follows. Let $\mathbf{V} = [v_{ij}]$ be an nxn matrix in which v_{ij} is the amount of resources (benefit) that actor i receives from an exchange with actor j, and $v_{ij} = 0$ if actor i cannot exchange with actor j. The weight of each relation, $\mathbf{W} = [w_{ij}]$, is determined by the relative values of the resources it provides to the actors in the relation—that is,

$$w_{ij} = \left(\frac{v_{ij}}{\sum_{j=1}^{n} v_{ij}} \right) \left(\frac{v_{ji}}{\sum_{i=1}^{n} v_{ji}} \right). \qquad (1)$$

A particular exchange relation will be weighted heavily only insofar as the exchange values it provides are relatively high for *both* actors in the relation. For example, if actor i may exchange only with actor j, then the relative value of the exchange relation will be at its maximum for actor i; at the same time, the relative value of the exchange relation for j may be slight. The relative value of an exchange with actor i will be slight if many alternative sources of resources are available to actor j, or if the resources that j derives from an exchange with i are substantially less than the resources j derives from an exchange with some other actor. An exchange between actors i and j is likely only when the exchange is relatively attractive to both parties.

Now the probability of a particular sequence of exchanges can be formulated as the product

[1] I refer to restrictions on social exchange such as e-exchange regimes, where e = 1, 2,.... The one-exchange regime limits actors to one transaction at most per trial of an experiment; the two-exchange regime limits actors to two transactions at most with different actors; the three-exchange regime limits actors to three transactions at most with different actors; and so forth. Application of the expected-value model is not limited to e-exchange regimes. Below I focus on such regimes to facilitate a comparison of this model with other approaches.

[2] In my previous work, the probabilities of R-nets were determined either with data on the observed relative frequencies of R-nets or with a baseline assumption of equally likely R-nets. Although the baseline assumption was grossly misleading in certain cases (Friedkin 1993), an alternative formal model had not been developed; thus I relied on observed relative frequencies for predictions. The present approach increases the explanatory power of the theory with a model of value-driven exchange sequences that eliminates the reliance on observed relative frequencies of R-nets.

of the *relative* weights of the exchanges that compose the sequence. The agreement of two actors on an exchange restricts the exchanges that still can form; hence the weight of an exchange is relative to the weights of those exchanges which are possible at a given point in the sequence. For example, under the one-exchange regime in the Kite there are 10 possible sequences of exchanges. Each of these sequences results in a particular R-net. (The sequences are shown in Table 1.) Consider the sequence {1–2} {3–4}, which results in the relational network R_5. For the first exchange of the sequence there are six possible exchanges; the probability of the {1–2} exchange is given by its relative weight among these possibilities. For the second exchange, there are only three feasible exchanges; the probability of the {3–4} exchange is given by its relative weight among these possibilities.

Initially, where the benefits of an exchange are unknown, equally likely exchanges are assumed. Therefore, at the start of the process, the relative weight of an exchange is expressed simply by the number of alternative exchanges that are possible at a given point in the development of the sequence. These probabilities are displayed in Table 1 for each of the possible sequences in the Kite.

Given the probability of each exchange sequence, the probability of a particular R-net, $P(R_i)$, is simply the sum of the probabilities of those sequences which result in the network. For example, the probability of R_1 is the sum of the probabilities of the {2–3} {4–5} and the {4–5} {2–3} sequences. Initially the probabilities of the relational networks are .222 for {R_1, R_2, R_4, R_5} and .111 for R_3. These initial probabilities, $P_0(R_i)$, determine a set of exchange payoffs which, in turn, govern the process of the social exchange.[3]

A Micromodel of Bargaining

The values of actors' payoffs are determined by a micromodel of the bargaining process. This model builds on the thesis that the sizes of actors' offers are related inversely to the relative frequency of the actors' exclusion from social exchange. It involves a set of assumptions concerning initial offers and the reconciliation of inconsistent offers.[4]

Let actor i's initial offer to actor j be a function of the dependency of actor i on actor j

$$f_{ij} = a_{ij} - b_{ij}(c_{ij})^{d_{ij}} \qquad (2)$$

where $a_{ij} > 1$ is the amount of available resources, f_{ij} is the amount of these available resources that actor i initially offers to actor j, $0 \leq d_{ij} \leq 1$ is the *dependency* of actor i on actor j, and $0 < b_{ij} < a_{ij}$ and $0 < c_{ij} < 1$ are coefficients that determine the shape of the curve. Actor i's dependency on actor j is defined as the probability that actor i is excluded from an exchange *and* that the two actors do not exchange with each other.[5]

