Network Architecture, Salience and Coordination*

Syngjoo Choi† Douglas Gale‡ Shachar Kariv§
UCL NYU UC Berkeley

Thomas Palfrey¶
Caltech

November 19, 2008

Abstract

This paper reports the results of an experimental investigation of monotone games with imperfect information. Players are located at the nodes of a network and observe the actions of other players only if they are connected by the network. These games have many

*This research was supported by the Princeton Laboratory for Experimental Social Science (PLESS). The paper has benefited from suggestions by the participants of seminars at several universities. We acknowledge The National Science Foundation for support under grants SBR-0095109 (Gale), SES-0617955 (Gale and Kariv), and SES-0617820 (Palfrey) and The Gordon and Betty Moore Foundation (Palfrey). Kariv is grateful for the hospitality of the School of Social Science in the Institute for Advanced Studies.

†Department of Economics, University College London, Gower Street, London WC1E 6BT, UK (Email: syngjoo.choi@ucl.ac.uk, URL: http://www.homepages.ucl.ac.uk/~uctpsc0).

‡Department of Economics, New York University, 19 W. 4th Street, New York, NY, 10012, USA (E-mail: douglas.gale@nyu.edu, URL: http://www.econ.nyu.edu/user/galed).

§Department of Economics, University of California, Berkeley, 508-1 Evans Hall # 3880, Berkeley, CA 94720, USA (E-mail: kariv@berkeley.edu, URL: http://econ.berkeley.edu/~kariv/).

¶Division of the Humanities and Social Sciences, California Institute of Technology, MC 228-77, Pasadena, CA 91125, USA (E-mail: trp@hss.caltech.edu, URL: http://www.hss.caltech.edu/~trp/).

sequential equilibria; nonetheless, the behavior of subjects in the laboratory is predictable. The network architecture makes some strategies *salient* and this in turn makes the subjects' behavior predictable and facilitates coordination on efficient outcomes. In some cases, modal behavior corresponds to equilibrium strategies.
JEL Classification Numbers: *D82, D83, C92.*
Key Words: *experiment, monotone games, imperfect information, networks, coordination, strategic commitment, strategic delay, equilibrium selection, salience.*

1 Introduction

A perennial question in economics concerns the conditions under which individuals cooperate and coordinate to achieve an efficient outcome. In a series of papers, Gale (1995, 2001) showed that, under certain conditions, cooperation arises naturally in the class of *monotone games*. A monotone game is like a repeated game except that actions are irreversible: players are constrained to choose stage-game strategies that are non-decreasing over time. This irreversibility structure allows players to make commitments. Every time a player makes a commitment, it changes the structure of the game and the incentives for other players to cooperate.

Choi et al. (2008), henceforth CGK, conduct a theoretical and experimental study of a class of simple monotone games that are naturally interpreted as step-level, or threshold, public good games. Each player has an endowment of tokens E that he can either keep for himself or contribute toward the cost of an indivisible public good. The good costs K tokens to complete. The players make irreversible contributions to the public good at a sequence of dates. At the end of T periods, the public good is provided if and only if the sum of the contributions is large enough to meet the cost of the good. Each player assigns the value V to the good, so his utility if the good is provided is equal to V *plus* his endowment *minus* his contribution. If the good is not provided, his payoff equals his endowment *minus* his contribution.

The main theoretical result in CGK is that, if the length of the game T is greater than the cost of the good K (and certain side constraints are satisfied), then players must cooperate and provide the good with positive probability (probability one in a pure-strategy equilibrium). Thus, the timing and irreversibility of decisions can help avoid uncooperative equilibria.

However, the coordination problem remains a potential obstacle to efficiency.

A central assumption in CGK is that information is *perfect*: every player is assumed to be informed about the entire history of actions that have already been taken. Perfect information and sequential choice can make it easier for players to coordinate their actions if they are so inclined. Clearly, *imperfect* information can be an obstacle to cooperation. In the extreme case where no player has any information about the other players' prior actions, the situation is essentially the same as in the one-shot game, where players contribute simultaneously to the provision of the public good. For intermediate cases, with partial information, the central result of CGK continues to hold: under weak conditions, sequential rationality implies provision of the public good with positive probability. Our motivation for the present study is to determine the effect of the information structure in this intermediate range – between the cases of zero and perfect information – on coordination, dynamics, and efficiency.

The imperfect information structure is represented by a directed graph that specifies the information flows in the group, i.e., who observes whose past actions. Each player is located at a node of the graph. Player i can observe player j if and only if there is an edge leading from node i to node j. The network architecture is common knowledge to the players. The experiments reported here involve all three-person networks with zero, one or two edges. We call the unique 0-edge network the empty network and the unique 1-edge network the one-link network. There are four 2-edge networks, called the line, the star-in, the star-out, and the pair network. The complete set of networks is illustrated in Figure 1, where an arrow pointing from player i to player j indicates that player i can observe player j. The set of networks illustrated in Figure 1 is essentially complete in the sense that any other network with less than three edges is simply a re-labeling of these networks.

[Figure 1 here]

Each of the networks obtained by adding edges has a variety of sequential equilibria, but the individual behaviors in these equilibria have their "counterpart" in one or more of the networks used in the experimental design. It is therefore adequate to restrict attention to the networks depicted in Figure 1 in order to focus on identifying which behaviors – and possibly equilibria – are plausible or salient. Nevertheless, this set of networks has several non-trivial architectures, each of which gives rise to its own distinctive information flows. Because of the imperfection of information and the lack of common knowledge about the number of past contributions, the theory suggests that even in three-person networks the process of coordinating on an efficient outcome can be complicated. The theory also involves a number of crucial elements, such as Bayesian updating and sequential rationality, which raise questions about how subjects will behave even in small experimental networks and confirm the importance of verifying the relevance of the theory empirically. For practical purposes, the networks with zero, one or two edges "span" the set of networks, reveal the important features of the game, and provide a reasonable test of the theory.

