CURRENT RESEARCH IN SOCIAL PSYCHOLOGY

Volume 3, Number 2
Submitted: February 9, 1998
Resubmitted: April 9, 1998
Accepted: April 15, 1998
Publication Date: April 28, 1998

THE EVOLUTIONARY STABILITY OF STRATEGIES IN EXCHANGE NETWORKS

Phillip Bonacich
University of California, Los Angeles

ABSTRACT

Different exchange networks may present different strategic problems to their members. Powerless positions faced with common exploiters may profit from exchange only if they can present a united front and avoid ruinous competition among themselves. In some networks the primary danger is not powerlessness but the risk of exclusion. Effective strategies may sacrifice short-term payoffs to increase the chances of inclusion. Evolutionary game theory offers some tools for determining effective strategies. Biologists and others interested in situations in which successful strategies propagate themselves with greater frequency have developed the concept of evolutionary stability of strategies. This concept can be modified to find stable configurations of strategies in exchange networks.

[12]
----
[13]

INTRODUCTION

I wish to propose a method and some preliminary results on the joint effects of strategy and structure in exchange networks. Previous studies in this area have focused on variations in structure, but in this paper the focus is on variations in strategy. By using computer simulations and mathematical models, I suggest some ways in which optimum strategies are different for different positions in networks.

This is a new area of exploration with two scientific payoffs to my approach. First, the search for optimal strategies may lead to the discovery of strategies more likely to be used by sophisticated, experienced bargainers in "real-life" situations outside the laboratory. Experimental results thus far have been based on undergraduate students unfamiliar with the situations they face. The second payoff is that new models for identifying positions of power may be developed. The power enjoyed by positions in exchange networks must be a function of the strategies used by actors. If sophisticated actors use more effective strategies, the power relations in the network may change. I will provide some examples later in the paper. The discovery of new strategies does not mean that existing approaches are incorrect, merely that there are conditions where their assumptions are not met.

I will be using concepts from non-cooperative and evolutionary game theory. Relations between *exchanging* actors in an exchange network can be described by cooperative game theory because they can come to binding agreements over the terms of their exchange. However, actors who do not bargain directly with one another (or who don't even communicate with one another) can nonetheless have a profound impact on each other's rewards, an impact that has been ignored in previous research on exchange networks. For example, a set of powerless actors faced with a common exploiter will have an interest in reducing competition among themselves. In experiments, these actors cannot communicate with one another. In "real life" they may not be able to form binding agreements (because of anti-trust regulations, for example). Non-cooperative game theory can provide the tools for analyzing these situations.

The topic of power within exchange networks has been an active one for the past fifteen years (for example, Bonacich and Bienenstock 1992; Burke 1997; Cook et al 1986; Cook and Yamagichi 1992; Markovsky, Willer, and Patton 1988; Markovsky et al 1993; Yamaguchi 1996). Researchers active in this area have attempted to identify the structural properties that identify powerful positions. One particular type of situation has generated a great deal of interest. Selected pairs of individuals, defined by a network, have opportunities and common interests in engaging in mutually profitable transactions. Actors can engage in a profitable transaction with only a limited number of other actors, so that individuals must make choices between alternative partners. For example, in an automobile market, each seller may have a limited number of automobiles and each buyer may want just one automobile. In a "dating" market, each individual may not be able to date more than one person in a given evening.

The typical experimental paradigm is even more specific. Transactions involve coming to an agreement about the distribution of a quantity of points (usually 24). Points are converted to money at the end of the experiment. Individuals can transact with at most one other individual per game. There is a fixed number of opportunities for transactions (games). Each transaction has the same value. In a game-theoretic sense, these are cooperative characteristic value games for dyads that bargain with one another.

Assumptions About Strategies

Structural assessments of power clearly are functions of the strategies that actors use. In computer simulations of exchange networks, the common assumption is that actors who have been excluded from an exchange in one game increase their possibilities of being included in future exchanges by increasing the value of their offers to others. Similarly, actors are assumed to try to test the strength of their partners by reducing their offers to others in the subsequent game if they are included in an exchange. Computer simulations based on these assumptions almost always agree with the experimental results, and, moreover, there is experimental evidence that subjects in experiments use these principles (Bonacich 1997; Thye, Lovaglia and Markovsky

1997). Bonacich (1996) has shown that the observed distribution of power in networks can be deduced mathematically from the first of these assumptions in conjunction with a few others.

This primitive strategy is, of course, applicable to many situations outside the laboratory. However, there may be situations in which it is not applicable. For example, "real" actors may know more about the network and about the rewards of other actors. They may have the experience to develop more sophisticated and knowledgeable strategies. Consider, for example, a set of automobile companies attempting to sell to the same consumers. Competition between companies is avoided through price leadership and other mechanisms, even though explicit and enforceable agreements are illegal.

Admitting other strategies suggests the possibility of new structural predictions of power. Markovsky (1987) makes this very point. The structural models that have been tested may be limited to situations in which only naïve strategies are used. Moreover, different strategies may develop in different kinds of positions within networks (Bonacich and Bienenstock 1997). The current assumption that all positions in networks use the same strategy may be unrealistic. Whitmeyer (1997) shows that coordinated strategies can produce greater gains for sets of weak positions and that this implies new structural measures of power. Strong positions may have more sophisticated strategies for maintaining their power.