The curve (Eq. 2) rises from $a_{ij} - b_{ij}$. When actor i is least dependent on actor j ($d_{ij} = 0$), the predicted initial offer of actor i is

[3] This formulation is similar to that of Markovsky et al. (1993) in that probabilities of exchange are derived from a stochastic model of the process by which exchanges occur. The approaches differ in two respects. First, in Markovsky et al. the stochastic process enters into the prediction of power differentials only when no power differentials are predicted by the GPI_2 index. In the present approach, there exist no separate theoretical foundations for the emergence of power differentials. Second, in Markovsky et al. the stochastic process contains no theory of value; that is, on the assumption that actors randomly seek an exchange among their potential partners, the probability of a transaction is the joint probability that two actors seek an exchange with one another. In the present approach, exchange relations have values for the actors involved (values that may change over the trials of an experiment); these values determine the likelihood of exchange.

[4] The determination of values (i.e., the microbargaining model) is secondary to the more fundamental social process 1) that transforms a social structure (i.e., the probability distribution of alternative patterns of exchange transactions) and 2) that distributes and constrains the distribution of values in a population. In experiments that require an agreement to divide resources, values are determined by the actors' vulnerability to exclusion and by their bargaining. Where different resources are exchanged (for example, advice for status, or good x for good y), or where a resource is transmitted through intermediaries, values may be determined quite differently. If the receipt of value is coupled with dyadic exchange, however, then the probability of such exchange may be determined by the joint value of the exchange to the actors.

[5] Let $\max(\deg_i)$ represent the maximum number of exchanges for actor i in any of the R-nets. Actor i is excluded from an exchange in R_k if the number of the actor's exchanges in R_k is less than $\max(\deg_i)$; let A indicate this event. Let B indicate the event that actor i and actor j do not exchange with each other in R_k. The dependency of actor i on actor j is the probability of the joint event $d_{ij} = P(A \cap B)$ in the sample space of the power structure. Under one-exchange regimes, this formulation of an actor's dependency simplifies to the actor's vulnerability to exclusion (i.e., the probability that the actor is excluded from an exchange); hence under this regime an actor makes the same initial offer to all possible transaction partners. Under multiple exchange regimes, an actor's dependency and initial offers may vary for different transaction partners.

Table 1. Exchange Sequences in the Kite

Exchange Sequence	*Probability of Sequence*
1. {1–2} {3–4} = R_5	$P(\{1-2\}\{3-4\}) = \left(\frac{w_{12}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{34}}{w_{34} + w_{35} + w_{45}}\right) = .0556$
2. {1–2} {3–5} = R_4	$P(\{1-2\}\{3-5\}) = \left(\frac{w_{12}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{35}}{w_{34} + w_{35} + w_{45}}\right) = .0556$
3. {1–2} {4–5} = R_3	$P(\{1-2\}\{4-5\}) = \left(\frac{w_{12}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{45}}{w_{34} + w_{35} + w_{45}}\right) = .0556$
4. {1–3} {4–5} = R_2	$P(\{1-3\}\{4-5\}) = \left(\frac{w_{13}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{45}}{w_{45}}\right) = .1667$
5. {2–3} {4–5} = R_1	$P(\{2-3\}\{4-5\}) = \left(\frac{w_{23}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{45}}{w_{45}}\right) = .1667$
6. {3–4} {1–2} = R_5	$P(\{3-4\}\{1-2\}) = \left(\frac{w_{34}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{12}}{w_{12}}\right) = .1667$
7. {3–5} {1–2} = R_4	$P(\{3-5\}\{1-2\}) = \left(\frac{w_{35}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{12}}{w_{12}}\right) = .1667$
8. {4–5} {1–2} = R_3	$P(\{4-5\}\{1-2\}) = \left(\frac{w_{45}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{12}}{w_{12} + w_{13} + w_{23}}\right) = .0556$
9. {4–5} {1–3} = R_2	$P(\{4-5\}\{1-3\}) = \left(\frac{w_{45}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{13}}{w_{12} + w_{13} + w_{23}}\right) = .0556$
10. {4–5} {2–3} = R_1	$P(\{4-5\}\{2-3\}) = \left(\frac{w_{45}}{w_{12} + w_{13} + w_{23} + w_{34} + w_{35} + w_{45}}\right)\left(\frac{w_{23}}{w_{12} + w_{13} + w_{23}}\right) = .0556$