The games that make up the various treatments in our experiments differ only with respect to their network architecture. The other parameters are the same for all treatments. There are three players, each of whom has an endowment consisting of a single token ($E = 1$). The cost of the public good is two tokens ($K = 2$), so that some coordination is required to provide the public good and there is an opportunity for one player to free-ride. There are three periods in the game ($T = 3$), so the crucial inequality identified by CGK is satisfied. The value of the public good is two tokens ($V = 2$), so it is always efficient for the good to be provided. We carefully chose the set of parameters that define the game based on the theoretical and experimental results of CGK. Most importantly, the theoretical properties of this configuration allow us to identify questions about the effect of differences in network architectures that can be explored using the experimental data.

Although the multiplicity of sequential equilibria means that standard theory makes only weak predictions about the outcome of the game, the actual behavior we observe is predictable and highly sensitive to the network architecture. We emphasize that even if subjects do not "play" an equilibrium strategy profile, the specific architecture of the network clearly induces some patterns of contributions more than others, both with regard to the identity of contributors and the timing of their contributions. Such coordination may not only lead to more predictable behavior, it can also improve the efficiency of the outcome. Note that even if the public good is provided, the outcome may be inefficient because subjects contribute too much. And, of course, if the good is provided with probability less than one, it is of considerable interest to know how often it is provided and why.

The main regularities we observe can be summarized under four headings:

- **Strategic commitment:** There is a tendency for subjects in certain network positions to make contributions early in the game in order to encourage others to contribute. Clearly, commitment is of strategic value only if it is observed by others. Strategic commitment tends to be observed among *observed* and *uninformed* subjects — subjects in positions where (i) they are observed by another position and (ii) they cannot observe other positions.

Among the positions where we observe strategic commitment in the laboratory are position B in the one-link network, position C in the line network, and position A in the star-in network. The effect is strongest for position C in the line network and appears to be associated with the high level of efficiency in that network.

- **Strategic delay:** There is a tendency for subjects in certain network positions to delay their decisions until they have observed a contribution by a subject in another position. Obviously, there is an option value of delay only if the decision depends on the information. Strategic delay tends to be observed among *informed* and *unobserved* subjects — subjects in positions where (i) they can observe other positions and (ii) they are not observed by another position.

We observe strong evidence of strategic delay among all subjects in positions where they can observe another subject, particularly in position A of the one-link network, position B of the line network, and position A of the star-out.

- **Mis-coordination:** We also identify situations in which there are problems coordinating on an efficient outcome. Mis-coordination tends to arise in networks where two players are *symmetrically* situated. In symmetric situations, it becomes problematic for two players to know either who should go first or, if only one is to contribute, which of two should contribute.

There is evidence of coordination failure in networks where two subjects, such as B and C in the star-out and star-in networks and A and B in the pair network, are symmetrically situated.

- **Equilibrium:** In some cases, the modal behavior corresponds with easily identifiable salient equilibria. This is not to claim that subjects are actually playing equilibrium strategies, just that the modal behavior corresponds to what some equilibria would predict.

The modal behavior of subjects in the line and star-out networks corresponds to the strategies that would be chosen in equilibria that involve strategic commitment by observed and uninformed players and strategic delay by informed and unobserved players.

Thus, our experiment finds empirical support for these ideas: there is strategic delay; there is strategic commitment; symmetry leads to mis-coordination (with some caveats); and in some networks where the degree of coordination is high, the modal behavior of the subjects corresponds to a single equilibrium or class of equilibria. There are anomalies, of course, and in those cases we investigate behavior at the level of the individual subject to determine whether these anomalies are systematic or attributable to only a few individuals.

The rest of the paper is organized as follows. A discussion of the core literature on salience and other related literatures is provided in Section 2. Section 3 describes the theoretical model and Section 4 outlines the research questions that we attempt to answer in the rest of the paper. Section 5 summarizes the experimental design and procedures. The results are gathered in Section 6 and Section 7 contains some concluding remarks. The paper also uses three data and technical online appendices for the interested reader. Sample experimental instructions are attached in Appendix I. In reporting our results, we focus on the subjects' behavior in the last 15 rounds of the experiment. The tables and figures based on the full 25 rounds of observations are presented in Appendix II.[1] Appendix III provides a more refined analysis and discussion of individual behavior.

2 Related literature

Our use of the term salient refers to structural properties of the game, particularly the dominance of strategic delay for some players and the effects of strategic commitment on other players. In somewhat related papers, Cooper et al. (1990), Van Huyck et al. (1990, 1991), and Straub (1995) studied coordination via payoffs-based notions, including risk- and payoff-dominance. Our concept of salience, based on structural properties of the game, is closer to the one explored by these authors than it is to the concept of "psychological" salience introduced by Schelling (1960) as part of his theory of focal

[1] Broadly speaking, the data from the full 25 rounds present a qualitatively similar picture, although there are some signs that subjects' coordination improved over time.

equilibria.

In Schelling's account, what makes an equilibrium focal is its psychological "frame," rather than its structural properties. He argued that, in the description of a pure coordination game with multiple equilibria, the labels of the strategies may have an effect on the players' behavior. When there is no other reason to choose among a set of strategies, players will choose the strategy with the most salient label. The resulting equilibrium is called a *focal point*. Lewis (1969) used the concept of salience as an element of his theory of conventions (see also Cubitt and Sugden, 2003). Tests of Schelling's notion of salience in the context of one-shot coordination games are provided by Crawford and Haller (1990), Mehta et al. (1994), Sugden (1995), Bacharach and Bernasconi (1997), Blume (2000), Bardsley et al. (2006) and Crawford et al. (2008).