In this paper I wish to suggest an approach to discovering effective and stable strategies in exchange networks. The method will be based on computer simulations and mathematical models.

METHOD

The method will make use of the fact that network exchange situations are iterated non-cooperative games for pairs of players who do not bargain with one another over the division of points but who may wish to coordinate their strategies. The fact that network exchange experiments always involve multiple rounds of exchanges means that pairs of positions that do not exchange with one another and cannot coordinate their strategies through binding agreements can, nonetheless, accommodate their strategies to one another. For example, weak players faced with a common exploiter cannot form a binding agreement to make only low offers to their common exploiter, but they can encourage low offers by following each others' leads in making low offers or punishing their partners for making high offers.

Consider, for example, the networks in Figures 1 and 2. In Figures 1 and 2, experimental evidence suggests that positions 4 and 5 tend to have power over positions 1, 2, and 3. Since these graphs are bipartite (all relations are *between* two disjoint sets of points), they could represent relations between buyers and sellers. Could the power distribution change with more effective strategies by players 1, 2, and 3?

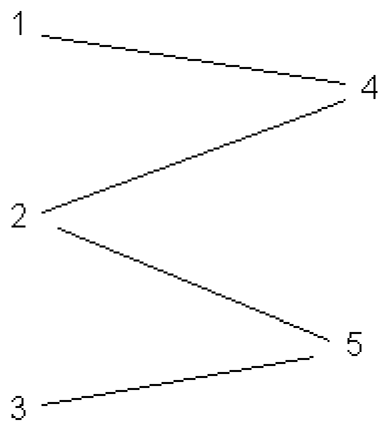

Figure 1: 5-Line Network

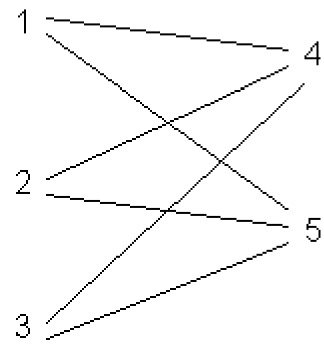

Figure 2: 3/2-Com Network

[15]
----
[16]

The method, evolutionary game theory, will assume a fiction: that there are populations from which the actors occupying a type of position are sampled and that effective strategies become more prevalent in those populations. Suppose we are interested in finding effective strategies for the positions in Figure 1. We will assume that there are two infinitely large populations, one for positions 4 and 5 and one for positions 1, 2, and 3. Actors from these populations are sampled to fill the positions in the network. If the actors holding positions are successful, their strategies become more prevalent in the populations from which they were sampled. A system of difference or differential equations will describe the development of the system over time, as favorable strategies in the two populations become more prevalent. Eventually, a stable distribution of prospering strategies may occur in both populations.

This assumed process of sampling from two populations is a fiction designed to find successful strategies, but one could imagine situations in which the assumption has some validity. In Figures 1 and 2 there are, implicitly, two types of actors, one of which is less frequent and should enjoy more power. The two populations *observe* the successes and failures of their population mates and *imitate* the more successful strategies. For example, suppose that there are more buyers than sellers in a market. Then buyers are likely to find themselves like those of positions 1, 2 and 3 in Figures 1 and 2. Buyers will observe the outcomes for their colleagues who follow various strategies. At the same time, sellers will be noting effective seller strategies. If members of the two populations imitate the more successful strategies, the dynamics described in this paper would have some reality.

Evolutionary Stability

The concept of *evolutionary stability* of strategies has been developed by biologists and others interested in situations in which successful strategies propagate themselves with greater frequency (Weibull 1995, Chapter 2). The concept is typically applied to symmetric non-cooperative games. Games are symmetric when there is only one kind of player and non-cooperative when binding agreements are impossible. Pairs of sampled strategies play independent games. Successful strategies are assumed to become more frequent in the population, and thus the whole distribution of strategies evolves over time.

The concept of evolutionary stability has been applied to the study of how cooperation develops in a prisoner's dilemma-like environment; interacting pairs of actors are each better off not cooperating, but both are better off if both cooperate. The players are assumed to be playing a *dyadic* 2-person game in which their rewards are a function only of their choices and the choices of their partners. The game is *symmetric* because all the actors face identical sets of choices with identical consequences.

The standard model for evolutionary stability cannot be used in network exchange games without modification. Network games are not 2-person; the strategies used by the players in a network can, potentially, affect all the other players. Network games are not symmetric (Cressman 1995). Players occupy different positions in networks and these positions define separate structures of interest. For example, in Figure 2 all the buyers (positions 1, 2, and 3) have 2 possible exchange partners and all the sellers (positions 4 and 4) have 3.

The standard definition of evolutionary stability describes the conditions under which a strategy cannot be invaded successfully by a small proportion of actors with another strategy. The standard definition is that x is an ESS strategy if (Weibull 1995: 37):

$$U(x,x) \geq U(y,x) \text{ for all alternatives y and if } U(x,x) = U(y,x)\text{, then } U(x,y) > U(y,y),$$

where $U(w,z)$ refers to the utility for strategy w when strategy z is its opponent.

In words, an ESS strategy x is a "best response" to itself; no other strategy does better against it than itself. Moreover, if some other strategy y does as well against x as x itself does, then x is a better response to y than y is itself.

