DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

NETWORKS: AN EXPERIMENTAL STUDY

Charles R. Plott

Steven Callander
Northwestern University

SOCIAL SCIENCE WORKING PAPER 1156

July 2002

Networks: An Experimental Study*

Steven Callander†
Northwestern University

Charles R. Plott‡
California Institute of Technology

Abstract

This paper reports on an experimental investigation of the evolution of networks and the individual decision making processes that guide it. Since there is no history of experimental work on network formation, part of the paper is devoted to the formulation of problems that can be examined experimentally. The results are that networks, composed of decentralized decision makers, are capable of overcoming complex coordination and learning problems and converge to stationary configurations. While stationarity is frequently observed, such an achievement is not guaranteed and when it doesn't occur significant and persistent inefficiencies can result. The models of equilibration based on the principle of Nash stability are more reliable than models based on the alternative principles of efficiency seeking or focalness of the network configuration. However, individual decision making within networks is not in accordance with the simple decision rule of Nash best response. Instead we observe complicated strategies that appear to trade off short term profits in order to signal to, and teach, other agents the strategies required for long term profit maximization.

*The financial support of the Laboratory of Experimental Economics and Political Science at Caltech is gratefully acknowledged.

†Assistant Professor of Managerial Economics and Decision Sciences, MEDS department, Kellogg School of Management, Northwestern University.

‡Professor of Economics and Political Science, Division of the Humanities and Social Sciences, California Institute of Technology.

1 Introduction

Patterns of economic and social phenomena that can be described as networks are pervasive. However, only in recent years have attempts been made to understand networks as special phenomena that are governed by their own set of principles. The emerging literature has suggested models of both strategic behavior within the contexts of networks in addition to broad models of the evolution of network formation.[1] Unfortunately, the inherent complexity of network models limits both the theory as well as intuition that can serve as a guide to further theory. We have based this study on the hope that from the examination of simple cases the principles of behavior at work can be identified and on them one can build realistic and useful models of network formation. It is in this pursuit that we turn to experiments to help us understand these phenomena.

At the outset it is important to emphasize that the experiments reported here are "exploratory". No previous network experiments exist on which to build. The number of variables is staggering and there is no obviously best configuration with which to start. The theory is itself at a most elementary state with no motivations or applications from the field that might narrow a set of questions. The state of the theory together with the number of variables suggest that a "measurement" approach to experiments will not work. It makes little sense to measure the effects of variables when neither the theory nor the importance of the variables are well established. So, in a sense part of the problem addressed in this study involves questions about where to start and in what directions one might push. The overall development of the paper is designed to explain the considerations that were made in the approach so the study can be used as a benchmark for others who feel that alternative directions might be more productive.

At the beginning of the research the most fundamental questions about behavior in the context of network formation were open. Do networks happen? Is their formation predictable? Can and do they settle down or converge to stable networks from which few or no changes take place? Is there a principle behind their evolution, and if so what is it? Indeed, it remained to be confirmed that game theory and standard economic techniques are appropriate in the analysis of networks. That is, it had not been established that networks are indeed an

[1] For static papers see Jackson and Wolinsky (1996), and Dutta and Mutuswami (1997). An application to industrial organization is Kranton and Minehart (2001). For dynamic papers see Bala and Goyal (2000), and Jackson and Watts (2002).

economic phenomenon governed by classical economic principles. The experiments reported in this paper have been designed to address these fundamental questions, and initiate an empirical foundation on which a rigorous theory of networks can be established.

The answers to some of the questions described above and investigated in this paper may seem obvious, or even appear axiomatic. However, they are not guaranteed, for as will be seen in many ways the network problem is similar to the public goods problem that can involve free riding and coordination failure.[2] In the theoretical literature on public goods many assumptions are the same as those employed in the network literature. Most notable being the assumption of full rationality and common knowledge on behalf of the agents, and consequently the use of Nash equilibria to predict outcomes.

We find that networks do spontaneously emerge and that they are capable of converging to configurations and remain stable from round to round. However network formation and behaviors exhibited while they are forming are complex. More specifically, we find that convergence to stable networks does not always occur, but that when it does occur the stable network is predictable. Significantly, the dynamics of network formation does not exhibit monotonically increasing efficiency. This implies that initial inefficiencies caused by miscoordination or misunderstandings may become institutionalized.

Approximately speaking, at certain critical points in network dynamics the coordination, bargaining and free rider aspects of individual decision making become aligned and stability is achieved. At these points it appears that all decision makers become aware of which network is best for them, and are aware that other agents are aware of this, and so on ad infinitum. In non-convergent networks this coordination of beliefs simply does not occur and significant inefficiency results.

We find the principle behind convergence and network dynamics to be Nash stability (though not necessarily strict Nash), and not efficiency or focalness. These findings allow us to answer one of our fundamental questions, that game theory and its concepts of equilibria are appropriate to be applied to the network problem.

The paper is developed as follows. Section 2 presents a brief introduction to the formal concept of networks. It also contains a summary of the design variables that have been used in the literature and the theoretical models of

[2] As Ledyard (1995, p. 172) was left to conclude in a review of public goods experiments, "if these experiments are viewed solely as tests of game theory, that theory has failed."

networks to be used in the experiments. Section 3 presents the overall experimental design and the particular features of both Series 1 and 2 of experiments. Section 4 presents the results for both series of experiments and section 5 concludes the paper.

2 Experimental Setting and Network Models

As mentioned in the introduction, the experimental design resides in the domain of "exploratory" methodology. The approach is dictated by both the lack of previous experiments together with the abundance of variables and a corresponding incompleteness of theory. The approach is to explore proposed general principles that the literature suggests might govern network development and evolution. Thus, the experimental environments are chosen to include environments in which the operation of suggested principles might be detected. Even if incomplete, theory aids in identifying variables that might interact in understandable ways. For these choices the limitations placed on experiments by technology must also be considered.

2.1 Environments

We study networks in which each node is a separate individual agent. Each agent unilaterally chooses the links they form between themselves and other agents.[3] At each node exists a 'piece' of information (this can be interpreted alternatively as a good or service) that has the capacity to flow through the network without 'decay'. All information that exists at the node to which a connection is being made is passed to the node that initiated the connection. The benefit is received by each node through which the information passes, including the node of origin (the benefit can be received only once). The absence of decay implies that the value of information is independent of the number of links that it passes through before reaching an agent. Links are assumed to be one way and are paid for by the connecting agent, who receives the benefit. Each agent is free to connect to any other agent, or combination of agents, that he chooses. The timing and knowledge on which agent decisions can be based are discussed later in this and the succeeding section.

[3] There are several other literatures that proceed under the banner of "networks" that differ significantly and will not be considered here. These include the literature on airline networks (see Hendricks, Piccione, and Tan (1999)) and network externalities.

Each of the experiments involved six agents, or nodes. An example of a six agent network is depicted in Figure 1.

** Figure 1 about here.

In the Figure 1 network each agent chose to implement only one link. The direction of the arrow points to the agent who constructed the link and receives the information flow. In this particular network each agent receives every available piece of information. This can be seen as the sequence of links traces continuously through each node in the network.

Agents within the network receive continuously updated information about the structure of the network in which they are operating. Thus, from a modeling perspective a natural beginning would be a model in which the structure of the network is common knowledge. Of course, the physical realities of presenting such information to subjects must be acknowledged. Exactly how that can be done and how the information must be organized will discussed in the experimental procedures section. Nevertheless, we proceed on the assumption that each agent has full information about the size and composition of the network, as well as the links selected by other agents. There is no communication of any kind between the agents other than through their link choice.

The experiments consisted of a series of rounds. In a round each agent selected links to connect. Links were reliable (i.e., never failed) though they lasted for only one round.[4] In each round therefore the network started anew. In one of the series of experiments decisions about links were made in real time with continuous feedback.

2.2 Network Structures and Models of Network Formation

In this section we introduce three broad principles of network formation that are suggested by the literature: Nash equilibrium, efficiency and focalness. The following section will address models that have a very micro origin with the behavior of the individual. The approach we take is to develop the models in terms of specific experimental environments, which allows us to simultaneously discuss the model and its predictions for the cases to be studied. Theory and models typically address a network configuration that is in "equilibrium" as a fixed point in some model or some solution concept from games. Operationally

[4] Bala and Goyal (2000a) explore the implications if links are not reliable.

and empirically, equilibrium itself is typically regarded as a stable configuration, a network that experiences no changes in connections due to individual actions.[5] Theory also suggests changes in network configurations as individuals have the opportunity to adjust connections. Natural questions to pose, therefore, are whether or not the movements of network changes are "in the direction of" the equilibrium of some theory. In this sense the theories can serve as both models of stable configurations as well as individual action and movements of configurations.

Nash equilibrium refers to the well known concept from the theory of games. If, given what other agents are doing, an individual can improve personal gains by some change then the particular network configuration is not a Nash equilibrium. We will consider both strict and weak definitions of Nash equilibrium.

Efficiency refers to the proportion of gains received by all agents relative to all potential gains, without regard to the individuals that receive the gains. If gains are the maximum possible then the system is at 100% efficiency. Such a calculation reflects both the distribution of information around the network and the cost of the formation of the network.

Focalness, is not usually considered in formal models since in the world of abstract reasoning there is not typically a sense of position that may be used as a coordinating device. In contrast, however, subjects in the experiments discussed here very much exist in geographical space and, as originally discussed by Schelling (1960), this space may be used as a coordinating device by the agents. The application of the concept for purposes of this paper reflect the positions that subjects might have been placed in the room, the positions in which data were put on the chalkboard or the positions in which individuals appeared in network representations on screens.

Four separate parameter sets are employed to illustrate the differences in these principles. Despite their varying motivations, these three principles are not always distinct in their network predictions. Table 1 describes four sets of network parameters in which the predictions of these principles converge and diverge. The Nash equilibrium, efficient, and focal networks for these different parameter values are described in Table 2.