Our paper is also related to the large literature on coordination games in experimental economics (see Crawford, 1997, Camerer, 2003, and Devetag and Ortmann, 2007 for comprehensive discussions). Of particular interest are several articles that examine the effect of sequential timing in resolving coordination problems under conditions of imperfect information. The first such paper of which we are aware is Cooper et al. (1993) who examine strategic behavior in a "blinded" sequential battle of the sexes game, where the first player's move is not revealed to the second player but the timing is common knowledge. They find much higher rates of coordination and higher efficiency compared to a simultaneous-play version of the game. Rapoport (1997) investigates similar issues in a three-person battle of the sexes game and Budescu et al. (1997) report results from a similar information treatment applied to common-pool resource dilemmas. Weber et al. (2004) provide a summary of experimental findings about these effects, which they call "virtual observability," and run some additional coordination-game experiments, in which they observe similar effects. These studies are different from ours in several respects. First, they consider only the no-information case, where previous moves were completely unobservable by later players. Second, a single player moves in each stage of the game, whereas, in our study, all players move simultaneously in each stage of the game.

There is a small body of work on monotone games *with perfect information*. Admati and Perry (1991) introduced the basic concepts and their work was extended by Marx and Matthews (2000). Gale (1995, 2001) developed the theory applied in this paper in two different environments. Duffy et al. (2007) investigated the model of Marx and Matthews (2000) experimentally and replicated efficient outcomes in a dynamic laboratory setting. Following the seminal paper of Erev and Rapoport (1990), a number of experimental papers analyzed the effect of the information structure on public-good provision. Most recently, Ngan and Au (2008) extended Erev and Rapoport (1990) to investigate the effect of information in a real-time step-level public good game. While the games we focus on share a number of features, we address very different questions. Most importantly, the previous literature mainly focuses on the sensitivity of provision in different information treatments, without invoking a theoretical perspective, whereas we focus explicitly on the correspondence between equilibrium behavior and empirical behavior.

Also related to our network design, but somewhat further afield, there is a large and growing literature on the economics of networks (see Jackson, 2008). Although network experiments in economics are recent, there is now a large experimental literature on the economics of networks. (see Kosfeld, 2004, Goyal, 2005, and Jackson, 2005, for excellent, if now already somewhat dated, surveys). To the best of our knowledge, all of the previous experimental work on networks have quite different focuses than ours.

3 Equilibrium properties

Next, we define the game and discuss the properties of the equilibrium set for the different networks, paying particular attention to incentives for strategic commitment, strategic delay, and mis-coordination and the existence of salient equilibria.

3.1 The game

We study a dynamic game in which there are three players indexed by $i = A, B, C$, and three periods indexed by $t = 1, 2, 3$. Each player has an endowment of one token that he can contribute to the production of a public good. The contribution can be made in any of the three periods, but the decision is irreversible: once a player has committed his token, he cannot take it back. Let x_{it} denote the amount contributed by player i at the end of period t. Then we can represent the state of the game in period t by a vector

$$x_t = (x_{At}, x_{Bt}, x_{Ct}) \in \{0, 1\} \times \{0, 1\} \times \{0, 1\}.$$

The fact that the players' decisions are irreversible implies that $x_{it+1} \geq x_{it}$ for each player i; or, in vector notation, $x_{t+1} \geq x_t$. The initial state of the game is defined to be $x_0 = (0, 0, 0)$.

The players' payoffs are functions of the final state of the game $x_3 = (x_{A3}, x_{B3}, x_{C3})$. We assume that the public good is indivisible and costs two tokens to produce. The good is provided if and only if the total contribution is at least two tokens. If the public good is provided, each player receives a payoff equal to two tokens *plus* his initial endowment of one token *minus* his contribution. If the public good is not provided, each player receives a payoff equal to his initial endowment *minus* his contribution. Then the payoff of player i is denoted by $u_i(x_3)$ and defined by

$$u_i(x_3) = \begin{cases} 2 + (1 - x_{i3}) & \text{if } X \geq 2 \\ 1 - x_{i3} & \text{if } X < 2, \end{cases}$$

where $X \equiv x_{A3} + x_{B3} + x_{C3}$ denotes the total contribution at the end of the game. Note that the aggregate endowment and the aggregate value of the public good are greater than its cost, so that provision of the good is always feasible and efficient. However, as will be shown below, the coordination problem cannot necessarily be solved if each player has imperfect information about the actions of players in the same network.

To complete the description of the game, we have to specify the information available to each player. The information structure is represented by a directed graph or network. The network architecture is common knowledge. A player i can observe the actions of another player j, if and only if there is a directed edge leading from player i to player j. If player i can observe player j then i will know, at the beginning of period $t+1$, the history of j's contributions up to period t.

The six networks we study are illustrated in Figure 1 above and are used as treatments in the experimental design. Each of these networks has a different architecture, a different set of equilibria, and different implications for the play of the game. The games defined by the networks possess multiple equilibria, so theoretical analysis alone does not tell us which outcomes are likely to be observed; for that we need experimental data. Nonetheless, thinking about the equilibrium set does help us make some intuitive guesses about which outcomes might "stand out" or "suggest themselves" to human subjects.

3.2 The empty network

To illustrate the implications for equilibrium behavior of the different networks and information structures, we consider a series of theoretical examples of the underlying game. We begin with the empty network, which serves mainly as a benchmark to which the other networks can be compared. In the empty network, no player can observe any other player. Although a player can make his contribution in any of the three periods, the fact that no one receives any information in any period makes the timing of the decision irrelevant. This game is essentially the same as the one-shot game in which all players make simultaneous, binding decisions. More precisely, for each equilibrium of the one-shot game, there is a set of equilibria of the dynamic game that have the same outcome (probability distribution over the vector x_3). Conversely, for every equilibrium of the dynamic game, there is an equilibrium of the one-shot game with the same outcome (probability distribution over the vector x_1).