An ESS strategy cannot be invaded by a small proportion , of any single other strategy. Let $T(x,\varepsilon)$ and $T(y,\varepsilon)$ be the utilities to strategies x and y if the proportions of the two strategies x and y are $1-\varepsilon$ and ε, respectively. Then

Equation 1:
$$T(x, \varepsilon) = (1-\varepsilon)U(x, x) + \varepsilon U(x, y)$$

Equation 2:
$$T(y, \varepsilon) = (1-\varepsilon)U(y, x) + \varepsilon U(y, y)$$

If $U(x,x) > U(y,x)$, then $T(x, \varepsilon) > T(y, \varepsilon)$ for all ε less than some critical value regardless of the values of $U(x,y)$ and $U(y,y)$. Or, if $U(x,x) = U(y,x)$, then $T(x, \varepsilon) > T(y, \varepsilon)$ if $U(x,y) > U(y,y)$. Thus, if a strategy is ESS and if growth is monotonically related to success in future generations, then x cannot be invaded by y. This is the motivation for the definition of evolutionary stability.

[17]
----
[18]

What we want is a definition that extends to network exchange games. To motivate the definition, consider the network in Figure 1. For the moment, we will assume that the strategies of positions 4 and 5 (who could be the buyers of a commodity) are fixed. We will assume that two strategies *A* and *B* compete at positions 1, 2, and 3 (the sellers of the commodity). $1-\varepsilon$ and ε are the proportions of the two strategies *A* and *B.*

There are eight patterns with the following probabilities. The subscripted values in the table refer to the expected rewards of the strategies. A_{ij} and B_{ij} are the rewards to strategies *A* and *B* respectively, if they occupy position *j* in pattern *i*.

Table 1:

Pattern	Position 1	Position 2	Position 3	Probability
1	A_{11}	A_{12}	A_{13}	$(1-\varepsilon)^3$
2	A_{21}	A_{22}	B_{23}	$(1-\varepsilon)^2\varepsilon$
3	A_{31}	B_{32}	A_{33}	$(1-\varepsilon)^2\varepsilon$
4	B_{41}	A_{42}	A_{43}	$(1-\varepsilon)^2\varepsilon$
5	A_{51}	B_{52}	B_{53}	$(1-\varepsilon)\varepsilon^2$
6	B_{61}	A_{62}	B_{63}	$(1-\varepsilon)\varepsilon^2$
7	B_{71}	B_{72}	A_{73}	$(1-\varepsilon)\varepsilon^2$
8	B_{81}	B_{82}	B_{83}	ε^3

This is a complex table. Inferences about the success of the two strategies may not be obvious. Computing the average rewards for the two strategies implicitly assumes that both are equally prevalent. Comparing the average rewards for the A strategy in configuration 1 and the B strategy in configuration 8 would also be an incomplete analysis; a strategy that does well if it monopolizes a type of position might be highly vulnerable to an invasion by another strategy. This paper suggests reasonable methods that take into account all the information in this table.

[18]
----
[19]

Let $T(A,\varepsilon)$ and $T(B,\varepsilon)$ be the expected rewards to strategies *A* and *B*, as functions of ε , the proportion of *B* strategies. Multiplying each occurrence of the *A* or *B* strategy in Table 1 by its probability, we find the expected rewards for the two strategies.[1]

Equation 3:

$$T(A,\varepsilon) = \frac{(1-\varepsilon)^2(A_{11}+A_{12}+A_{13})+\varepsilon(1-\varepsilon)(A_{21}+A_{31}+A_{22}+A_{42}+A_{33}+A_{43})+\varepsilon^2(A_{51}+A_{62}+A_{73})}{3}$$

Equation 4:

$$T(B,\varepsilon) = \frac{(1-\varepsilon)^2(B_{23}+B_{32}+B_{41})+\varepsilon(1-\varepsilon)(B_{52}+B_{53}+B_{61}+B_{63}+B_{71}+B_{72})+\varepsilon^2(B_{81}+B_{82}+B_{83})}{3}$$

First-order comparisons

If ε is small enough, then T(*A*, ε) > T(*B*, ε) if $A_{11}+A_{12}+A_{13} > B_{23} + B_{32} + B_{41}$. These are the corresponding terms multiplied by $(1-\varepsilon)^2$ in the numerators for T(*A*, ε) and T(*B*, ε). This means that, starting with a situation in which all three positions are using the *A* strategy, any positions adopting a *B* strategy in isolation are, on the average, worse off.

Second-order comparisons

Now suppose that $A_{11}+A_{12}+A_{13} = B_{23} + B_{32} + B_{41}$. Then the second set of comparisons, those involving ε (1-ε), can be used. T(*A*, ε) > T(*B*, ε) if $A_{21}+A_{31}+A_{22}+A_{42}+A_{33}+A_{43} > B_{52}+B_{53}+B_{61}+B_{63}+B_{71}+B_{72}$. This comparison involves starting with a situation in which just one position uses *B* and then adding a second position using *B* instead of *A*. The new user of *B* should be worse off, on the average. If the above is an equality, we could move on to the comparison involving the ε^2 terms in the numerators. No change would be required if the set of positions contained one, not three, positions. A strategy would be evolutionarily stable if its reward U was strictly greater than the rewards achieved by any other strategy at this position.