** Tables 1 and 2 about here.

[5] Some environments, such as Jackson and Wolinsky (1996), require bilateral action for links to form and employ the analogous concept of "pairwise stability" to describe equilibrium.

Parameter sets 1 and 2 involve symmetric costs and benefits and lead to identical predictions. As shown by Bala and Goyal (2000), the "wheel network" is uniquely efficient and strict-Nash for these parameters. As the name suggests, a wheel network requires each agent to connect only one link from another agent such that these links form one long chain. It is important to note that this chain need not appear as a wheel when depicted graphically. Examples of wheel networks are given in Figure 2 below and Figure 1 earlier. This architecture[6] is efficient as all agents receive maximum value for the cost of only one link. While these two configurations are equivalent with respect to efficiency and Nash equilibrium, focalness draws a distinction between them. Agents in our experiments are seated as depicted in the figures and are assigned consecutive numbers as indicated. Therefore, we assume that the wheel in Figure 2 is focal (the counter-clockwise wheel), whereas the wheel in Figure 1 is not.

** Figure 2 about here.

It is important to note that even though the wheel network is the unique strict Nash equilibrium, there exists many weak Nash equilibria. Figure 3 provides an example with eight links.

** Figure 3 about here.

Parameter Set 3 differentiates further between the potential equilibrating principles. By making connections between neighbors[7] more expensive than other connections, the focal wheels, which rely exclusively on neighborly links, are no longer efficient. The efficient configurations are wheels in which there are no neighborly links. An example of such a wheel is given in Figure 4. Note that Figure 1 depicted above is not efficient despite being a non-focal wheel as there are some neighborly links in this configuration.

** Figure 4 about here.

The final parameter set alters these predictions further and to a degree allows the predictions of the efficient and Nash equilibrating principles to be

[6]Two networks have the same architecture, as defined by Bala and Goyal (2000, p. 1182), if one network can be obtained from the other by permuting the strategies of agents in the other network.

[7]Neighbors are defined as geographically adjacent agents. For example, the neighbors for agent 6 are agents 5 and 1.

separated. The asymmetric cost structure implies that it is cheaper for agent 1 to connect a certain link than it is for any other agent. This incentive is so strong that the wheel architecture is no longer efficient, and instead a star network centered on agent 1 is the uniquely efficient network, as well as being a strict Nash equilibrium.[8] Significantly, however, the wheel network is still a strict Nash equilibrium. The star network is depicted in Figure 5.

** Figure 5 about here.

2.3 Models of Individual Behavior and Network Dynamics

Individual behavior is a compelling area to explore for not only understanding the dynamics of existing networks, but also predicting the existence and behavior of all networks. There are many theories of individual behavior, which become increasingly complex in the network environment. We will focus here on two such models, involving varying degrees of strategic choice: (Nash) best response and simple strategic behavior.

Best response, studied in a network context by Bala and Goyal (2000), assumes that agents naively and myopically respond to the network environment. More formally, in a model of simultaneous choice, this decision rule supposes that each agent chooses the set of links that maximizes his payoff given the current link selections of other agents. Thus, it is myopic in that future payoffs are ignored, and naive in that adjustments by other agents are not anticipated.

[8] The proofs of these claims are quite simple. To see that the star is a strict Nash equilibrium consider firstly agents 2-6. All of these agents are receiving all pieces of information at the cost of a single link from the cheapest source. Thus, they are playing a strictly optimal strategy. Now consider agent 1. He is receiving all pieces of information but at the expense of five links (recall he must pay the adjustment fee). However, if he dropped a link then he would lose a piece of information. Thus, he is also strictly optimizing and the star is a strict Nash equilibrium.

To see that this configuration is uniquely efficient suppose that there are links that do not include agent 1. Say agent 4 is connected from 5. This link costs $.20. Consider an alternative network in which this link is omitted and replaced by a link from 5 to 1 and from 1 to 4. These links cost at most $0.10 (as they may already exist). Therefore, this alternative network is cheaper and weakly increases information flow. Consequently, the original network cannot be efficient. It is easy to see that networks involving a subset of links in the star network are also inefficient (just add links of the star that are missing). Therefore, the star centered on agent 1 is uniquely efficient.

In the model of Bala and Goyal (2000) agents best respond though with a degree of inertia (i.e., with some probability they do not change their selection from one round to another).[9] In a remarkable result, Bala and Goyal show that, despite the myopic and naive behavior of agents, Nash equilibrium social communication networks evolve very rapidly. This result is perhaps best interpreted as a benchmark with respect to the evolutionary capabilities of networks: that with self-interested and boundedly rational agents convergence to stable networks is possible.

Simple Strategic Behavior is a model based on the possibility that agents act with a greater degree of sophistication than allowed for by the best response decision rule. It may be suspected that agents make choices with more foresight, as well as learning and even teaching optimal strategies to themselves and other agents. Unfortunately, given the complexity of network environments, even the simple structure studied here, the application of complex decision rules does not provide much insight or testing power. Therefore, we will consider here only one simple decision rule tailored to the network environment. **Simple Strategic Behavior (SS)** requires agents to connect only one link, and that this link be their part of a focal wheel network. We denote the behavior by **(SScc)** when the network is the counter-clockwise wheel, and **(SScw)** when the network is the clockwise wheel.

The logic behind the **SS** decision rule is the following. For many parameter values, including sets 1 and 2 from Section 2.2, the wheel network is not only optimal for the agents as a collective, but it is also optimal for every agent individually. Further, the clockwise and counter-clockwise wheels are in many respects focal. Therefore, a reasonable expectation would be that agents are moving towards these configurations even if the corresponding link selections are not in their short term interests. These choices would increase the chances of coordination on an optimal network, as well as teach other agents the 'optimal strategy.'[10] These calculation may not necessarily lead an agent to conform to SS behavior, as, for example, he may add an additional link for insurance purposes. However, simple strategic behavior captures the basic intuition of these arguments and intentions, and as we will see later, performs well in describing the choices of agents in network environments.

[9] Bala and Goyal (2000) also analyse environments with two way links and decay, features that for simplicity will not be considered here.

[10] This notion is similar in spirit to recent work on 'strategic teaching' by Camerer, Ho and Chong (2002).

3 Experimental Procedures

A total of twelve experiments were performed. Each experiment consisted of six inexperienced subjects recruited from the undergraduate and graduate population of the California Institute of Technology. As summarized in Tables 3 and 4 the experiments consist of five experiments in Series One and seven experiments in Series Two, and followed the design principles described in Section 2.1. The parameters and procedures of Series One were heavily influenced by the model presented by Bala and Goyal (2000). The design of Series Two reflects the experiences of Series One.

** Tables 3 and 4 about here.

Subjects were randomly assigned to locations so friends arriving together tended not to be sitting next to each other. Each subject was assigned an identification number from 1 to 6. Instructions were read to subjects (see Appendix) and the subjects were given a practice exercise (without payment) and tested before the experiment began. The experiments consisted of rounds during which subjects could make connections to any other subject at a cost. The profits to a subject were the value of the information received minus the cost of connection. A network was deemed to have "converged" if the same configuration was chosen in three consecutive rounds.

3.1 Series One

Series One experiments were performed manually and payoffs were calculated using a physical process. In each round every agent recorded their link selection and this was submitted to the experimenter. They then placed in front of themselves, in full view of all agents, physical signs corresponding to their selections. The benefits of connections from the networks were then easily computed with each individual adding the signs exhibited by each node to which the individual was connected. This process quickly iterated to an accurate computation of the information accruing to each node. The network chosen was then drawn on the board at the front of the room. Agents computed their earnings and the round was complete.

To establish the power and capabilities of networks, parameter set 1 was used in all Series One experiments (see Table 1), and thus the wheel was the unique efficient, and Nash equilibrium network architecture. A random stopping rule was employed whereby between 10 and 20 rounds were possible.[11] There was an increasing chance of stopping as more rounds were played. We refer to this rule as stopping Rule 1. The probabilities of stopping at any point, along with those for Rule 2 which was used in Series 2 experiments, are detailed in Table 5.

** Table 5 about here.

3.2 Series Two

Several changes were made to this design for Series Two. Firstly, the process moved to computers and agents were partitioned off in different segments of the laboratory. This change was to alter if not remove the focalness from the experimental environment, and test the robustness of convergence with isolated individual decision making. The second change was that decisions were made continuously over two minute rounds, and that the link choices could be adjusted repeatedly in real time. Further, all agents were continuously updated as to the current selections of their fellow agents. However, it was only at the end of periods that the link connection fee was charged and benefits accrued. Agents were also charged an adjustment fee of 5 cents each time they added or subtracted a link during each round.

These changes were made to facilitate convergence and avoid the coordination problems of Series One (often between the clockwise and counter-clockwise wheel). We expected that these coordination problems would become far more problematic without the focalness provided in Series One. The adjustment fee was included to ensure that link selections during the rounds were meaningful signals and not purely cheap talk.

Series Two experiments employed random stopping rule 2, with between 15 and 20 rounds taking place, again with an increasing probability of stopping as more rounds occurred. All Series 2 experiments commenced with parameter set 2, with the wheel again efficient and a unique strict Nash equilibrium. If convergence was achieved (the same configuration in three consecutive rounds) then the parameters were changed to set 3 and the experiment continued. If convergence was again achieved then parameter set 4 was adopted. Subjects

[11] The only exception is experiment 010528 that instead involved a fixed 10 rounds. This trial was included in the final analysis as it provided an additional 60 observations (6 agents, 10 rounds) of individual decisions for tests of behavioral strategies. Critically, the inclusion of this experiment does not favorably bias our results towards network convergence as this experiment did not converge to a stationary configuration.

were unaware of the potential change of parameters. The decision to implement parameter changes was made to avoid randomness that might be associated with boredom and also to test the robustness of the model and separate between the equilibrium predictions of Nash stability, efficiency and focalness. The details of both experimental series are detailed in Table 6.