The one-shot game has multiple equilibria: There are three pure-strategy Nash equilibria in which two players contribute and one does not and the good is provided with probability one. To see that this is an equilibrium strategy profile, note that the players who contribute would be worse off choosing not to contribute (since the public good would not be provided) and the one player who does not contribute would be worse off contributing (since his contribution would not increase the provision of the public good). Conversely, there exists a pure-strategy Nash equilibrium in which no player contributes and the good is not provided. Obviously, if a player thinks that no one else will contribute, it is not optimal for him to contribute. Finally, the one-shot game possesses a symmetric mixed-strategy equilibrium where each player contributes with probability $1/2$ because each player is indifferent between contributing and not contributing.[2]

Each of the equilibria of the one-shot game has its counterpart in the dynamic game. For example, consider the pure-strategy equilibrium in which A and B contribute and C does not. In the dynamic game A and B can choose different periods in which to contribute or even randomize over periods. But as long as they contribute with probability one before the end the game, their strategies constitute an equilibrium of the dynamic game. Theory alone can-

[2]Positive provision of the good in equilibrium depends crucially on the fact that each contributing player is *pivotal* in the sense that, *at the margin*, his contribution is necessary and sufficient for provision (see, Bagnoli and Lipman, 1992 and Andreoni, 1998).

not provide convincing guesses about which of these multiple equilibria will occur.

3.3 The one-link network

In the empty network all players are symmetrically situated. Adding one link to the empty network creates a simple asymmetry among the three players. Now A can observe B's past contributions and condition his own decision on what B does, while B and C observe nothing. The addition of a single link eliminates one of the equilibrium outcomes present in the empty network. The pure-strategy sequential equilibrium with zero provision is not an equilibrium in the one-link network. To see this, suppose to the contrary that there exists an equilibrium in which no one contributes and consider what happens if B deviates from this equilibrium strategy and contributes in period 1. At the beginning of period 2, A knows that B has contributed and he knows that C does not know this. Then A knows that C will not contribute (C believes he is in the original equilibrium) and it is a dominant strategy for A to contribute. Anticipating this response, B will contribute before the final period of the game, thus upsetting the equilibrium.

The remaining equilibria of the one-shot game have their counterparts in the dynamic game with the one-link network (as well as in dynamic games with two-link networks, discussed below). These equilibria can be implemented if players simply wait until the final period and then use the strategies from the one-shot game. In addition to these simple replications of the one-shot equilibria, there are variations in which the players choose to contribute in different periods or randomize their strategies. Nevertheless, the salient feature of the one-link network is the fact that A observes B. Therefore, although there are many sequential equilibria, those in which B contributes first and A contributes after observing B contribute, seem salient. Whether or not we observe equilibrium play in the laboratory, the natural asymmetry suggests that A has an incentive to delay in order to observe whether B contributes and, conversely, B has an incentive to commit in order to encourage a contribution by A.

3.4 The two-links networks

The remaining networks can each be obtained by adding a single link to the one-link network. Each of these networks has a variety of sequential equilibria, but all of them are characterized by a positive probability of the provision of the public good. Besides the one-link network, the line network is the only network where all players are asymmetrically situated. The difference between the line and the one-link networks is that B can now observe C. As a result, A is now forced to make inferences about what B has observed, which makes the reasoning required to identify the optimal strategy quite subtle. As in the one-link network, there is an incentive for one player to contribute in order to encourage the player observing him, but there are two possible pairs that can do this: either B contributes first to encourage A or C contributes first to encourage B. Both possibilities are consistent with equilibrium. Among others, there are pure-strategy equilibria in which B contributes first, A contributes second and C does not contribute. There are also pure-strategy equilibria in which C contributes first, B contributes second and A does not contribute. Hence, the asymmetry alone cannot fully identify which of the many equilibria are likely to emerge.

In the other two-link networks, two players are symmetrically situated, which may intensify the coordination problem. In the star-out network, A is the center of the star and observes the behavior of the two peripheral players, B and C, while the peripheral players observe nothing. For A, it is weakly dominant to wait until the last period of the game to see whether his contribution is necessary to provide the public good. For B and C, there is a tension between their desire to contribute in order to encourage A and the desire to be a free rider and let the other peripheral player contribute. It is only necessary for one of the peripheral players, B or C, to encourage A. If both contribute, there is no need for A to contribute at all. The tension between encouragement and free-riding presents B and C with a coordination problem. In general, mis-coordination can result in either *under-provision*, where total contribution is less than two tokens and the good is not provided, or *over-provision*, where the total contribution is strictly greater than the cost of the good. In the star-out network, over-provision is less likely to occur, since player A will not contribute if he sees both B and C contribute. Under-provision may well occur, however, since B and C might each expect the other to contribute.

The star-in network is like the preceding one, but with the direction of the edges reversed. Now A observes nothing and is observed by B and C. A has an opportunity to encourage contribution by B and C, but this puts B and C in a quandary. Only one of them needs to contribute. Which one should it be? Alternatively, A might feel that if he refuses to contribute, it will be common knowledge and the two peripheral players will be forced to contribute. Either way, the difficulty of coordinating when neither peripheral player can observe the other may result in a coordination failure, leading either to over-provision or under-provision. Finally, in the pair network, A and B observe each other, while C neither observes nor is observed by A and B. This network is obtained by adding the edge leading from B to A to the one-link network. This may cause a kind of coordination problem which is different from those in the star-in and -out networks. Because A and B observe each other, each has an incentive to go first (to encourage the other) and to delay (to see what the other will do). This could result in under-contribution and non-provision.

4 Research questions

In this section, we use the equilibrium properties described in the previous section to identify questions that can be explored using the experimental data. Because each of the networks we study has a large number of equilibria, the theory does not make strong predictions. Which of these equilibria is the most plausible and whether equilibrium play is observed in the laboratory are empirical questions. Even if the experimental data do not conform exactly to one of the multiple equilibria, the data may suggest that some equilibria are empirically more relevant than others.