[19]

----

[20]

We assumed that the strategies of positions 4 and 5 were fixed. Now let us relax that assumption. The stability of a strategy in a set of positions will be conditional on the strategies being used by all the other positions in the network. One route would be simply to make all statements about ESS strategies conditional upon the strategies being used by all other positions in the network. The other direction, which seems to me more useful, is to define the evolutionary stability not of a single strategy used at a single position or set of positions but to talk about sets of strategies being evolutionarily stable, where one strategy is assigned to each type of position. From now on, the definition of evolutionary stability will refer to sets of strategies, one strategy assigned to each type of position. We will look at the stability of an assignment rather than the stability of a single strategy. Each type of position will be assumed to have a competing set of possible invaders. These sets of invaders may be different for different sets of positions.

Definition of Evolutionary Stability in Network Exchange Games

Let $\mathbf{P} = \mathbf{P_1}/\mathbf{P_2}/\ldots/\mathbf{P_m}$ be a partition of a set **V** of n vertices $\{v_i \mid 1 \leq i \leq n\}$ in an exchange network. Let g(i) = the number of the set for vertex v_i: $v_i \in P_{g(i)}$. Let $\mathbf{S}_k$ be the set of strategies available to all vertices in the set $\mathbf{P_k}$.

Consider an assignment of strategies to vertices such that that all vertices in the same set $\mathbf{P_k}$ are assigned the same strategy f(k) $\in$ $\mathbf{S_k}$.

Let U_i(fg(1), fg(2), . . . , fg(i), . . ., fg(n)) be the reward to v_i under this assignment of strategies.

The assignment is ESS if the following conditions hold true:

Equation 5:

$$\text{For every } k \text{ and every } x \in S_k,$$
$$\sum_{g(i)=k} U_i(fg(1), fg(2), \ldots, fg(i), \ldots fg(n)) > \sum_{g(i)=k} U_i(fg(1), \ldots, fg(i-1), x, fg(i+1), \ldots fg(n))$$

IF THE ABOVE IS AN EQUALITY

Equation 6:

$$\text{For every } k \text{ and every } x \in S_k,$$
$$\sum_{g(j)=k} \sum_{g(i)=k} U_i(\ldots, fg(i-1), x, \ldots, fg(j-1), fg(j), \ldots) > \sum_{g(j)=k} \sum_{g(i)=k} U_i(\ldots, fg(i-1), x, \ldots fg(j-1), x \ldots)$$

And so on.

[20]
----
[21]

Equation 5 says that assignment f is ESS if single positions of any type are worse off deviating from the assignment. Equation 6 says that if single positions of any type are no worse off deviating, then the second positions to deviate must be strictly worse off in total.

Replicator Dynamics for Exchange Networks

Replicator dynamics are systems of differential or difference equations designed to model the development of populations of strategies. In this paper I will be using the difference equation form. The dynamics are based on the assumption that successful (profitable) strategies increase their proportion in the population at the expense of less successful strategies. The replicator dynamics also require a modification to be suitable for network exchange games.

First let me describe the usual form for replicator dynamics. Let $x(t)$ be a vector of proportions in the population that pursue each strategy at time t. Let $u(e^i,x(t))$ be the rewards to strategy i at time t, conditional on $x(t)$. Let $u(x(t),x(t))$ be the average reward to all elements in the population at time t. Then (Weibull 1995: 123),

Equation 7:

$$x_i(t+1) = \frac{\alpha + u(e^i, x(t)}{\alpha + u(x(t), x(t))} x_i(t)$$

We have to adapt equation 7 to the special situation in exchange networks, in which there are two or more separate but interconnected competitions. Let $x^{(j)}(t) = (x^{(j)}_1, x^{(j)}_2, \ldots)$ be the probability distribution of strategies in population *j*, representing the *j*th type of position in the graph. Let $x(t) = (x^{(1)}(t), x^{(2)}(t), \ldots)$ be a vector containing all the separate probability vectors. $u(x^{(j)}(t), x(t))$ refers to the expected rewards in population j, which is going to be conditional on the distributions of strategies in all the other populations. Then, corresponding to equation 7,

Equation 8:

$$x^{(j)}_i(t+1) = \frac{\alpha + u(e^i, x(t))}{\alpha + u(x^{(j)}(t), x(t))} x^{(j)}_i(t)$$

[21]
----
[22]

In other words, strategies increase in proportion to the ratio of their gain to the average gain in their population, given the particular distributions of strategies in the other populations. Suppose, for example, that we have two strategies competing at each of the two sides in a 3-Line network.

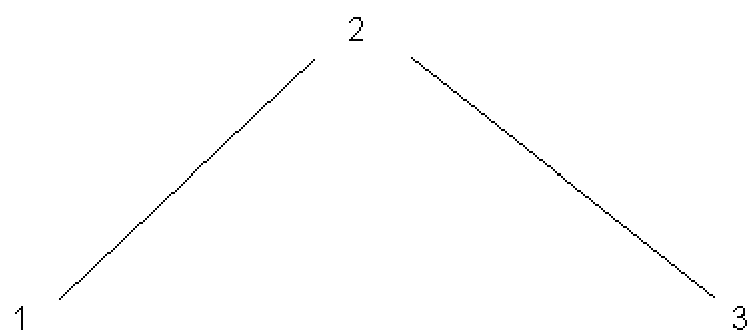

Figure 3: 3-Line Network

Let the two strategies available at positions 1 and 3 be *A* and *B*, while the two strategies available to position 2 are *C* and *D*. Let α be the probability of strategy *A* in the former population and γ the probability of *C* in the latter. The following table gives the values of all the strategy combinations and their probabilities.