an amount less than the available resources. As actor i's dependency on actor j increases (i.e., as $d_{ij} \rightarrow 1$), the predicted offer steadily approaches the asymptote a_{ij}, the total available resources.

A priori values for the coefficients $\{b_{ij}, c_{ij}\}$ may be derived under two assumptions. First, actors who are minimally dependent offer only one unit of their available resources. Second, actors who are maximally dependent offer all but one unit of their available resources. From the first assumption, $a_{ij} - b_{ij} = 1$. From the second assumption, $f_{ij} = a_{ij} - 1$ when $d_{ij} = 1$. Hence $f_{ij} = a_{ij} - (a_{ij} - 1)^{1-d_{ij}}$. For example, when 24 units of a resource are at stake, actor i's initial offer to actor j is $f_{ij} = 24 - 23^{1-d_{ij}}$.

Two actors' offers are inconsistent if their personal claims on available resources ($a_{ij} - f_{ij}$ and $a_{ij} - f_{ji}$ respectively) do not sum to the amount of the available resources. The following three assumptions deal with such cases. 1) If the sum of the actors' personal claims exceeds the amount of available resources, they split the difference and settle on the average of their two offers; that is, i gets one-half the sum of his or her claim and j's offer. 2) If each actor wants less than one-half of the available resources, they divide the available resources evenly. 3) If one actor wants one-half the resources or more, if the other wants less than one-half the resources, and if the sum of their personal claims is less than the amount of available resources, then the unclaimed portion of the available resources is allocated to the actor with the lower of the two personal claims.

For the Kite structure under the one-exchange regime, the initial dependencies of the actors, $\mathbf{D}_0 = [d_{ij(0)}]$,

$$\mathbf{D}_0 = \begin{bmatrix} 0 & 0.222 & 0.222 & 0 & 0 \\ 0.222 & 0 & 0.222 & 0 & 0 \\ 0.111 & 0.111 & 0 & 0.111 & 0.111 \\ 0 & 0 & 0.222 & 0 & 0.222 \\ 0 & 0 & 0.222 & 0.222 & 0 \end{bmatrix},$$

are derived from the baseline probability distribution, $P_0(R_i)$. The initial values (payoffs) of the exchanges, $\mathbf{V}_0 = [v_{ij(0)}]$,

$$\mathbf{V}_0 = \begin{bmatrix} 0 & 12 & 9.61 & 0 & 0 \\ 12 & 0 & 9.61 & 0 & 0 \\ 14.4 & 14.4 & 0 & 14.4 & 14.4 \\ 0 & 0 & 9.61 & 0 & 12 \\ 0 & 0 & 9.61 & 12 & 0 \end{bmatrix},$$

are derived from the bargaining assumptions. These initial values, $\mathbf{V}_0$, represent the exchange ratios for transactions based on dependencies arising from a purely random sequence of exchanges (i.e., Table 1).

The $\mathbf{V}_0$ values, in turn, affect the probability of subsequent exchanges according to the formulation described previously, and result in a $P_1(R_i)$ probability distribution; in this new distribution the probabilities of the relational networks are .166 for $\{R_1, R_2, R_4, R_5\}$ and .338 for R_3. This $P_1(R_i)$ probability distribution determines new dependencies and payoffs. For time $t = 1$, the set of dependencies, $\mathbf{D}_1 = [d_{ij(1)}]$, is

$$\mathbf{D}_1 = \begin{bmatrix} 0 & 0.166 & 0.166 & 0 & 0 \\ 0.166 & 0 & 0.166 & 0 & 0 \\ 0.338 & 0.338 & 0 & 0.338 & 0.338 \\ 0 & 0 & 0.166 & 0 & 0.166 \\ 0 & 0 & 0.166 & 0.166 & 0 \end{bmatrix},$$

and the set of payoffs, $\mathbf{V}_1 = [v_{ij(1)}]$, is

$$\mathbf{V}_1 = \begin{bmatrix} 0 & 12 & 13.68 & 0 & 0 \\ 12 & 0 & 13.68 & 0 & 0 \\ 10.32 & 10.32 & 0 & 10.32 & 10.32 \\ 0 & 0 & 13.68 & 0 & 12 \\ 0 & 0 & 13.68 & 12 & 0 \end{bmatrix}.$$