A subject who observes one or more other subjects is called *informed*; otherwise he is called *uninformed*. We have suggested that a subject who is uninformed and observed by one or more other subjects, has an incentive to contribute early in order to encourage the other subjects to contribute. In the one-link network, the subject in position B can contribute early in order to encourage A to contribute. Similarly, in the line network, C can encourage B. In the star-in network A can encourage B and C, and vice-versa in the star-out network. Hence, we are led to the following question:

Question 1 (strategic commitment) *Do subjects who are uninformed and observed by one or more subjects make a contribution early in the game to encourage other subjects to contribute?*

An informed subject has an incentive to delay his contribution until the final period of the game in order to gain information about the contributions of the subjects he observes. In the one-link network, it is a (weakly) dominant strategy for the subject in position A to wait to see whether B has contributed. In the line network, A has an incentive to wait until he has observed B contribute, but B has a similar incentive to wait until he has observed C. In the star-out network, A has an incentive to wait until he has observed whether the subjects in position B and C contribute, and vice-versa in the star-in network. This raises the following question:

Question 2 (strategic delay) *Do informed subjects delay their contributions until they have observed another subject contribute?*

When two subjects are symmetrically situated in a network, the symmetry may give rise to coordination problems. In the star-out network, the subjects in positions B and C are symmetrically situated; as a result, each has an incentive to commit early in order to encourage A, but each subject also has an incentive to be a free rider. In the star-in network, B and C are symmetrically situated; as a result, they both have an incentive to delay in order to to observe A, but each of them also has an incentive to be a free rider if A contributes. In the pair network, B and C are symmetrically situated; as a result, both have an incentive to commit in order to encourage the other and both have an incentive to delay; the difficulty of deciding who goes first may lead to a coordination failure. This raises our next question:

Question 3 (mis-coordination) *Do subjects who are symmetrically situated in a network have difficulty coordinating on an efficient outcome?*

The first two questions above are based on *local* properties of the networks, that is, on the edges into and out of a particular position. The third question is based on *global* properties. In general, one expects global properties to matter. For example, position C is locally the same in the empty network and the one-link network, but we expect different behavior for a subject in these positions precisely because the network structures differ with respect to positions A and B. This observation can be applied to any pair of networks and leads to the following question:

Question 4 (global properties) *Do subjects who are otherwise similarly situated behave differently in different networks?*

Finally, we raise a question about the relationship between equilibrium and empirical behavior. It is very difficult to establish that subjects are behaving consistently with equilibrium, partly because there are so many equilibria and partly because individual behavior is heterogeneous. However, it is worth asking whether, in some cases, the modal behavior in each position constitutes an equilibrium strategy profile. This raises the following question:

Question 5 (equilibrium behavior) *Does the profile of modal behaviors constitute an equilibrium strategy profile for some networks?*

5 Design and procedures

The experiment was run at the Princeton Laboratory for Experimental Social Science (PLESS). The subjects in this experiment were Princeton University students. After subjects read the instructions (see Appendix I), the instructions were read aloud by an experimental administrator. Throughout the experiment we ensured anonymity and effective isolation of subjects in order to minimize any interpersonal influences that could stimulate cooperation. Each experimental session lasted about one and a half hours. Payoffs were calculated in terms of tokens and then converted into dollars, where each token was worth \$0.50. A \$10 participation fee and subsequent earnings, which averaged about \$22, were paid in private at the end of the session.

Aside from the network structure, the experimental design and procedures described below are identical to those used by CGK. We studied the six network architectures depicted in Figure 1 above. The network architecture was held constant throughout a given experimental session. In each session, the network positions were labeled A, B, or C. A third of the subjects were designated type-A participants, one third type-B participants and one third type-C participants. The subject's type, A, B, or C, remained constant throughout the session.

Each session consisted of 25 independent rounds and each round consisted of three decision turns. The following process was repeated in all 25 rounds. Each round started with the computer randomly forming three-person groups by selecting one participant of type A, one of type B and one of type C. The groups formed in each round depended solely upon chance and were independent of the networks formed in any of the other rounds. Each group played a dynamic game consisting of three decision turns.

At the beginning of the game, each participant had an endowment of one token. At the first decision turn, each participant was asked to allocate his tokens to either an x-account or a y-account. Allocating the token to the y-account was irreversible. When every participant in the group had made his decision, each subject observed the choices of the subjects to whom he was connected in his network. This completed the first of three decision turns in the round.

At the second decision turn, each subject who allocated his token to the x-account was asked to allocate the token between the two accounts. At the end of this period, each subject again observed the choices of the subjects to whom he was connected in his network. This process was repeated in the third decision turn. At each date, the information available to subjects included the choices they had observed at every previous date.

When the first round ended, the computer informed subjects of their payoffs. The earnings in each round were determined as follows: if subjects contributed at least two tokens to their y-accounts, each subject received two tokens plus the number of token remaining in his x-account. Otherwise, each subject received the number of tokens in his x-account only. After subjects were informed of their earnings, the second round started by having the computer randomly form new groups of participants in networks. This process was repeated until all the 25 rounds were completed.

There were two experimental sessions for each network.[3] Each session comprised either 12, 15, 18, or 21 subjects. The diagram below summarizes the experimental design and the number of observations in each treatment (the entries have the form a / b where a is the number of subjects and b the number of observations per game). Overall, the experiments provide us with a very rich dataset. We have observations on 1525 games in a variety of different networks.

	Networks					
Session	Empty	One-link	Line	Star-out	Star-in	Pair
1	12 / 100	15 / 125	15 / 125	18 / 150	15 / 125	18 / 150
2	15 / 125	12 / 100	21 / 175	15 / 125	15 / 125	12 / 100
Total	27 / 225	27 / 225	36 / 300	33 / 275	30 / 250	30 / 250

[3]The two sessions for each treatment were identical except for the number of participants and the *labeling* of the nodes of the graphs, which we changed in order to see whether the labels were salient. As far as we could tell, they were not. We observed "session effects" in the one-link treatment, but these were caused by three subjects in one session.