Table 2:

Pattern	Position 1	Position 2	Position 3	Probability
1	A_{11}	C_{12}	A_{13}	$\alpha^2\gamma$
2	A_{21}	C_{22}	B_{23}	$\alpha(1-\alpha)\gamma$
3	B_{31}	C_{32}	A_{33}	$\alpha(1-\alpha)\gamma$

4	B_{41}	C_{42}	B_{43}	$(1-\alpha)^2\gamma$
5	A_{51}	D_{52}	A_{53}	$\alpha^2(1-\gamma)$
6	A_{61}	D_{62}	B_{63}	$\alpha(1-\alpha)(1-\gamma)$
7	B_{71}	D_{72}	A_{73}	$\alpha(1-\alpha)(1-\gamma)$
8	B_{81}	D_{82}	B_{73}	$(1-\alpha)^2(1-\gamma)$

[22]
----
[23]

In this situation, $x^{(1)} = (\alpha, 1-\alpha)$, $x^{(2)} = (\gamma, 1-\gamma)$, and $x = (\alpha, 1-\alpha, \gamma, 1-\gamma)$. Equations 9 and 10 give the expected rewards for strategies *A* and *B* respectively, and equation 11 gives the mean reward in the first population, the one from which positions 1 and 2 recruit. The expected rewards of the two strategies are conditional on the distribution of strategies in the other population of strategies for position 1.

Equation 9:

$$u(e^A, x(t)) = \frac{\alpha\gamma(A_{11}+A_{13})+(1-\alpha)\gamma(A_{21}+A_{33})+\alpha(1-\gamma)(A_{51}+A_{53})+(1-\alpha)(1-\gamma)(A_{61}+A_{73})}{2}$$

Equation 10:

$$u(e^B, x(t)) = \frac{\alpha\gamma(B_{23}+B_{31})+(1-\alpha)\gamma(B_{41}+B_{43})+\alpha(1-\gamma)(B_{63}+B_{71})+(1-\alpha)(1-\gamma)(B_{81}+B_{83})}{2}$$

Equation 11:

$$u(x^{(1)}(t), x(t)) = \alpha u(e^A, x(t)) + (1-\alpha)u(e^B, x(t))$$

Equations 12 and 13 hive the expected rewards for strategies C and D in the second population, and equation 14 gives the expected reward in the second population.

Equation 12:

$$u(e^C, x(t)) = \alpha^2 C_{12} + \alpha(1-\alpha)(C_{22}+C_{32}) + (1-\alpha)^2 C_{42}$$

Equation 13:

$$U(e^D, x(t)) = \alpha^2 C_{52} + \alpha(1-\alpha)(C_{62}+C_{72}) + (1-\alpha)^2 C_{82}$$

Equation 14:

$$u(x^{(2)}(t), x(t)) = \gamma u(e^C, x(t)) + (1-\gamma)u(e^D, x(t))$$

[23]
----
[24]

RESULTS

Some Preliminary Results

The approach used in this paper is to base the mathematical analysis (evolutionary stability and replicator dynamics) on the "data" generated from computer simulations. In the simulations, there are three stages to each game. In the first stage, positions make offers to all those to whom they are connected. All offers are equal and in the first game are randomly sampled from a rectangular distribution with integer values between 9 and 15 inclusive. Strategies have different ways of adjusting offers on subsequent games. In the second stage, offers are accepted, and in the third stage the acceptances are confirmed by the initial offerers. If there are any remaining connected positions that are not involved in transactions, the process is repeated with the remaining positions. Each set of actors plays twenty games. Each simulation result is based on thirty sets of actors.

Example I

Strong Power Networks

In some exchange networks there are, implicitly or explicitly, two unequally large categories of positions. The members of the two categories exchange only with members of the other category (the graph is *bipartite*). Because there are surplus members of the larger category, members of the smaller category have an advantage.[2] The networks in Figures 1 and 2 are of this type; positions 1, 2, and 3 exchange only with positions 4 and 5, and visa versa.

I begin with the assumption that members of the larger and weaker category would have an interest in restricting their competition with one another; it is the competition between themselves that reduces their bargaining power. Therefore, I contrasted the following two strategies within some examples of this type of network.

Up/Down (U/D or U)

This is the standard strategy used in computer simulations of exchange networks. If a position is left out of exchanges in one game, it raises its offers to others by two points in the next game. If it is included in an exchange in one game, it lowers its offers to others by one point in the next game.

Tit-For-Tat (TFT or T)

This strategy is sensitive to the magnitude of offers made by selected other positions, which we will call *alters*. The initial offers made by TFT are low, between one and eleven points (out of 24). TFT raises its offers by one point when the alter has made a higher offer than TFT in the last game, it lowers its offers by two points if the alter has made a lower offer than TFT in the last game, and it leaves its offers unchanged if it and alter have made equal offers. If TFT and the alter are the only positions competing to be included in an exchange with one common exploiter, then TFT is quite close to U/D, the major difference being that TFT responds more strongly to inclusion and less strongly to exclusion. Two TFT players facing one common exploiter will make lower and lower offers, whereas two U/D players will make higher and higher offers. The alters can be located anywhere in the network; they do not have to be positions with which TFT shares potential partners.