The $\mathbf{V}_1$ values, in turn, affect the probability of subsequent exchanges according to the formulation described previously and result in a $P_2(R_i)$ probability distribution. In this new distribution the probabilities of the relational networks are now .199 for $\{R_1, R_2, R_4, R_5\}$ and .203 for R_3. The iterative process is continued until an equilibrium, if such exists, is attained. The equilibrium is a state in which the effects of the social exchange process on the probability distribution of the exchange networks have become negligible. For the Kite structure, the equilibrium probabilities of the relational networks are .192 for $\{R_1, R_2, R_4, R_5\}$ and .232 for R_3; the equilibrium value matrix, $\mathbf{V}_\infty = [v_{ij(\infty)}]$, is

$$\mathbf{V}_\infty = \begin{bmatrix} 0 & 12 & 12.6 & 0 & 0 \\ 12 & 0 & 12.6 & 0 & 0 \\ 11.4 & 11.4 & 0 & 11.4 & 11.4 \\ 0 & 0 & 12.6 & 0 & 12 \\ 0 & 0 & 12.6 & 12 & 0 \end{bmatrix}.$$

Note that at each point in time in this iterative process, $t = 1, 2, \ldots$, the matrix $\mathbf{V}_{t-1}$ describes the payoffs to actors at time t and determines the probabilities of exchange at that time; that is, $P_t(R_i) = f(\mathbf{V}_{t-1})$.

Expected Values

The expectation for an actor's net payoff is

$$E(\text{PAYOFF}_{i(t)}) = \sum_{k=1}^{K} P_t(R_k)\ \text{PAYOFF}_{ik(t-1)}$$

where $P(R_k)$ is the probability of R_k in a sample space composed of K networks and PAYOFF_{ik} is the resource outcome for actor i in R_k. That is, for each R-net a product is formed of the probability of the R-net and the outcome. These products are summed to form the expectation of the outcome.

Table 2 describes the predictions for the Kite structure. I present predictions for both the initial and the equilibrium states of the exchange process.[6]

ANALYSIS OF THE KITE STRUCTURE

The expected-value model predicts an initial advantage for B in exchanges with A (B/A = 14/10) that disappears at equilibrium (B/A = 11/13). Three sets of experimental results on the B/A exchange ratio support the conclusion that B has a slight advantage in exchanges with other actors: 14/10 (Skvoretz and Willer 1993), 13/11 (Markovsky et al. 1993), and 13/11 (Bienenstock 1992). These aggregate statistics, however, do not indicate the change (if any) that occurs in the exchange ratios over trials of the experiment.

Bienenstock's (1992) study permits a limited analysis of the exchange dynamics in the Kite.[7] Her study was based on 10 groups of actors, who were allowed to complete as many trials as they could in an allotted time.[8] These data are useful for the present purposes, with two warnings. First, when actors know the overall structure of exchange and their positions in the structure, initial exchange transactions may be affected by the actors' more or less accurate assessments concerning the implications of the structure. Hence the "blind" exploratory transactions at time $t = 0$ that are assumed by the expected-value model may be short-circuited or distorted by such knowledge. Second, in this experiment, although all of the groups completed four trials and all but two of the groups completed eight, most of the groups did not complete more than 10 trials; only four groups completed as many as 12, and only one group completed as many as 15. Therefore, beyond eight trials, these data became increasingly scanty and concern the ongoing adjustments of a rapidly diminishing number of groups.