** Table 6 about here.

4 Results

To understand the nature of networks we study how network structures evolve dynamically and also how individuals make decisions within a network environment. As such, we have divided our results into the three sections. Firstly we present results regarding the macro features of network structures and then the second section presents an investigation of the strategies employed by individual agents. We conclude with synthesis results on how individual behavior critically impacts the evolution of dynamic networks.

4.1 Macro: Network Configurations

Table 7 contains a summary of data from all experiments. Eight of the twelve networks converged to Nash equilibrium configurations (two of five from Series One and six of seven from Series Two). The convergent state was achieved as early as round 4 and as late as round 17. All convergent states were Nash equilibria of the one-shot game, though not always strict Nash. The remaining four experiments did not converge to any stable configurations, Nash or otherwise but three of these experiments temporarily achieved Nash configurations (either weak or strict) that did not prove stable. At no point in any experiment was the empty network chosen.

** Table 7 about here.

The first result to be drawn from the data is that networks can occur and evolve. It is significant in that it is not a negative result as this would leave us unable to reject the hypothesis that social and economic networks are mere historical accidents. It indicates that the agents can appreciate the unique characteristics of networks as opposed to reflecting arbitrary or random choices.

Result 1 Networks happen. Not only are links formed but an appreciation of the externalities inherent in networks is incorporated into agent decisions.

Support: In each experiment a network instantaneously formed. At no point was the empty network chosen. Significantly, at no point was the complete, point-to-point network chosen (i.e., everyone connects to everyone else).

This basic evidence suggests that given the appropriate conditions a social or economic network will emerge. The simple observation provides initial confirmation that networks can arise by economic forces. In the remainder of the paper we attempt to understand the nature of these economic forces.

In a sense Result 2 is central by establishing two important facts. First, the process of network formation tends to stop - a type of equilibration. Secondly, the final configuration tends to be at a Nash equilibrium. Thus, there is a convergence process and the forces at work in the process are captured by game theory in general and the Nash equilibrium in particular. Network formation is not simply a random process.

Result 2 (a) Networks tend to converge and (b) the convergent state is predicted by Nash equilibrium and (c) a greater tendency toward convergence is exhibited by institutions that allow continuous adjustment (Series 2).

Support: (a) See Table 7. Eight of the twelve networks converged to Nash equilibrium configurations (two of five from Series One and six of seven from Series Two). After convergence, the parameters were changed in three of the Series Two experiments and convergence to different networks was achieved in all three. The convergent state was first achieved in rounds 9 and 11 of the Series One experiments, and in rounds 17, 16, 5, 7, 4, and 15 of the Series Two experiments.

With six agents there are $(2^5)^6 = 1,073,741,824$ possible networks. The probability of convergence with random selection in an n round experiment (the same network in three consecutive periods) is then strictly less than $\frac{n-2}{(2^{30})^2}$.[12] Therefore, the hypothesis that network dynamics are ran-

[12] This simple expression is the probability that any three consecutive networks are identical in n periods. It is used here for analytical simplicity. The exact probability that the experiment ceases because of convergence is strictly less than this.

dom can be rejected with an extremely high level of confidence.[13]

(b) All eight convergent networks (and the three re-convergent networks after parameter changes) are Nash equilibria of the one shot game. In no experiments did a network exhibit stability at non-Nash equilibrium configurations.

(c) One institutional adaptation, continuous decision making, was employed in the second series of the experiments presented here. Roughly speaking, this alteration seemed to aid convergence (convergence in six out of seven experiments converged, versus 2 out of 5 that converged for discrete decision making institutions).

While convergence is not guaranteed, the predictability of convergence, in addition to the convergence itself, should be interpreted as strong evidence that something systematic is driving network dynamics. Clearly the Nash equilibrium is a useful concept for capturing what is observed and that fact suggests questions about other features of the model and other principles that might be used in conjunction or as substitute principles for modeling and understanding the process. Three concepts surface immediately: efficiency, strict Nash and focalness of the network. The next result asks if "efficiency seeking" alone, which is closely related to the Nash equilibria, could be driving the results to Nash. By looking at the non convergent examples and asking if they are efficiency improving even if they do not converge to some stable configuration, the question is answered negatively.

Result 3 Non-convergent networks do not exhibit increasing efficiency.

Support: See Figures 6 a, b, and c, and Figure 7b. These graphs represent measures of network efficiency throughout experiments 981106, 990115, 990128, and 01067a, respectively (the non-convergent networks). The measure of efficiency in a network is the amount of information earned per link paid for in the network as a whole. So if the network is at an efficient wheel configuration the measure of efficiency is 6: each agent receives every available piece of information at the cost of only one link.

[13] It should be noted that the claim that network dynamics are not random is quite robust. Even if we restrict agents to choose only one link at a time (what they would need to choose in the efficient Nash network) then randomness can still be rejected at a high level of significance. In this case there are $5^6 = 15625$ possible networks. Thus the probability of convergence with random selection in an n period trial is strictly less than $\frac{n-2}{15625^2}$.

> The slope parameters of these graphs were estimated using ordinary least squares and the t-statistics of these estimates are, respectively, 0.47, 1.74, 1.33, and 1.08. So for all four experiments we fail to reject, even at the 10% level, the null hypothesis that efficiency is not increasing.

** Figures 6 and 7 about here.

Result 3 makes an important point. It tells us that networks are not efficiency seeking phenomena. Thus, our understanding of the convergence process must look beyond the simple property of efficiency or inefficiency to understand how networks will evolve. To do this we turn to a study of the networks that resulted in convergence and the concept of **strict Nash equilibria**. In the many applications of game theory it is well known that the concept of Nash equilibrium is a somewhat weak condition. There exist many suggested refinements of this concept, introduced to make the equilibrium prediction more precise and, if possible, unique. These same concerns apply to the study of networks also, and in a strong way. Bala and Goyal (2000, p.1194) calculate that there exist in excess of 20,000 Nash networks for the environment studied in this paper. They suggest the refinement of **strict Nash equilibrium** and show that this reduces the equilibrium set to a unique architecture (the wheel) which has 120 possible configurations. Our next result produces evidence that this refinement, at least in its pure form, is not entirely appropriate for the study of real networks.

Result 4 Stable configurations may not be in the set of **strict Nash equilibria**.

Support: In experiment 010607b the network converged to a weak Nash equilibrium configuration. This convergent network is depicted in Figure 8. In this network agent 5 is indifferent between connecting a single link from agents 1, 2, 3 or 6, and agent 3 is indifferent between connecting a single link from agents 4, 5, or 6.

** Figure 8 about here.

Though this evidence consists of only a single counterexample, it indicates that the refinement to strict Nash equilibria in network theories is premature.

It also illustrates possible delicate and difficult coordination issues that exist in network formation dynamics as well as possible fragility in networks. Such observation suggests that we look further at the dynamics to determine the robustness of network configurations that conform to various solution concepts.

Further complicating the question of an appropriate model, especially one with roots in equilibrium selection concepts, is that Nash configurations did not always prove stable in the different experiments. Network formation can "pass through" a Nash equilibrium. We observe that Nash configurations, even strict ones, are not always stable. That is, the system need not stop evolving if it happens to achieve the configuration of a Nash Equilibrium.

Result 5 Nash configurations, even strict Nash configurations, are not necessarily stable.

Support: Weak Nash configurations that did not prove stable were played in experiments 981106 (two weak Nash), 010528, 010607a, 010607b, 010613a (three weak Nash), and 010614b. Further, strict Nash configurations (the wheel) were played in experiments 990115, 010607a (in rounds 4, 11-12, and 15), and 010614b (in rounds 9-10, and 12-13). In experiment 010614b the same strict Nash configuration played in rounds 9-10 and 12-13 ultimately proved stable in rounds 15-17.

In view of Result 2, these deviations, particularly from the strict Nash configurations, are surprising and naturally lead to speculation and conjectures about how the model might be modified to account for the phenomena. The most obvious candidates are that these deviations resulted from mistakes, boredom, or confusion. However, this would not seem to be the complete story for the following reasons. Firstly, all participants successfully completed the example calculations in the instructions. Secondly, at least in Series 2 the participants had the opportunity to rectify any mistakes.[14] And, thirdly, no participants indicated any of these three factors in their comments at the end of the experiments.[15] Deeper speculations lead to the idea of common knowledge upon

[14] Assuming they were not making their choices at the last second. This was only the case for one agent in experiment 010614b and this agent, in fact, was not the one to deviate from the Nash equilibrium.

[15] These comments were only elicited after the experiments of Series Two. As a result of the deviations from Nash configurations in the experiments of Series One we began asking participants to describe, at the completion of the experiment, the strategy they employed as well as how they thought their fellow participants were behaving.

which the notion of equilibrium is built. With respect to equilibrium this concept says that every agent knows that every agent is maximizing, and that every agent knows every agent knows every agent knows, and so on. Consequently, it is possible that a group of agents is not in a stable network even though the focal, efficient and strict Nash wheel configuration is being played. Evidence, though weak, can be found for this in experiment 010614b. In this experiment, agent 3 waited until the very last moment before making his decision in each round of play. He persisted with this strategy even when the counter-clockwise wheel was played in rounds 9-10 and all other agents were making their decisions relatively early. This behavior may suggest that agent 3 was not completely aware of the strategic situation faced by himself and his fellow participants, and may have been a factor in agent 5 deviating in round 11. Of further interest is that, if this was in fact how agent 3 was playing, then efficiency and coordination were still achieved with individually optimizing behavior. This possibility is consistent with the intuition behind Bala and Goyal's (2000) main result, that convergence can be achieved despite the presence of self-interested myopic agents.