For example, in the 3-Line network in Figure 3, suppose that positions 1 and 3 are using the TFT strategy while position 2 is using the U/D strategy. Suppose, moreover, that position 2 makes initial offers of 12 points to positions 1 and 3, while positions 1 and 3 make offers of 8 and 6 points to position 2, respectively. Position 2 accepts position 1's higher offer. Position 1, of course, accepts 2's offer, and the compromise agreement (the average of the offers) is 10 points for 2 and 14 points for 1. In the next game, position 2, using U/D, reduces his offers to 11. Position 1 reduces his offer by 2 points to 6 and position 3 increases his offer by one point, to 7. This time position 2 accepts 3's offer, resulting in a compromise deal of 10 points to 2 and 14 points to 3. The long-term consequence is that position 2 never enjoys his potential positional advantage.

Using the 3/2-Com network and other similar networks, we found a surprising result. Table 3 shows the average outcomes for the positions in the 3/2-Com network in a series of simulations using all the distinct combinations of the U/D and TFT strategies. Let $X_1X_2X_3|X_4X_5$ mean that each position *i* was using strategy X_i.

Table 3:

Number		1	2	3	4	5
1	UUU|UU	93	93	93	342	342
2	UUT|UU	139	139	56	313	313
3	UTT|UU	215	121	121	251	251
4	TTT|UU	112	112	112	245	245
5	UUU|UT	74	74	74	369	369
6	UUT|UT	109	109	40	348	354
7	UTT|UT	263	268	268	263	268
8	TTT|UT	229	229	229	131	143
9	UUU|TT	53	53	53	401	401
10	UUT|TT	72	72	26	395	395

11	UTT|TT	113	66	66	358	358
12	TTT|TT	160	160	160	240	240

Table 3 shows that there were just two evolutionarily stable configurations of strategies for 3/2-Com network. The first was TTT/TT (row 12). Row 11 shows that an isolated U/D strategy among the less powerful positions would do worse (from 160 to 113), and row 8 shows that an U/D strategy among 4 or 5 when the alter was continuing to use TFT would also hurt (from 240 to 131). As we expected, this combination was considerably more advantageous for the less powerful positions than UUU|UU - 160 in the former and only 93 in the latter. If this strategic configuration were used, the existing theories, which appear to be based on actors who use the U/D strategy, would seriously overestimate the power of actors 4 and 5.

However, there was another ESS combination: UUU|TT, in which the powerful positions restricted competition with one another while the less powerful positions continued to compete with one another for inclusion. The evolutionary stability of this combination can be seen by comparing row 9 with rows 10 and 4. In this case, the members of the less powerful class were even worse off. If this were the outcome, the standard experiments would have underestimated the potential power of the more powerful actors.

Exploring this further, the replicator dynamics were used to infer the history of this system under various assumptions about the initial distribution of TFT and U/D strategies among the two classes of actors. Initial proportions of the U/D and TFT strategies for the powerful players 4 and 5 and the weak players 1, 2, and 3 were set between .10 and .90 and increments of .1 (this is 9 times 9, or 81, possibilities). Roughly speaking, no matter what the proportion of U/D strategies was among players 4 and 5, the surviving strategy among players 1, 2, and 3 was determined by the initial distribution of strategies in their population; if more than half the weak players initially used TFT, TFT won out in the end, but if more than half the players used U/D, all players were U/D in the end. However, TFT was *always* the winning strategy among players 4 and 5, no matter what the initial proportions of strategies in the two populations.

For example, Figures 4 and 5 show the proportions of U/D strategies among the strong players 4 and 5 and the weak players 1, 2, and 3 over 30 generations of play when the initial proportions of the U/D strategy among the strong and weak players are .80 and .50, respectively. Figure 4 shows that the proportion in the strong population using U/D regularly declines, while Figure 5 proportion of strong players using U/D increases.

In short, there is an unconditional drift in favor of the TFT cooperative strategy among the strong players. On the other hand, which strategy wins out among the weak players ill depend on the initial proportions using the two strategies, and the success of TFT is not guaranteed.

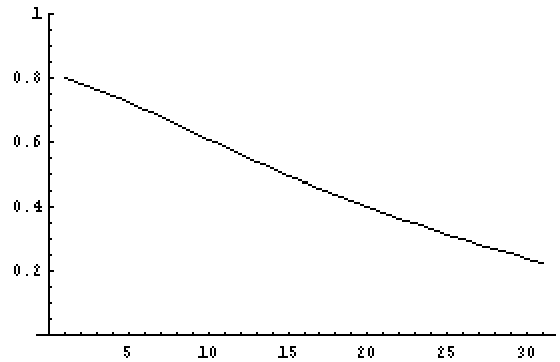

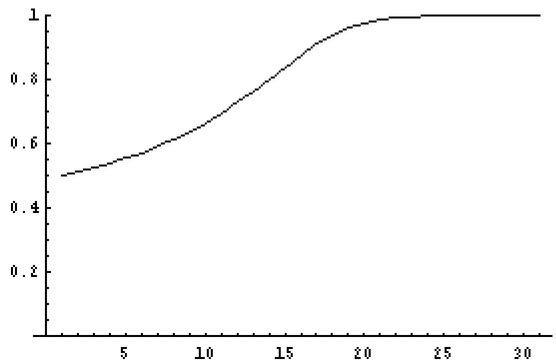

Figure 4: Powerful Positions in 3/2-Com **Figure 5: Weak Positions in 3/2 Com**
Network Using Up/Down Strategy **Network Using Up/Down Strategy**