Over successive trials, a rolling average of the payoff for B in B/A exchanges is 13, 14, 14, 14, 13, 13, 13, 13, 13, 12, 12, 12, 12, 12, 12.[9] In Trials 1–4 the average payoff for B is 13.6, based on 21 exchanges; in Trials 5–8 the average payoff for B is 13.3, based on 26 exchanges; in Trials 9–17 the average payoff for B is 12.5, based on 15 exchanges. This pattern indicates that the initial advantage for B is diminished, as predicted. The model also predicts that at equilibrium, B will be at a slight disadvantage in exchanges with A. The present data neither support this prediction nor strongly disconfirm it, on the basis of the scanty available data on B/A

[6] The equilibrium state indicates the long-run outcomes of the system. In some structures the initial state is an equilibrium state. Where the two states are not the same, as in the Kite, the expected values for the initial state should be interpreted as indicators of the status of the system at an initial point in its development, rather than as predictions of long-run outcomes.

[7] I am indebted to Elisa Bienenstock, who provided me with the data for this analysis.

[8] In Bienenstock's experiment, subjects sat in separate rooms and interacted through computer terminals. The task required connected pairs of subjects to bargain over the division of a set number of points. Subjects were assigned to a fixed position; that is, they were not rotated through different positions. They were aware of the structure of the network and of their position in the structure. Subjects were limited to no more than one exchange per trial. They were allowed to complete as many trials as they could in one hour. Each trial consisted of one to five rounds of bargaining (offers and counteroffers). Each round consisted of three stages: the offer, the acceptance, and the confirmation. In the first stage, subjects could make offers to all the actors with whom they might exchange. In the next stage, subjects were informed of the offers that had been made to them, and were reminded of the offers that they had made to others. Subjects either allowed the offers they had made to stand, or accepted an offer. If they accepted an offer, their standing offers to others became void. In the final stage of a round, subjects either accepted an offer made to them or did not accept any offers. If an agreement was reached, the parties to the agreement did not participate in subsequent rounds. Multiple rounds, up to five, were possible; at each round, subjects were informed of their possible exchange partners. All but six of the trials produced maximally complete patterns of exchange transactions. My analysis deals with the 97 maximal R-nets that emerged from this experiment.

[9] Let t indicate an experimental trial. For $t = 1,2,\ldots,15$ the rolling averages are computed by averaging the data for t, $t+1$, and $t+2$. The numbers of cases on which these averages are based are respectively 24, 27, 26, 28, 26, 25, 22, 19, 17, 15, 13, 9, 6, 4, and 3.

Table 2. Initial and Equilibrium Outcomes for the Kite Structure

Initial State					
	Actors				
$P_1(R_k)$	1(A)	2(A)	3(B)	4(A)	5(A)
R_1 .166	0	9.61	14.4	12	12
R_2 .166	9.61	0	14.4	12	12
R_3 .338	12	12	0	12	12
R_4 .166	12	12	14.4	0	9.61
R_5 .166	12	12	14.4	9.61	0
Expected Value	9.62	9.62	9.53	9.62	9.62
Equilibrium State					
	Actors				
$P_\infty(R_k)$	1(A)	2(A)	3(B)	4(A)	5(A)
R_1 .192	0	12.6	11.4	12	12
R_2 .192	12.6	0	11.4	12	12
R_3 .232	12	12	0	12	12
R_4 .192	12	12	11.4	0	12.6
R_5 .192	12	12	11.4	12.6	0
Expected Value	9.81	9.81	8.76	9.81	9.81

R-nets: {2–3, 4–5} = R_1; {1–3, 4–5} = R_2; {1–2, 4–5} = R_3; {1–2, 3–5} = R_4; {1–2, 3–4} = R_5; $P_0(R_3) = 0.111$ and $P_0(R_1) = P_0(R_2) = P_0(R_4) = P_0(R_5) = 0.222$.

exchanges after eight experimental trials. The A/A exchanges are uniformly 12/12, also as predicted.

The model helps us understand the combination, in the *initial* state of the process, of a B/A exchange ratio that favors B, a higher rate of exclusion for B than for A, and nearly identical net expectations for B and for A. Before any effects of exchange payoffs on the probabilities of exchange—that is, in the start-up ($t=0$) state, where the possible exchange transactions are equally likely—an actor who occupies the A position is twice as likely to be excluded from exchange as an actor who occupies the B position (.211 versus .111); hence the prediction of an initial B/A exchange ratio that favors B. As a consequence of these unbalanced exchanges, A–A exchanges are preferable to A–B exchanges for the actors in the A positions. Thus we get a predicted initial ($t=1$) state in which the B/A exchange ratio of 14/10 reflects B's past advantage, but in which B is excluded at a higher rate (.388) than in the past. The high price of doing business with B lowers the demand for transactions with B; although B's exchange ratios are favorable, A's preference for exchange with another A lowers the net returns for B.