The possibility of errors and the possible role of common knowledge of rationality lead us once again to the concept of efficiency and to the concept of focalness. Randomness that is biased toward efficiency could enhance the stability of a Nash equilibrium once attained and focalness could help agents understand what other agents might be attempting to do. Thus, these concepts might help to identify equilibrium facilitating features of networks if not prove to be independent principles themselves. The following result demonstrates that their role must be secondary if there is any role at all.

Result 6 Neither focalness or efficiency is the primary determinant of stable configurations.

Support: Experiment 010607b converged to a non-focal and inefficient configuration in rounds 16-18. Experiment 010528 converged to a non-focal wheel in rounds 17-19. Further, after the parameter changes in experiments 010613a, 010613b, and 010614a, the networks diverged from the focal wheel (that was no longer efficient) and reconverged to non-focal wheels.

Combined with Result 2, this result indicates that Nash equilibrium is the guiding principle of network dynamics and convergence, and that focalness and efficiency aren't. This result confirms, if nothing else, that networks are a real economic phenomenon, and should be looked at from an economic perspective. Networks exhibit the classic economic tension between individual incentives and inefficient outcomes. Significantly, this inefficiency is not only transitory but can, in fact, become institutionalized.

It should be noted that these results do not imply that efficiency and focalness should be completely disregarded when analyzing networks. These results merely claim that they are not the primary driving principle. It is still possible that efficiency and focalness play an important role in determining network evolution in situations where stability does not provide a definitive prediction. Indeed, that six out of eight convergent networks were to the focal wheel is indication that this is most likely the case.

4.2 Micro: Individual Decision Making

In our attempt to understand the evolution of networks we turn now to individual behavior. The complexity of networks and the relatively few observations we obtain for each agent make the job difficult but we are able to construct significant tests of behavioral rules from the theoretical literature, which we reject, as well as tests of the conjecture described in section 2.3, which is supported by the data. The conjecture is that in a dynamic environment agents will use link choice to signal, and teach, other agents. In the following section we attempt to piece together behavioral findings with the dynamics of Section 4.1 to further understand the evolution of networks.

In this section we restrict attention to the decisions of individuals in the experiments of Series One. This is done for several reasons. Firstly, this series most closely resembles the theoretical model of Bala and Goyal (2000) and, therefore, provides the more appropriate test of their theory of individual behavior. Secondly, it provides, in a sense, cleaner data. In Series Two agents made decisions continuously and it is difficult to infer the information available to, or the intentions, of each agent at the time the decision was made. The final reason is simply the sheer number of choices available to each agent. Fortunately, the focalness of the clockwise and counter-clockwise wheels in Series One (that are dampened by design in Series Two) permit a more powerful test of the Simple Strategic behavior conjecture.

The first test is of the best response decision rule employed by Bala and Goyal (2000). In the previous section it was documented that networks are indeed capable of converging to the efficient and Nash equilibrium wheel network.

This systemic behavior is predicted by the Bala and Goyal model so it is only natural is ask if the micro behavior required by their model is supported by the data as well. The following result reports that the best response decision rule is not well supported by the data from individual agents decisions. This result leads to a paradox frequently observed in economic experiments that the models work well when applied at the systemic level but the exact behavior of the agents is at odds with the behavioral principles at the foundation of the model.

Result 7 Agents do not act in accordance with the best response decision rule.

Support: As a deterministic decision rule best response is rejected immediately as no agent followed its requirements every round, even allowing for any degree of inertia. As such, no agent can be said to follow this rule. To test for the possibility that agents follow this rule but are prone to error we construct tests in accordance with the techniques introduced by El-Gamal and Grether (1995). The possibility for error allows all possible observations to have positive probability. For a given error level we calculate the likelihood that best response with the given error level could produce the observed sample. If this likelihood is too small then we reject the hypothesis that this decision rule with error is used by the agent. Table 8 presents the findings for the thirty agents observed in Series One of the experiments. As inertia was originally included to model non-optimal behavior it has been included as an error in this analysis.

** Table 8 about here.

As can be seen from the table, even with a 5% possibility of error the hypothesis that best response is being employed can be rejected at the 1% level for all 30 agents. Even with a 25% chance of error the hypothesis can be rejected at the 10% level for 28 out of 30 agents. Indeed, the lack of rejection for 14 agents at the 1% level given a 25% chance of error is more reflective of the lack of power of the tests than confirmation of the decision rule.

If we aggregate the data we find an even more overwhelming rejection of the best response decision rule. We find that even with a 50% chance of error the hypothesis that agents employ the best response decision rule with error can be rejected at the 1% level of significance. Such an aggregated test is appropriate if, as mentioned previously, agents are assumed to be homogenous. Therefore, either no agent uses the rule or the assumption of homogeneity across agents is inappropriate.

This negative result leads naturally to models of individual behavior that are somewhat more sophisticated but also a bit more ad hoc. The following result indicates that such strategy and foresight is evident within networks.

Result 8 Many agents exhibit significant *simple strategic behavior*

Support: Despite the fact that a wheel network appeared in only nine out of the 67 rounds played in the five experiments of Series One, five out of the thirty agents exhibited "simple strategic" behavior in every single round. Two of these agents persisted with this strategy despite participating in experiments in which a wheel network never occurred. Three of the five agents in fact selected the same link, the counter-clockwise wheel (SScc), every period of their experiment. The other two changed between the clockwise and the counter-clockwise wheels at some point in the experiment in what appears to be attempts to coordinate with other agents on which of the focal wheels will actually be chosen.

The remaining agents do not act consistently and uniformly in accordance with simple strategic behavior. So, it is a question of frequency and propensity. To test whether these agents employ simple strategic behavior we shall allow for the possibility of error in their decision making and construct Tables 9 and 10 to test whether simple strategic behavior could have produced the observed sample, as we did above for the best response rule.

** Tables 9 and 10 about here.

As can be seen, many agents appear to act in accordance with simple strategic behavior. For example, with only a 10% error level the hypothesis that agents are acting in accordance with SS behavior cannot be rejected for 10 agents, fully one third of the participants, even at the 10% level of significance.

Significantly, stronger evidence in support of simple strategic behavior is obtained when it is compared directly to best response. Not only can it be seen that many more agents act in accordance with simple strategic behavior than best response but, in fact, upon closer inspection many of the agents for whom best response can't be rejected are, in fact, more likely to be using simple strategic behavior. The lack of rejection of best response is simply a consequence of the requirements of best response and simple strategic behavior coinciding in many situations. That is, what appears as best response behavior is in fact simple strategic behavior.

Result 9 More agents employ the simple strategic decision rule than the best response decision rule. Further, agents using the simple strategic rule adhere to it more frequently than agents best responding.

Support: Using again the techniques from El-Gamal and Grether (1995), we can determine by maximum likelihood which decision rule is the most likely to be used by each agent. For a given probability of error we determine which decision rule was the most likely to have produced the data. In a comparison of best response and counter-clockwise simple strategic behavior (SScc) we find that for 17 agents the most likely decision rule is SScc, and for twelve agents best response is the most likely rule. One agent is equally likely to have used either rule. However, the most striking result from this demarcation arises when we consider how many using each decision rule actually use the rule to a significant degree. We find that most of the agents for whom we were unable to reject the best response rule are, in fact, more likely using SScc. Tables 11 and 12 present these findings.

** Tables 11 and 12 about here.

These results indicate that most agents engage in strategic signaling and coordination efforts through their link selections, and do not necessarily maximize their current payoffs. In contrast, relatively few attempt to optimize current payoffs by using the best response decision rule. Further, the results imply that agents behave in some sense strategically and with foresight in network environments. Many agents employ a strategy that seems to be an attempt to teach, signal, and coordinate all agents within a network and in doing so facilitate movement toward Pareto optimality.

Not all agents employ the same strategy and both the tendencies toward a certain type of behavior together with the lack of uniformity have implications for understanding and modeling networks. Models that assume homogenous agents have striking powers of prediction nevertheless are based on a model of individual behavior that is in need of improvement.[16] We find that agents behave in more complex ways than the simple best response rule allows, and ways that might significantly impact the dynamics of the networks. It would appear that agents are making decisions on more analysis than simply their immediate payoff or their own actions.[17]

4.3 Interdependence: Micro and Macro

The findings of the previous section leave two prominent questions. Firstly, how instrumental are the strategies of individuals, particularly those employing simple strategic behavior, in achieving convergence? And, secondly, how do agents behave who neither best respond nor use simple strategic behavior? We produce an answer to the first question, and present some aggregate evidence to provide insight into the second, which will be explored with the final result of this section.

With respect to the first question, the behavior of individual agents appears to be crucial to whether or not convergence to the wheel network is achieved. Specifically, individual agents appear to be capable of influencing the evolution of networks by signaling to their fellow agents an optimal strategy. By playing SS behavior agents can teach other agents the structure of the game and the nature of payoffs. This induces these agents, who begin with some other strategy, to switch to SS behavior and as a consequence the probability of convergence increases.

This individual capability is best seen by considering how convergent and non-convergent networks differ. Surprisingly, they differ by only a small, but significant, amount. For all experimental networks, convergent or not, the majority of agents exhibited simple strategic behavior throughout and several agents did not. In the convergent networks the remaining agents learned the optimal configuration and began coordinating with their fellow agents on counter-clockwise simple strategic behavior. In the non-convergent networks

[16] It is possible that all agents are employing the same strategy and that the heterogeneity is at the level of understanding and/or knowledge. This view is perhaps more consistent with the (frequent) stability of Nash configurations. Either way, however, it would seem heterogeneity is an appropriate assumption for models of network formation.

[17] These conclusions should not be interpreted as conflicting with the results of Bala and Goyal (2000). As mentioned previously, the results of Bala and Goyal should be interpreted as the benchmark capabilities of network evolution.

this learning did not occur. These findings are captured by the following two results.