6 Results

In this section, we present the experimental results. To organize the results, we draw a distinction between three types of analysis. We begin our analysis with a purely descriptive overview of some important features of the experimental data, which we summarize by reporting contribution rates in the various network treatments. In this stage of the analysis we are interested in describing the data rather than providing a formal test of predictions or of equilibrium behavior. The experiment was not designed to compare behavior across network treatments, and the multiplicity of equilibria makes such comparative static properties scarce. We then move to a non-parametric analysis of the relationship between the strategic behavior suggested by our discussion of the theory (in the form of strategic commitment and strategic delay) and the data. The analysis is mainly focused on qualitative shifts in subjects' behavior resulting from changes in the network architecture. Finally, to better understand the mechanisms underlying our subjects' decisions we turn to investigating behavior at the level of the individual subject. Not surprisingly, there is some heterogeneity across subjects. Nevertheless, the choices made by most of our subjects reflect cleanly classifiable strategies which are stable across decision-rounds. To economize on space, the individual level analysis is provided in Appendix III.

6.1 Overview

The top panel of Table 1 reports the total contribution rates networks. The efficiency of behavior depends on the total number of contributions, not just the provision rate. More precisely, inefficiency can arise from under-contribution ($X < 2$) and from over-contribution ($X > 2$). In order to highlight the differences in efficiency across networks, we also tabulate the rates of under-contribution, efficient contribution ($X = 2$), and over-contribution. In the bottom panel of Table 1, the outcomes of each pair of networks are compared using the Wilcoxon (Mann-Whitney) rank-sum test. In the last row, the provision rate in each network is compared to the empty network.[4]

[4]These tests assume independence, which would be satisfied, for example, if the subjects in a given session use identical mixed strategies. If there is heterogeneity among subjects, however, the outcomes of games in which the same subjects appear will not be independent. This biases the standard errors downwards, increasing the likelihood of finding a significant treatment effect. There is no simple adjustment to the standard test that will take care of

[Table 1 here]

The highest provision rate (0.839) is observed in the line network and the smallest (0.570) is observed in the empty network. The empty network is isomorphic to the one-shot game in which players choose their strategies simultaneously. The provision rate in the symmetric mixed strategy equilibrium of the one-shot game is 1/2, which is similar to the empirical provision rate in the empty network. There are also considerable variations in efficiency across networks. The line and the star-out networks are the most efficient, whereas the empty, one-link, and pair networks are the least efficient. This suggests that there is something about the structure of the line and star-out networks that allows subjects to coordinate efficiently. We return to this question later.

The highest rate of under-contribution is observed in the empty network (0.430). Again, the predicted under-contribution rate in the symmetric mixed strategy equilibrium of the one-shot game is 1/2, which is similar to the empirical under-contribution rate in the empty network. The highest over-contribution rate (0.170) is found in the one-link network, which also has a high under-contribution rate (0.393). We also observe high under-contribution and over-contribution rates in the pair network (0.407 and 0.107), which appears to indicate a mis-coordination problem, discussed further later in the paper.

Next, Table 2 presents the timing of contributions across networks for uninformed and informed subjects. The contribution rates are defined as the ratio of the number of contributions to the number of uncommitted subjects, i.e., the number of subjects who still have a token to contribute. We sometimes refer to these as *conditional* contribution rates. The number in parentheses in each cell represents the number of uncommitted subjects (subjects who have an endowment left for contribution). The last column of Table 2 reports total contribution rates.

[Table 2 here]

For uninformed subjects (top panel), most contributions were made in the first period. The tendency of uninformed subjects to make early contributions is found in all networks, but the contribution rates in the first period

the possible dependence so we have used the null of independence while recognizing that it may not be satisfied in this case.

and the total contribution rates vary considerably across networks and positions. Most strikingly, in the line network position-C subjects contributed in the first period in most rounds. For informed subjects (middle panel), by contrast, there is a general tendency to delay. The modal behavior of subjects in position B of the line network is to contribute in the second period. Given the early contribution behavior of position-C subjects in this network, this indicates that position-B subjects delay their contribution until they observe that C has contributed. Finally, the isolated subjects (bottom panel) in the one-link and pair networks maintained low contribution rates across the three periods of the game.

6.2 Strategic commitment

Result 1 (strategic commitment) *There is a strong tendency for subjects who are uninformed and observed by others to contribute early. Specifically, subjects in positions B (one-link), C (line), and A (star-in) exhibit strategic commitment. This effect is strongest for position C (line) and is associated with a high level of efficiency in that network.*

In Section 3 we suggested that an uninformed subject observed by one or more other subjects has an incentive to make an early contribution in order to encourage the observer(s) to contribute. In particular, subjects occupying position B (one-link), C (line), and A (star-in) should, according to this reasoning, tend to contribute in the first period. Figure 2 shows the frequencies of contributions across time by uncommitted subjects occupying these positions in the three networks. We also include subjects in position B (line). This position is different from the others included in Figure 2, because it is both observed by position A and observes position C. Thus, in the line network, subjects in position B may be torn between the incentive to contribute early and the incentive to delay. The number above each bar in the histogram represents the number of observations.

[Figure 2 here]

The histograms in Figure 2 show that subjects in positions B (one-link), C (line), and A (star-in) all exhibit a tendency toward early contributions, but the actual contribution rates vary. Most noticeably, C (line) has a much higher contribution rate than the other two positions: the contribution rate in the first period is 0.900 for C (line), whereas the corresponding rates for B (one-link) and A (star-in) are 0.570 and 0.620, respectively.[5] This is another reflection of the greater efficiency of the line network. Given the strategic commitment of C (line), we note that subjects in position B (line) have more in common with informed subjects than with subjects who are uninformed and observed: most subjects in position B (line) contribute in the second and third periods, although there are a few subjects contributing in the first period.

One puzzling feature of the data is the similarity of the contribution rates at positions B (one-link) and A (star-in). Unlike B (one-link), A (star-in) may have an incentive to delay if he thinks that he can signal to B and C (star-in) that he is determined to be a free rider and force the other two to contribute. Thus, coordination in the star-in network would appear to be more difficult than in the one link. Nonetheless, we observe similar contribution rates at the two positions. In fact, the inefficiency (lack of coordination) in the one-link network is in general a puzzle to us. A priori, one would expect that the contribution rates of subjects at position B (one-link) and at position C (line) should be reversed. We address this question in Appendix III by investigating behavior at the level of the individual subject.