The predicted initial exclusion rate of .388 for B is close to the rate of .412 that Skvoretz and Willer (1993) report and to the rate of .382 that occurs during the first four trials of Bienenstock's experiment. (The overall rate of exclusion for B is .357 in Bienenstock's data.) In Skvoretz and Willer's experiment, actors were rotated through all the positions of the Kite and remained in each position for four trials; thus the overall rate of exclusion for a position should reflect the initial stages, rather than the equilibrium, of an exchange process that evolves over experimental trials. In Bienenstock's data I find no evidence of "blind" exploratory trials in which the possible exchanges among the actors are equally likely. In her study, as in the study by Skvoretz and Willer, actors were fully informed about the structure and their positions; hence they may have grasped some of the implications of the structure at the very start of the process.

The expected-value model predicts a modest fluctuation of B's rate of exclusion and exchange ratios as actors iteratively adjust their preferences for exchange partners in response to their previous payoffs. For t = 1,2,3,4,5,6,..., ∞ the predicted exclusion rates for B are .338, .203, .248, .228, .233, .232,..., .232 respectively. In the first eight trials of Bienenstock's data, the rolling average of B's exclusion is consistent with the predicted development: .458, .333, .192, .214, .231, .240 (i.e., these rates are calculated as in note 9).

Beyond eight trials, the rolling average of B's exclusion rises and falls in an unexpected fashion: .273, .316, .529, .533, .692, .556, .667, .250, .000. Overall, beyond eight trials B's rate of exclusion is .483. This relatively high rate of exclusion for B should have the long-run consequence of further disadvantaging B in exchanges with A. As I stated previously, 15 transactions occurred between A and B during Trials 9–17, for which B's average payoff was 12.5; hence any further disadvantaging of B should bring A and B into complete parity or even reverse the exchange ratio in A's favor. Although the expected-value model predicts an exchange ratio that modestly favors A at equilibrium, the observed substantial *rise* in B's rate of exclusion after the ninth trial is surprising. Unfortunately an assessment of this finding is made difficult by the small and diminishing number of groups that completed more than eight trials of the experiment.

Over the trials of this experiment on the Kite, the changes in B's circumstances are manifested in the exchange ratios and the exclusion rates, and less so in the average per-trial payoffs to the actors. Although B has a predicted initial advantage in exchanges with A, B has approximately the same predicted average per-trial payoff as A—that is, 10. At equilibrium, B's predicted average per-trial payoff is only slightly lower than A's (9 versus 10). In Bienenstock's study, the overall average per-trial payoffs are 9 for the B position and 10 for the A positions. The rolling average of payoffs for B (see note 9) suggests no systematic trend: 7, 9, 12, 11, 10, 10, 10, 9, 6, 6, 4, 5, 4, 9, 12.

The expected-value model predicts that in the Kite at equilibrium, the differences in the actors' exchange ratios and net expectations are modest. More important, the model predicts that this relatively uninteresting equilibrium state is attained via microdynamics that entail substantive changes in actors' interpersonal dependencies. These dynamics reconcile the anomaly reported by Skvoretz and Willer (1993). The combination of B's advantaged position in B-A exchanges and a relatively high rate of exclusion appears to be a transitional state in the social exchange process. The evidence from Bienenstock's study suggests that over successive trials, the system of exchange in the Kite settles into an equilibrium in which power differences are slight.

DISCUSSION

In definitions of power in exchange networks, it has been conventional to describe power in terms of unbalanced exchange ratios, so that B has more power than A if B derives more benefit than A when the two actors exchange. Such a viewpoint on power has merits, as we have seen in the analysis of the Kite, but it also has potential pitfalls. Favorable exchange ratios are not necessarily consistent with the maximization of rewards through a social exchange process. Unbalanced exchanges may lower an actor's *net receipts* when the disadvantaged trading partners can turn elsewhere for an exchange. We saw something like this arise at an initial point in the exchange process in the Kite, where the structurally central actor B had favorable exchange rates with A but slightly lower average per-trial payoffs than A because of a relatively high rate of exclusion from exchange. By the same token, it is not difficult to imagine a situation in which an actor who is allowed multiple transactions on each experimental trial undercuts competitors by accepting unfavorable exchange ratios while amassing more resources over trials than any other actor through a large volume of transactions. Such circumstances suggest that no single measure of interpersonal power may be relied upon, and that the mechanisms of social exchange may be as important as the outcomes.