Result 10 Network convergence is critically dependent on the behavior of all agents. Moreover, agents can learn to choose optimal strategies and enable networks to converge to stable outcomes.

Support: The results are presented in Table 13. All six agents in the networks that converged exhibit counter-clockwise simple strategic (SScc) behavior. By contrast, for the non-convergent networks such consistency of behavior was not observed, though there were at least four agents for whom randomness is rejected in favor of SScc behavior.

** Table 13 about here.

Upon closer inspection, however, it can be seen that all six agents in convergent networks do not exhibit SScc behavior consistently throughout the experiment. To expose this shift in networks that converged, Table 14 performs the same analysis as above but omits the periods after convergence has occurred. It can now be seen that, surprisingly, convergent networks look very similar to the networks that didn't converge.

** Table 14 about here.

Thus, all of the experiments appeared similar up until a critical point in two of the networks at which the remaining agents learned to play the SScc strategy and convergence to efficiency was achieved. In the non-convergent networks this learning simply did not occur, and consequently inefficiency was the result.

The next result shows that the devotion of agents to a particular strategy, in this case SScc, is also critical to network dynamics. Agents in the ultimately stable networks are significantly more committed to the simple strategic strategy, and this dedication appears to be pivotal in achieving convergence.

Result 11 Network convergence depends on the commitment to simple strategic behavior of individual agents. Thus, all agents can impact the probability of convergence.

Support: This result is exposed by the relationship between whether a network converged and the rate at which counter-clockwise simple strategic behavior is played (before convergence was achieved in the convergent networks).[18] Table 15 details these variables.

** Table 15 about here.

Unfortunately, because of the small number of observations, the power of this test is limited. However, despite this constraint, a significant relationship emerges. Estimating the equation $y = \alpha + \beta\left(\frac{\gamma}{n}\right) + \varepsilon$ by OLS we find that the estimate of β is positive and significantly different from zero (β estimate of 0.61, t-statistic of 2.38, and p-value of 0.14; α estimate of -0.97, t-statistic of -1.51, and p-value of 0.26).[19,20]

These results begin to expose the integral role of individual decisions in network formation. Agents can educate their fellow agents to play the optimal strategy but this requires the educator agents to play their component of the focal wheel network, and to play it consistently. Unfortunately the precise nature of the resulting learning cannot be clearly ascertained from this data. The final two agents may learn to play the simple strategic strategy because the repeated play of SScc by the other agents has taught them the common benefits of such play. Alternatively, they may eventually play SScc because it is a best response to the choice of SScc by their fellow agents.

The next result indicates that there is still more to the story of individual behavior and how agents are making decisions if not using SScc. Unfortunately, the possibilities are far too complex for meaningful tests of individual behavior to be conducted on the relatively few observations reported here. As a result, the analysis turns to the aggregate data. In an admittedly crude test, and employing data from both experimental series to increase test power, we are

[18] If instead the number of observations by the leading four agents is used here, as may be interpreted from Result 10, the conclusions are not affected (the number of observations in this instance are 22, 36, 20, 23, and 26, respectively).

[19] For this regression we omitted the rather special case of experiment 990115 that achieved the focal wheel but immediately diverged. Including it in the estimation as a non-converged network (or even a converged network) does not change the results substantially.

[20] As we are interested in establishing the existence of a significant relationship between these two variables, we will not make a further distributional assumption to produce estimates of the probability of convergence (though, of course, if the uniform distribution is assumed the probability estimates are equal to the constrained OLS estimates).

able to infer group responses to network situations which exhibit best response-like behavior.

The logic behind the group analysis, which is reported as Result 12 below, is the following. If in a network the total number of information pieces collected is less than 36 (six pieces per agent) then the network is inefficient. Such a network is inefficient, regardless of the number of links present, because additional links can be added that would add to individual and group profit. Therefore, the best response of at least one individual, assuming all other links remain unchanged, is to add more links.

Group level quasi-best responsiveness We say that there is "group level quasi-best responsiveness" if an inefficiency caused by a shortage of information is followed at the group level by the connection of additional links. Likewise, if all information is collected and there are too many links then this is followed by a reduction in the number of links.

Analysis from aggregate data cannot be used to make inferences about specific individuals but it can be used as a guide to the nature and process of network evolution, and as a indication of the type of model that might be useful for uncovering precise estimates of individual behavioral strategies in network environments. There are many possible explanations for individual behavior that could produce the aggregate effects documented here. Arguably, however, the main lesson suggested is that agents not employing the SScc strategy are still exhibiting some kind of rationality. A detailed exploration of this brand of rationality and its ability to induce efficient outcomes should form the basis of future work.

Result 12 Network dynamics exhibit strong evidence of "group level quasi-best" responsiveness.

Support: Evidence of this group dynamic can best be seen graphically. Figure 9 depicts the relationship between the total number of links selected and the average number of pieces of information received by each agent through the 18 rounds of experiment 010607b.

** Figure 9 about here.

The group level quasi-best responsiveness is evident in the correlation of these two measures. In all but one period (the final period before convergence) an inefficient accumulation of information by the agents (less than 6) is followed by an increase in the number of connected links in the following rounds. Similarly, in all rounds (other than after convergence was achieved) an efficient accumulation of information by the agents (average of 6) is followed by a decrease in the number of connected links in the following rounds.

To test this idea more formally we regress the average information value on the change in total links from period to period. This is written,

$$\Delta_t = \alpha + \beta I_t$$

Where I_t is the average information level in round t, $\Delta_t = L_{t+1} - L_t$, and L_t is the total links chosen in round t. The estimates for this equation for all experiments combined are given in Table 16.[21]

** Table 16 about here.

These findings indicate that the causal relationship between information accumulation and changes in the number of links selected is negative and statistically significant (at the 1% level). The data produce strong support for group level quasi-best responsiveness.

5 Conclusion

This research has attempted to present some key characteristics and principles of network evolution. The principles studied are theoretically general, with potential applications beyond the particular environments studied here. The research made no attempt to present a comprehensive study of networks but hopefully it serves to highlight the importance and uniqueness of networks, and will encourage further investigation of the characteristics and questions raised here.

From a theoretical point of view there are many reasons why network development and evolution might fail. Networks involve many features that make efficient and decentralized development problematic. The decisions facing agents within a network involve components of many well known problems. The flows within networks create externalities and, not surprisingly, free rider issues that

[21] For consistency, rounds in Series Two after convergence (when the parameters were changed) are excluded.

relate to the public goods problem. Similarly the development of these links involve coordination problems and, as they are typically costly, implicit bargaining. Indeed, the asymmetry of payoffs within networks, even efficient and Nash equilibrium configurations, is reminiscent of ultimatum games (all agents force another into a low payoff network).

The data reported here demonstrate that decentralized development of networks can occur in spite of the inherent potential for problems and complete failure (Result 1). Not only do they develop but they often, though not always (Result 4), converge to a stable configuration that has the properties of a Nash equilibrium (Result 2). The dynamics of change do not exhibit properties of a simple model in the sense that efficiency increasing paths do not seem to emerge (Result 3). Further, the paths can move through Nash equilibrium configurations without becoming stable (Result 5). Generally, neither the property of focalness or efficiency seem to be the primary determinants of stable configurations (Result 6). While the pattern of these results is suggested by existing models, model inaccuracy is certainly present. In part the models fail to capture the complexity of individual decisions. Individuals tend not to follow a Nash best response (Result 7). Instead, a more forward looking pattern of decisions exists (Results 8 and 9), possibly reflecting attempts to coordinate or even teach other agents about strategies that facilitate coordination. Typically the failure of a stable network to emerge is accompanied by a few agents who are not following the same strategies as others (Result 10 and Result 11). While it is impossible to characterize or even identify all of the individual decision rules employed, as a group the decisions reflect properties of the SScc (Result 12). Thus, since evidence exists that individuals are forward looking, system efficiency or efficiency improvements may still provide part of the motivation for these dynamics, though not the principle force.

Basic lessons emerge. The first is that the recent models of network emergence have predictive power but the source of that power is not entirely clear. Perturbations of the parameters of otherwise converged and efficient networks suggest that the predictive power of the models might be fragile. The environments studied here would seem to be very supportive of convergence but many parameter changes including asymmetric costs and information, decay, lack of public knowledge of links or moves, multiple directional flows, etc. can produce networks of much greater complexity than we have studied. This possibility suggests a need for two parallel investigations. On the one hand there is a need for analysis of the identification of efficient network configurations, their stability and their evolution under the hypothesis of the Nash dynamic. Existing literature strongly suggests that networks and network connections have a place as an economic form of organization subject to efficiency analysis. On the other hand there is a need for an understanding of the types of institutional arrangements within which network formation might take place and that might facilitate efficient, decentralized network developments and evolution. The results reported here suggest that institutions that facilitate an understanding of the intentions of players, a means for formation of and communication of rationality and a common knowledge of it, could be important. Questions to be answered include the following. What type of decision making and institutional structure might best facilitate efficient network evolution? What types of mechanisms might be best when the economic problem is related to network configuration? Indeed the motivation for Series Two was that a somewhat different information flow and decision making cost would improve the convergence and efficiency of the process. The data suggest that this was successful.

6 Appendix

NETWORK EXPERIMENT INSTRUCTIONS[22]

This is an experiment of network formation. If you follow these instructions and make appropriate decisions, you can earn an appreciable amount of money. At the end of the experiment your earnings will be paid to you privately, and in cash. In this experiment each person holds some private information. This information is valuable to you, and to every other person who can access it. It has a value of 25 cents for every person, including yourself, who holds it. You can access someone else's information directly by forming a link from them. Each link costs 15 cents. You may form as many links as you like. A piece of information can be passed along multiple times.

Say in a three person network that you connect a link from Person 1. Then if Person 1 is also linked to Person 2 you receive the information from both Person 1 and Person 2, but you only have to pay for the link from Person 1. This is shown in the following diagram. Note that the tip of the arrow points to the person who is paying for the link and so receiving the information.