6.3 Strategic delay

Result 2 (strategic delay) *There is strong evidence of strategic delay by informed subjects. In particular, subjects at position A (one-link), B (line), and A (star-out), tend to delay their decisions until another subject has contributed.*

As we argued in Section 3, informed subjects have an incentive to delay making a decision to contribute until they observe that another subject has contributed. According to this argument, subjects in positions A (one-link), A and B (line), and A (star-out) should exhibit strategic delay. Informed subjects in positions B and C (star-in) and A and B (pair) also have an incentive to delay but, because of the symmetry of these positions in their respective network structures, the incentive to delay is confounded with the

[5] To test for the difference of means, one could estimate a probit or logit model that accounts for the statistical dependence of observations caused by the repeated appearance of the same subjects in our sample. Instead, we conducted an extensive individual-level analysis of the data, which is reported in Appendix III.

coordination problem. For this reason, we deal with these positions separately in the following section.

For the network positions of interest here, we present the subjects' contribution rates, conditional on their information states, in Figure 3 below. The information state is 1 if a contribution has been observed and is 0 otherwise. The number above each bar of the histogram represents the number of observations. There is a strong incidence of strategic delay for subjects in positions A (one-link), B (line) and A (star-out). Observing a contribution increases the subject's contribution rate by a factor of four. By contrast, the contribution rates for position A (line) are low in both states. This suggests that the behavior of subjects in position A (line) can be best described as free riding. But note that given the tendency of subjects in positions B and C (line) to contribute, the behavior of position-A subjects is optimal and efficient.

[Figure 3 here]

6.4 Mis-coordination

Result 3 (mis-coordination) *There is evidence of coordination failure in networks where two subjects, such as B and C (star-out, star-in) and A and B (pair), are symmetrically situated. Coordination failure explains the majority of inefficient outcomes in the star-out, star-in and pair networks.*

We have delayed the discussion of positions B and C (star-out, star-in) and A and B (pair), because they involve a coordination problem that complicates the analysis of incentives for strategic delay and strategic commitment. The common feature of these pairs of positions is that they are symmetrically situated in their respective networks. In the star-out network, B and C have an incentive to encourage A but, at the same time, they have an incentive to be free riders and let the other encourage A. In the star-in network, B and C have an incentive to delay in order to see whether A contributes but, once A has contributed, they have an incentive to be free riders and let the other provide the public good. In the pair network, A and B have both an incentive to encourage the other and an incentive to delay. This conflict may lead to inefficient outcomes. We next investigate the coordination problem in each of these networks. We begin with the star-out network.

6.4.1 The star-out network

We first investigate the coordination problem by revisiting the efficiency results presented in Table 2 above. The star-out network has the lowest rate of over-contribution (0.055) among all networks. This result is not surprising, since A plays the role of a central coordinator in the star-out network, waiting to see if the peripheral positions, B and C, contribute and only contributing himself, if necessary, in the last period. It is less obvious how much of the under-contribution rate (0.321) is attributable to mis-coordination between B and C. To answer this question, Figure 4 depicts the total contributions made by subjects in positions B and C in each period. The number above each bar of the histogram represents the frequency of contributions by position-A subjects in the corresponding state in the next period.

[Figure 4 here]

It is interesting that the frequency of no contribution by subjects in positions B and C during the first two periods is very close to the rate of under-contribution (0.321). This suggests that the under-contribution outcomes in the star-out network are mainly caused by a coordination failure between position-B and position-C subjects. We can check this by focusing on the 47 (out of 165) games in which neither B nor C contributed by the end of the second period. The public good was not provided in any of those games. This implies that 88.8 percent (= 0.285/0.321) of the total under-contribution rate is attributable to a failure by subjects in positions B and C to coordinate their contributions.

6.4.2 The star-in network

In the star-in network, we distinguish two types of coordination failures, one that occurs when position-A subjects contribute first and one that occurs when they try to free ride. We divide the sample according to the timing of contributions of position-A subjects, and re-calculate the efficiency results. The new results are presented in Figure 5 below. The numbers represent the total number of observations. One interesting feature of the data presented in Figure 5 is that, even when the subjects in position A contribute in the first two periods, the under-contribution rate is relatively high (0.180) purely because of a coordination failure between the subjects in positions B and C. On the other hand, when position-A subjects do not contribute, the under-contribution rate is very high (0.900), which strongly suggests that the coordination between B and C becomes more difficult when A does not contribute. Of course, the failure to coordinate depends on A's refusal to commit, so this could be interpreted as a failure of A to coordinate with B and C. In any case, the under-contribution rate when position-A subjects do not contribute (0.900) is much higher than the under-contribution rate in the benchmark empty network (0.430).

[Figure 5 here]

6.4.3 The pair network

In the pair network, the salient solution to the coordination problem is for A and B to contribute. According to this hypothesis, under-contribution should be attributed to coordination failure between the subjects in positions A and B, whereas over-contribution is attributable to contributions from subjects isolated in position C. In order to investigate the coordination failure between subjects in positions A and B, we simply compute the relative frequency that positions A and B subjects fail to contribute two tokens. This turns out to be surprisingly high (0.407). The uncoordinated contributions of position-C subjects sometimes lead to over-contribution and sometimes compensate for under-contribution by subjects in positions A and B. On average, as one would expect, these contributions have no effect on efficiency. In fact, the under-contribution rate (0.407) is identical to the frequency of under-contribution by subjects in positions A and B. So we can argue that under-contribution in the pair network is driven by the coordination failure between subjects in positions A and B. Over-contribution, on the other hand, is clearly the result of uncoordinated contributions by position-C subjects.

6.5 Global properties

Result 4 (global properties) *There is strong evidence that the global properties of the networks, as well as the local properties, are important determinants of subjects' behavior. One example is the behavior of isolated individuals in the empty, one-link and pair networks.*

Strategic interaction between subjects is often influenced by the local properties of the networks, that is, by the links into and out of a particular position.