The present study raises an important issue concerning the stage at which the social exchange process is studied. If the process entails shifts in the pattern of interpersonal power, then summary statistics of outcomes that occurred throughout the process can be misleading. The Kite structure illustrates the problem, where it appears that exchange ratios are modified over the course of successive trials of an experiment. A sufficient number of trials must be conducted to allow such dynamics to emerge. We need studies that involve a large number of experimental trials on structures for which we expect different outcomes (e.g., exchange ratios) at different stages in the exchange process. With studies that reveal multiple stages of an exchange process we may assess whether a theory accurately represents not only the initial and equilibrium states of the process, but also the typical pattern of changes in groups over experimental trials.

Shifts in the pattern of exchange outcomes may occur when actors can improve on their expected rewards by altering the relative frequency of transactions with certain actors. A change in such a frequency in one exchange relation may have ramifications for transactions in other exchange relations; the present model seeks to describe the behavior of this system of adjustments. To examine such dynamics we require structures that are not at equilibrium in their initial states, and actors who do not enter into the process with an accurate assessment of the equilibrium implications of their structural positions. This latter condition may be satisfied with an experimental design in which subjects have no information about the structure of exchange opportunities in which they are situated, or about the positions that they occupy in the structure. The former condition is more difficult to satisfy because it involves the a priori identification of a power structure that will manifest important dynamical changes over the trials of an experiment. I conclude with a discussion of this problem.

When the expected-value model is applied to a particular power structure, it offers a description of the dynamics of the exchange process. Without applying the model to a particular power structure, there is no way of predicting whether the structure has important dynamical properties. Structures with important dynamical properties may be discovered by examining a large number of randomly constructed networks, but at present we do not know enough to design power structures that will surely manifest particular exchange dynamics and patterns of payoffs. There are different approaches to the development of such a theory. One approach is via input-output models that examine the relationships between structural variables and exchange outcomes; such an approach, however, does not elucidate the mechanisms by which social structure determines outcomes. In a more powerful approach, but one that is not always feasible, a process model is developed and the desired theory of structure is derived formally from that model. Such an approach avoids ad hoc structural hypotheses and constructs a coherent theory on simple foundations. I hope that the present expected-value model will encourage this sort of theoretical development.

REFERENCES

Bienenstock, Elisa J. 1992. "Game Theory Models for Exchange Networks: An Experimental Study." Doctoral dissertation, Department of Sociology, University of California, Los Angeles.

Emerson, Richard M. 1962. "Power-Dependence Relations." *American Sociological Review* 27:31–40.

———. 1972. "Exchange Theory, Part I: A Psychological Basis for Social Exchange." Pp. 38–57 in *Sociological Theories in Progress*, Vol. 2, edited by Joseph Berger, Morris Zelditch Jr., and Bo Anderson. New York: Houghton Mifflin.

French, J.R.P., Jr. 1956. "A Formal Theory of Social Power." *Psychological Review* 63:181–94.

Friedkin, Noah E. 1992. "An Expected Value Model of Social Power: Predictions for Selected Exchange Networks." *Social Networks* 14:213–29.

———. 1993. "An Expected Value Model of Social Exchange Outcomes." *Advances in Group Processes* 10:163–93.

Markovsky, Barry, John Skvoretz, David Willer, Michael J. Lovaglia, and Jeffrey Erger. 1993. "The Seeds of Weak Power: An Extension of Network Exchange Theory." *American Sociological Review* 58:197–209.

Molm, Linda D. and Karen S. Cook. 1993. "Social Exchange and Exchange Networks." Pp. 209–35 in *Sociological Perspectives on Social Psychology*, edited by Karen S. Cook, Gary A. Fine, and James S. House. New York: Allyn and Bacon.

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.

Noah E. Friedkin *is Professor of Education and Sociology at the University of California, Santa Barbara. His current research is concerned with the development of network theories of social exchange and opinion formation.*