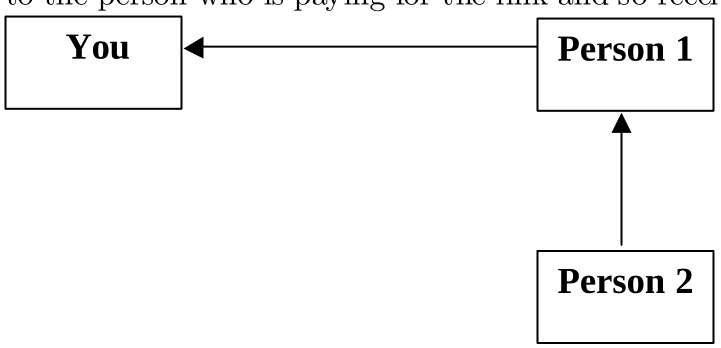

Information you now hold: You, Person 1, and Person 2
Value = 3 x 25 cents
Costly links you have paid for: From Person 1
Cost = 15c
Net profit for you for this round = 75c – 15c = 60 cents

The information links are only one way. So in the above example, if you have paid for the link from Person 1 then you receive Person 1's information, but Person 1 does not receive your information. Person 1 would have to pay

[22]These are the instructions used in Series Two of the experiments. Other than minor changes to allow for the differences between the series, the instructions used in Series One are identical.

for a link from you if he wanted your information. Note that it is permitted for any two people to pay for links from each other simultaneously.

Let's try a simple example. In the following diagram what would be the profits for persons 1,2 and 3? Write your answers in the space provided. Represent links paid for by circling the person number from whom each link is connected. Circle 'N' if a person has not paid for any links. Remember, you also receive value from holding your own information.

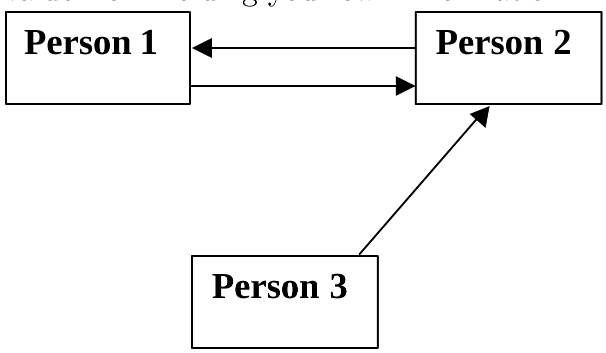

	Links Chosen	Cost	Information Received	Value	Profit
Person 1	N 1 2 3	15c	1,2,3	3x25c	60c
Person 2	N 1 2 3				
Person 3	N 1 2 3				

The experiment will involve multiple rounds. Each round will close after two minutes. In each round you will select from which people you wish to pay for a link. You will mark your selection(s) in the box marked 'Link Submission Form' and submit by clicking on 'submit links.' You may change your selections as often as you like during each round. However, you will be charged an adjustment fee of 5c every time you add or subtract a link and click on 'submit links.' (5c for each addition or subtraction) You can record these charges on your Record Sheet in the column marked 'Cost of Changes.'

The box marked 'Connections' represents the links currently selected by each person (read horizontally), with your expected payoff below. This record will be updated continuously as people connect and disconnect links during each round.

The 'Link/Total Value' box describes the value you will receive by connecting a single link from each other person. At the end of each round you will pay the 15c connection fee for each currently selected link. These connections will then be used to calculate the information you accumulate and the earnings you receive. Note that you will be paid for your performance in each round of play.

This process is then repeated in each subsequent round. In each round connections start anew, so you will pay for any links you hold in that round, regardless of whether you have held that link previously. In each round you may connect any link, or combination of links, that you desire.

The exact number of rounds to be conducted will be determined randomly. We will conduct at least fifteen rounds. At the end of the fifteenth round and after every subsequent round a pair of dice will be rolled. If the sum of the roll exceeds a certain number, specified in the table below, then the experiment will stop. Otherwise we will continue with another round and repeat the process. You will notice that the probability of stopping after a given round increases as we play more rounds.

Round	End if Sum $\geq$	Prob End
15	12	1/36
16	10	6/36
17	8	15/36
18	6	26/36
19	4	33/36
20	2	1

After the experiment is completed you will be paid your profits. Are there any questions before we begin? Please do not talk or communicate with anyone else during the experiment. We will insist that everyone remain silent until the end of the last period. If we observe you communicating with anyone, other than the experimenter, we will ask you to leave without completing the experiment. We are now ready to begin round one. Please choose your desired connections for round one on your screens.

References

[1] Bala, V. and S. Goyal (2000). "A Noncooperative Model of Network Formation." *Econometrica* 68 (September): 1181-1230.

[2] Bala, V. and S. Goyal (2000a). "A Strategic Analysis of Network Reliability." *Review of Economic Design* 5: 205-228.

[3] Camerer, C.F, T-H. Ho, and J-K. Chong (2002). "Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games." *Journal of Economic Theory* 104 (May): 137-188.

[4] Dutta, B. and S. Mutuswami (1997). "Stable Networks." *Journal of Economic Theory* 76: 322-344.

[5] El-Gamal, M. and D. Grether (1995). "Are People Bayesian? Uncovering Behavioral Strategies." *Journal of the American Statistical Association* 90 (December): 1137-45.

[6] Hendricks, K., M. Piccione and G. Tan (1999). "Equilibria in Networks." *Econometrica* 67 (November): 1407-1435.

[7] Jackson, M. and A. Watts (2002). "The Evolution of Social and Economic Networks." *Journal of Economic Theory* 106 (October): 265-295.

[8] Jackson, M. and A. Wolinsky (1996). "A Strategic Model of Economic and Social Networks." *Journal of Economic Theory* 71: 44-74.

[9] Kranton, R. and D. Minehart (2001). "A Theory of Buyer-Seller Networks." *The American Economic Review* 91 (June): 485-508.

[10] Ledyard, J. (1995). "Public Goods: A Survey of Experimental Research." In J. Kagel and A. Roth (ed.), *The Handbook of Experimental Economics*. Princeton University Press (Princeton, New Jersey).

[11] Schelling, T.C. (1960). *The Strategy of Conflict*. Harvard University Press (Cambridge, Massachusetts).

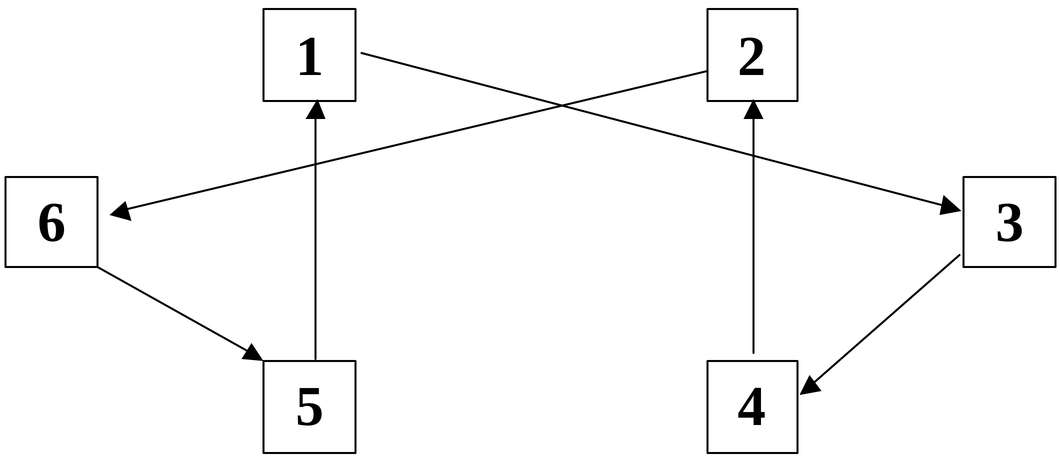

Figure 1: A Six Agent Network

Parameter Set Number	**Link Connection Cost**	**Info Value per unit**
1	\$0.15	\$0.20
2	\$0.15	\$0.25
3	\$0.30 from neighbors, \$0.15 from others	\$0.25
4	free connection in/out 1, \$0.15 from others	\$0.25

Table 1: Parameter Sets

Parameters	**Strict Nash**	**Weak Nash?**	**Focal**	**Efficient**
Set 1	wheel	many (e.g., star)	(counter-)clockwise wheel	wheel
Set 2	wheel	many (e.g., star)	(counter-)clockwise wheel	wheel
Set 3	wheel	many (e.g., star)	(counter-)clockwise wheel	non-focal wheel
Set 4	wheel/star	many	(counter-)clockwise wheel	star centred on 1