The implications are ironic. TFT, a type of cooperation among the members of a class, might be the only way in which the members of the weaker class can improve their outcomes. However, the strategy is actually more easily inaugurated by the stronger players, because they are a smaller class and thus find it easier to establish cooperative relations. The larger class of players must have an initial dominance of cooperation over the self-interested U/D strategy for TFT to predominate.

[28]
----
[29]

Example II
The following network, called an "hourglass" or a "kite," has interested researchers. Experimental evidence suggests that the middle position has a slight advantage in terms of earnings per inclusion, but tends to be excluded from exchanges. Bonacich and Bienenstock (1992) note that this is a *coreless* game in which all outcomes can be undermined by excluded members. It would seem that in such a network, where all positions are vulnerable to exclusion and no positions are powerless, there would be a premium for forming stable, committed exchange relations. Therefore, I contrasted two strategies in this network: UP/DOWN and LOYALTY. The LOYALTY strategy encourages stable relations between trading pairs by making higher initial offers to trading partners in the previous game.

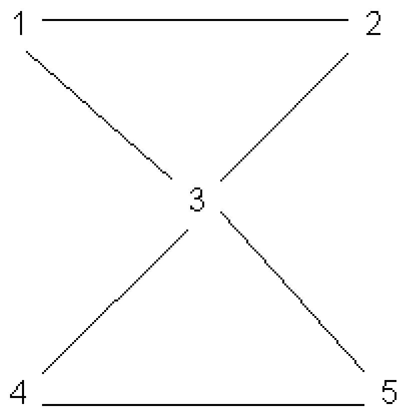

Figure 6: Hourglass Network

In this example, let there be just one equivalence class: **P** = /12345/. There are only twelve essentially different patterns. The following table shows the results for all these patterns. The patterns are described in terms of what strategies, U or L, are played by positions 1 through 5.

[29]
----
[30]

Table 4:

	1	2	3	4	5
UUUUU	195.2	195.2	177.9	195.2	195.2
LUUUU	192.13	146.73	164.07	228.21	228.21
UULUU	182.78	182.78	228.9	182.78	182.78
LLUUU	224.04	224.04	158.73	176.60	176.60
LULUU	206.17	132.47	244.04	188.45	188.45
LUULU	107.44	254	237.23	107.44	254
ULLUL	175.3	175.06	259.27	175.3	175.06
ULULL	146.87	196.9	143.53	236.35	236.35
UULLL	142	142	226.5	142	142
LLULL	211.38	211.38	114.47	211.38	211.38
ULLLL	127.6	193.07	193.97	222.68	222.68

LLLLL	199.44	199.44	162.23	199.44	199.44

[30]
----
[31]

The only ESS set in this situation is for all positions to use the Loyalty strategy. For position 3 to change would result in a reduction in earnings from 162.23 (row 12) to 114.47 (row 10). For any other position to change would result in a reduction from 199.44 (row 12) to 127.6 (row 11). Thus, the expected earnings would also be less for a single U/D in any position. The replicator dynamics also show the superiority in this network of the Loyalty strategy. When more than about one third the population uses the Loyalty strategy, this strategy predominates in the long run.

The superiority of the LOYALTY over the UP/DOWN strategy present in the hourglass network does not appear to hold in all coreless networks. Consider, for example, the following network.

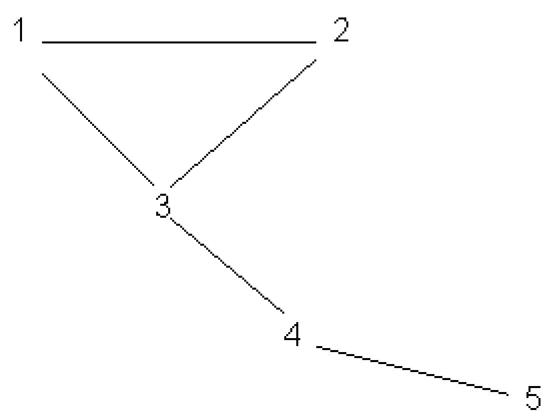

Figure 7: Another Unstable Network

This network is also coreless; any agreement pattern can be disrupted by excluded positions. Unlike the Hourglass network, there is one position, 4, which should never be excluded. Contrasting the LOYALTY and UP/DOWN strategies in this network, the only ESS configuration is one in which positions 1, 2, and 3 use LOYALTY while 4 and 5 use UP/DOWN.

In summary, it appears that in coreless networks, where many positions face exclusion but no positions lack at least some power, a strategy like LOYALTY, which encourages commitment in order to avoid exclusion, should play a role. However, more work remains to be done on which positions will use LOYALTY.

[31]
----
[32]

Example III
Consider the following 4-line network.