Table 2: Model Predictions

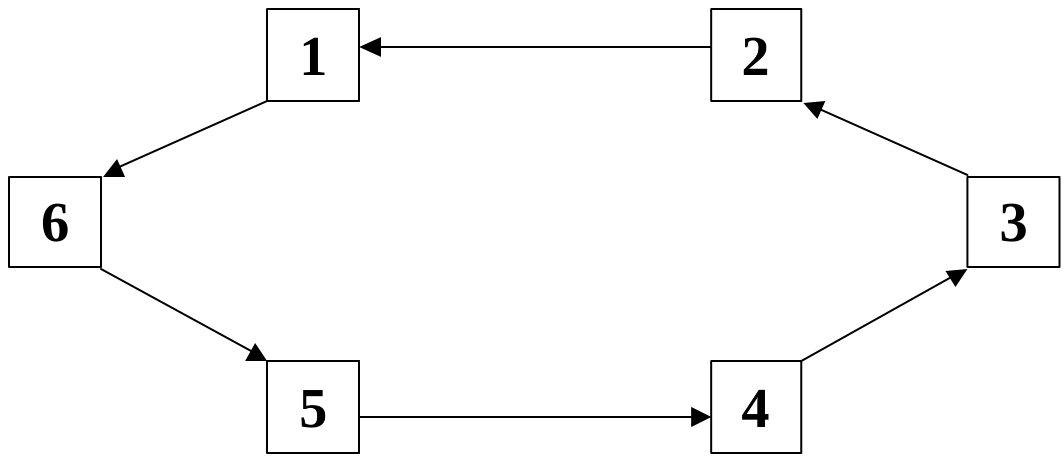

Figure 2: The Counter-Clockwise Wheel

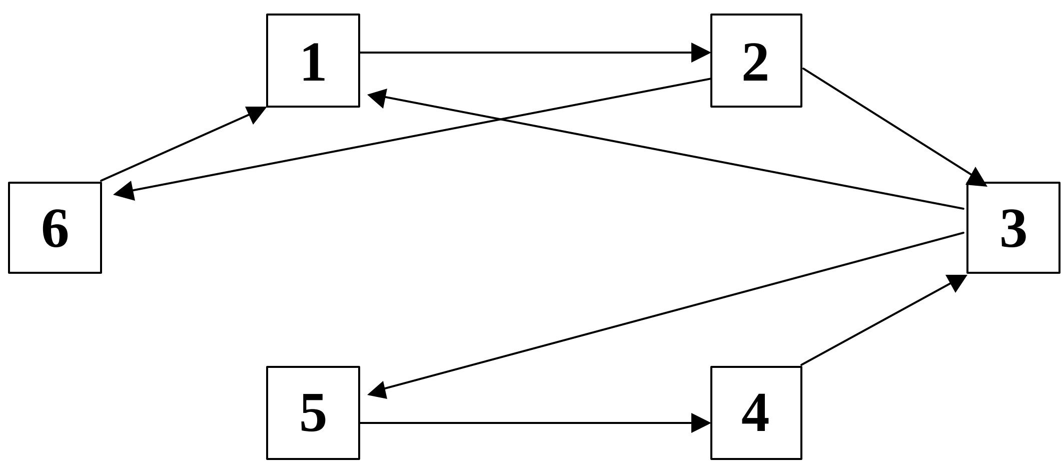

Figure 3: A Weak Nash Configuration with 8 links

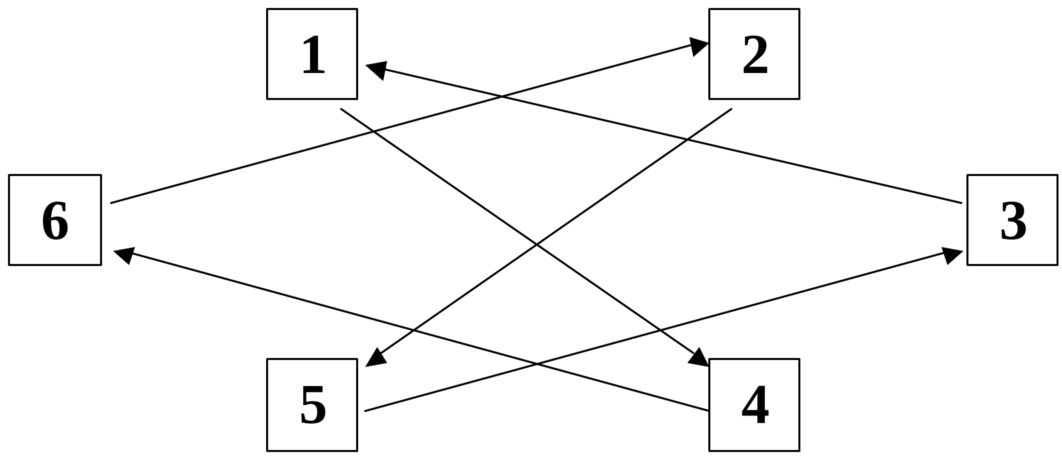

Figure 4: An Efficient Non-Focal wheel (parameter set #3)

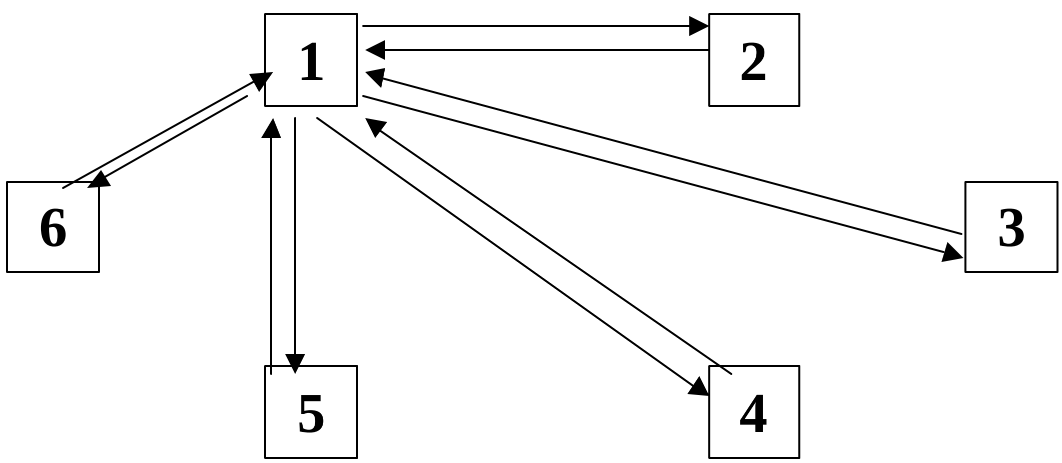

Figure 5: An Efficient Star (parameter set #4)

Property	**Value of Property**
· number of agents	6
· flow quality	no decay
· flow direction	one way
· actors	individuals at nodes

Table 3: Experimental Design: Common Features Series One and Series Two

Series 1	Series 2
· 5 experiments	· 7 experiments
· conducted manually	· conducted over computer
· simultaneous moves	· continuous opportunity to move
· no cost of adjustment	· adjustment cost imposed
· random stopping between 10-20 rounds	· random stopping between 15-20 rounds

Table 4: Experimental Design: Specific Features Series One and Series Two

Stop if dice roll $\geq$					
Round	**Rule 1**	**Rule 2**	**Round**	**Rule 1**	**Rule 2**
10	12	-	15	7	12
11	11	-	16	6	10
12	10	-	17	5	8
13	9	-	18	4	6
14	8	-	19	3	4
15	7	12	20	2	2

Table 5: Stopping Rules

	Series 1			**Series 2**	
Experiment	**Parameters**	**Rounds**	**Experiment**	**Parameters**	**Rounds**
981106	Set 1	1-10	010528	Set 2	1-19
990115	Set 1	1-15	010607a	Set 2	1-17
990128	Set 1	1-16	010607b	Set 2	1-18
990212a	Set 1	1-13	010613a	Set 2	1-7
990212b	Set 1	1-13		Set 3	8-12
				Set 4	13-16
			010613b	Set 2	1-9
				Set 3	10-16
			010614a	Set 2	1-6
				Set 3	7-17
			010614b	Set 2	1-17

Table 6: Experimental Design: Parameters

Series 1		
Experiment	**Rounds**	**Result**
981106	10	No convergence
990115	15	No convergence
990128	16	No convergence
990212a	13	Converged to focal wheel in rounds 9-13
990212b	13	Converged to focal wheel in rounds 11-13
Series 2		
010528	19	Converged to non-focal wheel in rounds 17-19
010607a	17	No convergence
010607b	18	Converged to inefficient weak Nash in rounds 16-18
010613a	16	Converged to focal wheel in rounds 5-7
		Converged to efficient non-focal wheel in rounds 10-12
		No convergence in rounds 13-16
010613b	16	Converged to focal wheel in rounds 7-9
		Converged to efficient non-focal wheel in rounds 14-16
010614a	17	Converged to focal wheel in rounds 4-6
		Converged to efficient non-focal wheel in rounds 15-17
010614b	17	Converged to focal wheel in rounds 15-17

Table 7: Summary Data: All Experiments

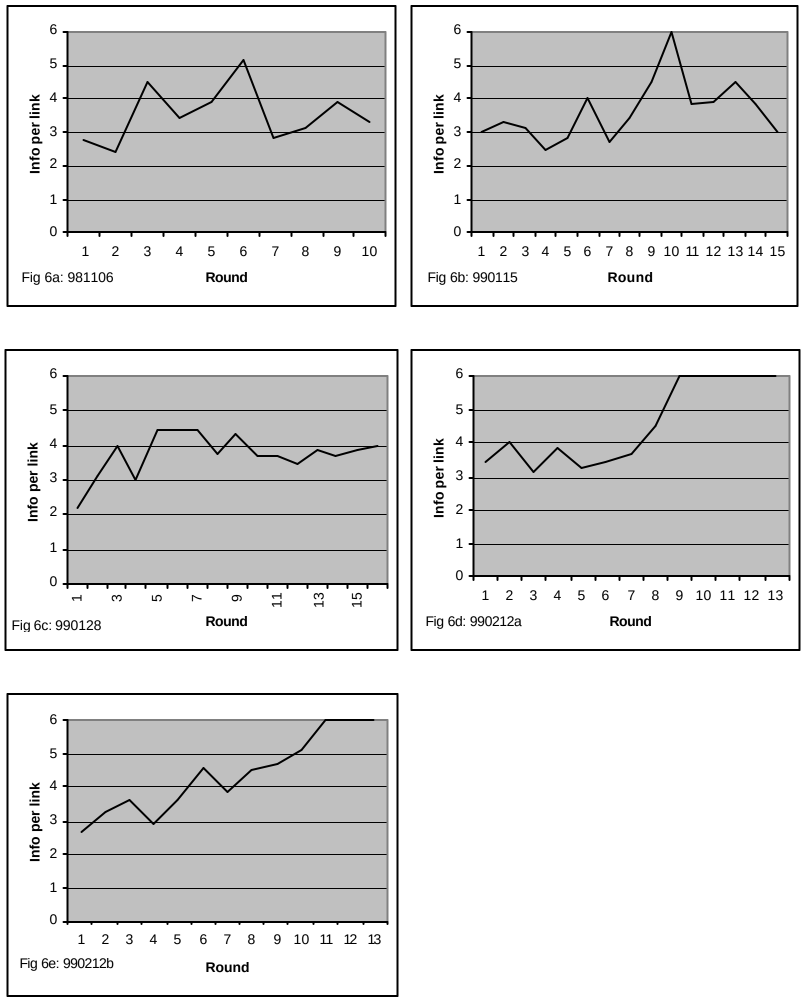

Figure 6: Series One Network Efficiency

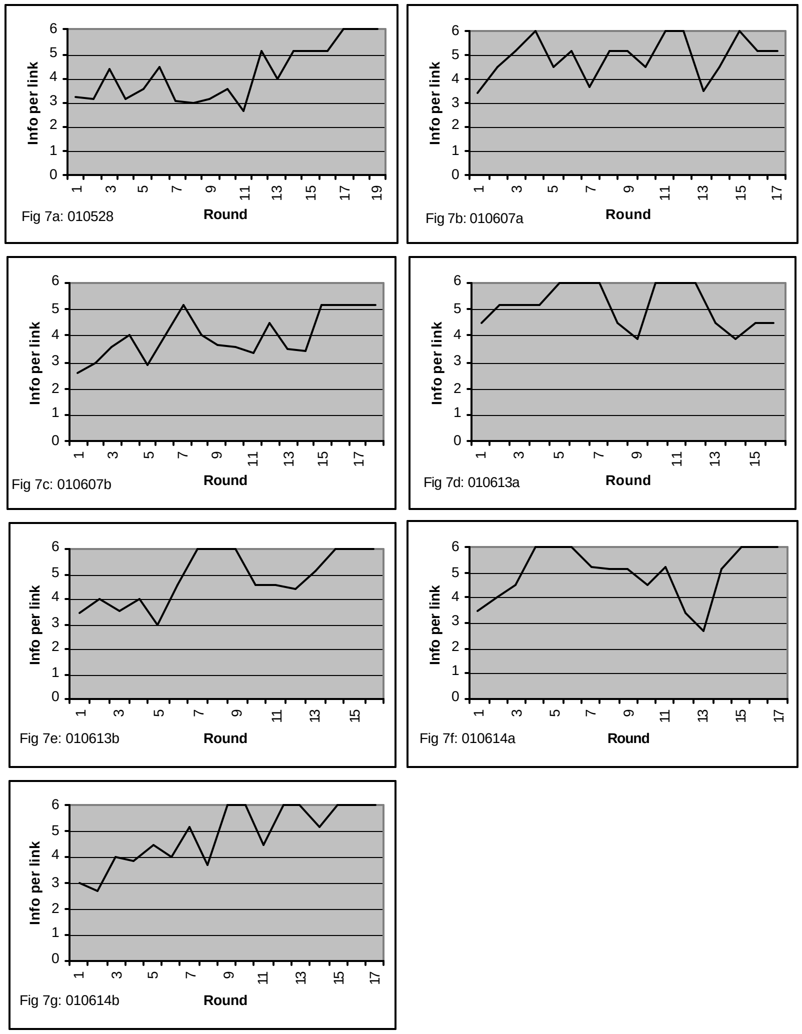

Figure 7: Series Two Network Efficiency

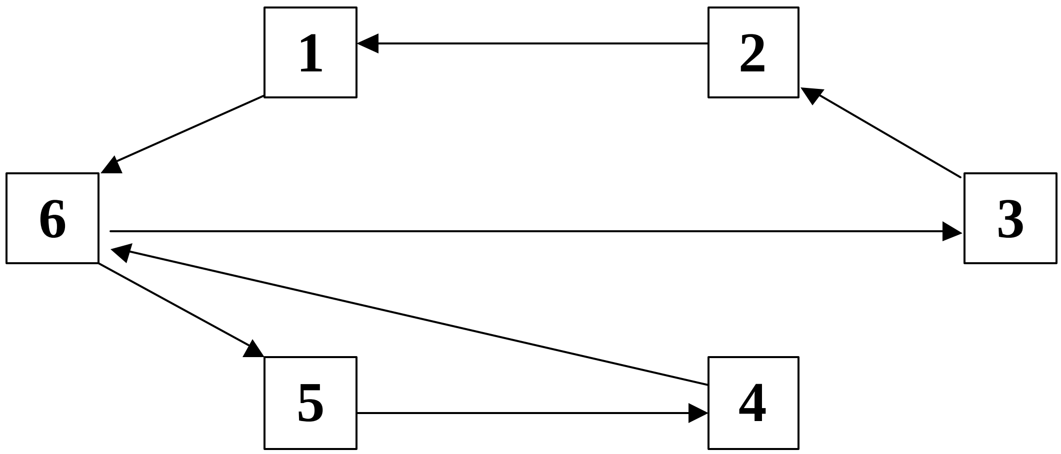

Fig 8: Stationary Weak Nash Configuration in Experiment 010607b

	Error Level				
Level of Significance	1%	2%	5%	10%	25%
1%	0	0	0	1	14
2%	0	0	0	1	12
5%	0	0	0	0	8
10%	0	0	0	0	2

Table 8: Number of Agents for Whom the Best Response with Error Decision Rule Cannot be Rejected

	Error Level				
Level of Significance	1%	2%	5%	10%	25%
1%	6	7	8	12	20
2%	6	6	8	12	18
5%	6	6	7	10	16
10%	6	6	6	10	15

Table 9: Number of Failures to Reject Simply Strategic Behavior

	Error Level				
Level of Significance	1%	2%	5%	10%	25%
1%	3	5	5	10	17
2%	3	5	5	9	15
5%	3	3	5	5	13
10%	3	3	5	5	11

Table 10: Number of Failures to Reject Counter-Clockwise Simple Strategic Behavior (SScc)

	Error Level				
Level of Significance	1%	2%	5%	10%	25%
1%	3	5	5	10	14
2%	3	5	5	9	12
5%	3	3	5	5	12
10%	3	3	5	5	11

Table 11: Number of Failures to Reject Counter-Clockwise Simple Strategic Behavior (SScc): $17\frac{1}{2}$ Agents

	Error Level				
Level of Significance	1%	2%	5%	10%	25%
1%	0	0	0	1	5
2%	0	0	0	1	4
5%	0	0	0	0	1
10%	0	0	0	0	1

Table 12: Number of Failures to Reject Best Response: $12\frac{1}{2}$ Agents

	Experiment[23]				
Level of Significance	981106	990115	990128	990212a	990212b
1%	4	4	3	6	6
2%	4	5	4	6	6
5%	4	5	4	6	6
10%	4	6	4	6	6

Table 13: Agents For Whom Randomness is Rejected in Favor of Counter-Clockwise Simple Strategic Behavior (SScc)

	Experiment	
Level of Significance	990212a	990212b
1%	4	4
2%	4	4
5%	4	5
10%	4	5

Table 14: Agents for Whom Randomness is Rejected in Favor of Counter-Clockwise Simple Strategic Behavior Before Convergence Occurs

	Experiment				
Before Convergence:	981106	990115	990128	990212a	990212b
Observations of ssa(γ)	23	41	21	25	29
Rounds (n)	10	15	16	8	10
Rate of ssa $\left(\frac{\gamma}{n}\right)$	2.3	2.733	1.313	3.125	2.9
Converged (y)	No	No	No	Yes	Yes

Table 15: Observations of Counter-Clockwise Simple Strategic Behavior

[23] The hypothesis tested here is that agents choose among all strategies (including SScc) randomly. The results, therefore, reject the hypothesis for a majority of agents that a random selection generated the observed sample of frequent SScc selection. Thus, this leads to the conclusion that agents chose according to the SScc strategy with greater than random probability.

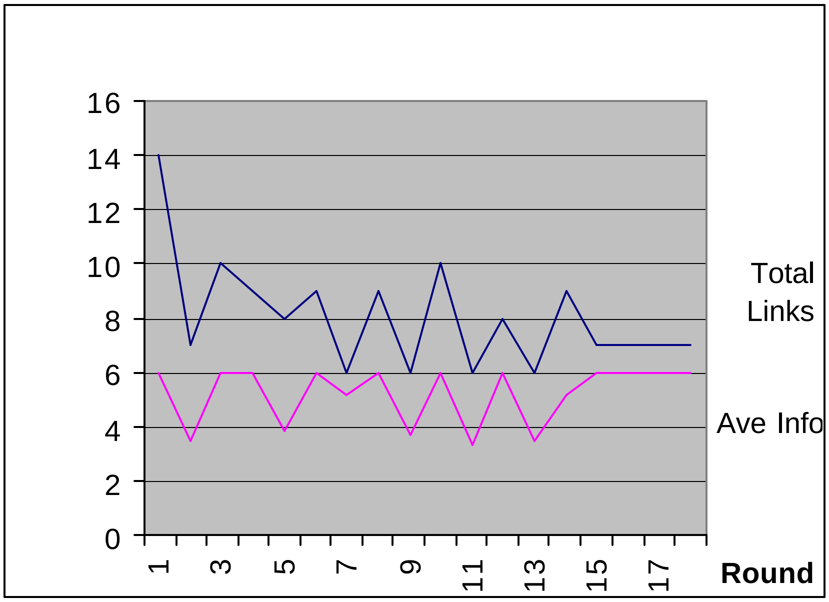

Figure 9: Total Links and Average Information Received

Variable	Estimate	t-statistic
α	4.86	7.66
β	-1.00	-8.29
$R^2 = 0.32$	$n = 148$	

Table 16: "Group Level Quasi-Best Responsiveness"