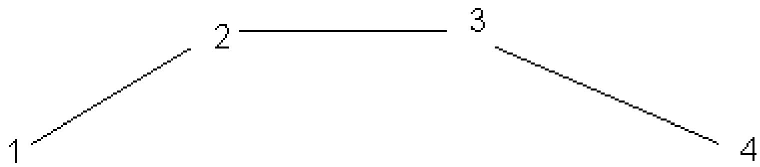

Figure 8: 4-Line Network

Experimental evidence, and existing theories, suggests that positions 2 and 3 should enjoy a slight (or "weak") power advantage over positions 1 and 4. Using a simulation in which all positions use UP/DOWN, the rewards for positions 1 and 4, over 20 games, average 207.7 while the middle positions 2 and 3 average 262.4. Is there anything that positions 1 and 4 can do to reduce this power disadvantage? One possibility would be that positions 1 and 4 play TFT using each other as alters; each position would then copy the offers made by its alter while particularly rewarding low offers to the middle and more powerful positions. The results of a computer simulation show that this is not to their advantage even if both positions play TFT. If both 1 and 4 use TFT, they average just 96.52 points. In fact, it is not clear what positions having a weak power disadvantage can do to improve their lot.

DISCUSSION

In this paper I have suggested an approach toward exploring the relation between position and a network and optimal strategy. Ever since Cook and Emerson (1978), researchers have assumed in their theories or in their simulations that all actors, regardless of position in the network, use the same simple strategy: if they are excluded, they raise the value of their offers to others in the succeeding game, and if they are included, they then lower the value of the offers they make to others. This simple strategy, although suitable for inexperienced laboratory subjects, may be overly simple for some situations outside the laboratory. If more experienced actors use other bargaining strategies, current theories may be inaccurate.

Non-cooperative game theory permits us to incorporate into our models relations between actors who do not exchange with one another but whose choices affect one another. For example, weak positions have an interest in limiting their offers to a common exploiter. The concept of evolutionary stability, adapted to network games, provides a criterion for finding stable configurations of strategies and possible insights into how sophisticated strategists may play these games.

FOOTNOTES

1. Equations 3 and 4 are easy to calculate. The probability that an actor using an A strategy occupies position 1 is 1/3 (the probability that he occupies positions 2 or 3 is also 1/3). The probability that A strategies fill positions 2 and 3 also is $(1-\varepsilon)^2$. When this happens, the reward to the actor using the A strategy is A_{11}. This accounts for $(1-\varepsilon)^2 A_{11}/3$ in equation 3. The other are accounted for similarly.

2. For more complete and accurate characterizations of networks with large power differences, see Markovsky, Willer, and Patton (1988) or Bonacich (1996).

REFERENCES

Bonacich, Phillip (1996). "A Psychological Basis for a Structural Theory of Power in Exchange Networks." American Sociological Association Annual Meetings. New York.

Bonacich, Phillip and Elisa Jayne Bienenstock (1992). "The Core as a Solution to Exchange Networks." *Social Networks*, 14:231-243.

Bonacich, Phillip and Elisa Jayne Bienenstock. (1997). "Strategy in Exchange Networks: Accommodation versus Exploitation." Pp. 351-361 in Jacek Szmatka, John Skvoretz, and Joseph Berger, eds., Status, Network, and Structure: Theory Development in Group Processes. Stanford: Stanford University Press.

Burke, Peter J (1997). "An identity model for network exchange." *American Sociological Review*, 62:134-150.

Cook, K.S., R.M. Emerson, M.R. Gillmore, and T. Yamagishi (1983). "The Distribution of Power in Exchange Networks: Theory and Experimental Results." *American Journal of Sociology,* 89:275-305.

[33]
----
[34]

Cook, Karen S., Mary Gillmore, and Toshio Yamagishi (1986). "Point and Line vulnerability as Bases for Predicting the Distribution of Power in Exchange Networks: Reply to Willer." *American Journal of Sociology,* 92:445-48.

Cook, Karen S. and Toshio Yamagishi (1992). "Power in Exchange Networks: A Power-dependence Formulation." *Social Networks*, 14:245-265.

Cressman, Ross (1995). "Evolutionary game Theory With Two groups of Individuals." *Games and Economic Behavior*, 11:237-253.

Markovsky, Barry (1987). "Toward Multilevel Sociological Theories: Simulations Of Actor and Network Effects." *Sociological Theory,* 5:101-117.

Markovsky, Barry, David Willer, and Travis Patton (1988). "Power Relations in Exchange Networks." *American Sociological Review,* 53:220-236.

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

Thye, Shane, Michael J. Lovaglia, and Barry Markovsky (1997). "Responses to Social Exchange in Networks." *Social Forces*, 75:1031-1047.

Weibull, J` rgen (1995). *Evolutionary Game Theory.* Cambridge: MIT Press.

Whitmeyer, Joseph (1997). "Applying General Equilibrium Analysis & Game Theory to Exchange Networks." *Current Research in Social Psychology*, 2:13-23, http://www.uiowa.edu/~grpproc.

Yamaguchi, Kazuo (1996). "Power in Networks of Substitutable and Complementary Relations: A rational Choice Model of Power Centralization." *American Sociological Review*, 61:308-332.

AUTHOR BIOGRAPHY

Phillip Bonacich is Professor of Sociology at UCLA. His current research interest is power and coalitions within exchange networks. Email address: bonacich@soc.ssnet.ucla.edu