On the other hand, there are instances where the global properties of the network have a large influence on the behavior of subjects. One easy test of the importance of global properties is a comparison of the behavior of *isolated* subjects, that is, subjects in positions A, B, or C (empty), C (one-link), and C (pair). These positions have no inward or outward links, so if only the local properties matter, the behavior of the subjects in these positions should be identical in all three networks. However, we observe considerable differences in contributions across the three networks at both the aggregate level and the individual level. From the last column of Table 2 above, the contribution rate in the empty network is over twice as high as that of subjects in positions C (one-link) and C (pair) (0.531 compared to 0.230 and 0.167, respectively).

Next, we compare the patterns of contribution behavior of subjects in positions A and B (one-link) with either positions A and B (line) or B and C (line). It is interesting to observe how subjects' behavior changes as the result of one additional link from B to C. In Figures 2 and 3 above, we observe that subjects in position C (line) have qualitatively the same behavior as the subjects in position B (one-link), since in both positions subjects make their contribution in the first period. Similarly, subjects in position B (line) exhibit strategic delay just as subjects in position A (one-link) do. Nonetheless, the line network achieves much higher efficiency than the one-link network, even though the presence of two informed positions, A and B, and two observed positions, B and C, in the line network suggests the possibility of coordination problems.

6.6 Equilibrium

Result 5 (equilibrium) *The modal behavior of subjects in the line and star-out networks corresponds to what some equilibria would predict. In addition, there are considerable differences in the modal behavior of subjects in these networks, indicating that different equilibria might be plausible or salient.*

Because of the large number of equilibria in the games we studied, the theory does not have much to say about the kinds of behavior we should expect to see in the laboratory. Instead, we have emphasized the usefulness of experimental data for identifying which equilibria might be plausible or salient. Now we consider three cases in which the subjects' behavior approximates a salient equilibrium. One has to be very careful in making claims that individual subjects are playing equilibrium strategies. Given the multiplicity of equilibria and the heterogeneity of individual behavior, it is unlikely that all subjects coordinate on a single equilibrium. The most that we can claim is that the modal behavior of the subjects bears a striking resemblance to a particular equilibrium, while noting that there are considerable deviations from equilibrium on the part of some subjects. We have already alluded to the coordination problems found in the pair and star-in networks. We will thus not attempt to reconcile subjects' behavior in these networks with equilibrium behavior. Instead, we focus on the one-link, the line, and the star-out networks. We begin by considering the line network.

6.6.1 The line network

In the line network, the degree of coordination reflected by the efficiency of outcomes appears to be very high. The frequencies of contributions in different positions and information states are tabulated in Table 3. The states 0 and 1 in the table refer to the number of contributions observed by subjects in positions A and B in periods 2 and 3. Note that in order to reduce the number of states, we pool the data corresponding to a given number of contributions, regardless of when the contributions were made. The number in parentheses in each cell represents the number of observations.

[Table 3 here]

The first thing to note is the very high contribution rate (0.900) of subjects in position C in period 1. Secondly, subjects in position B contribute mainly after they observe a contribution by the subject in position C. More precisely, the contribution rate in position B, conditional on observing no contribution by C, is 0.077 in period 2 and 0.182 in period 3. By contrast, the contribution rate in position B, conditional on observing a contribution by C, is 0.632 in period 2 and 0.686 in period 3. Finally, the total contribution rate by subjects in position A is only 0.106. This regularity suggests (an equivalence class of) equilibria in which C contributes in period 1, B contributes after observing C contribute, and A does not contribute at all. There are deviations from this equilibrium pattern, notably the contributions by subjects in position B when they have not observed any contribution by the subject in position C. But these deviations are not large and the behaviors of subjects in positions C and A are very close to those predicted by this class of equilibria.

There are some interesting cases in the data where subjects deviate from the suggested equilibrium behavior. Position-A subjects are most likely to contribute if they observe that the subject in position B has contributed in period 1, that is, before he can observe the subject in position C contribute. Subjects in position A may have reasoned that this behavior was intended to encourage them to contribute and, in any case, preempts any possible revelation of the behavior of the subject in position C. Given the high probability that subjects in position C contribute in period 1, such reasoning by subjects is faulty, but it is interesting nonetheless. In period 3, we notice that subjects in position A are less likely to contribute if the subject in position B has contributed in periods 1 or 2; most of these observations are cases in which B contributed in period 2, thus signaling an earlier contribution by C. These observations suggest some rationality, even if they do not correspond exactly to the proposed equilibrium.

6.6.2 The star-out network

The next case we consider is the star-out network. The frequencies of contributions in different positions and information states are summarized in Table 4 below. The states 0, 1, and 2 refer to the number of contributions observed by subjects in periods 2 and 3. Again, in order to reduce the number of states, we pool the data corresponding to different histories that lead to the same information state. The number in parentheses in each cell represents the number of observations.

Here we see an extreme illustration of strategic delay by position-A subjects: out of 165 observations, there are only 19 contributions in the first two periods. Most of these occur in period 2 after one of the peripheral subjects in positions B or C has contributed. Although further delay would be optimal, the deviation from rational behavior seems small. By contrast, subjects in positions B and C have an incentive to contribute early to encourage the subject in position A, and on average they contribute in the first two periods 0.455 of the time. In the last period, their contribution rate falls precipitously to 0.044. The patterns here suggest (an equivalence class of) equilibria in which B and C contribute in the first two periods with probability 1/2 and contribute with probability 0 in the last period, while A waits until the last period and contributes only if he observes exactly one contribution by B or C in the preceding periods.

[Table 4 here]

Also notice that the timing of contributions by the subjects in positions B and C matters only to the extent that the total probability of contribution in the first two periods must be $1/2$ in equilibrium; the contribution probability in *individual* periods is immaterial. Thus, the fact that subjects contribute in those two periods with probability 0.455 is what matters; the contribution rates in period 1 and in period 2 are irrelevant. Position-A subjects match the prescribed behavior very closely in period 1 and period 3. Only in period 2 is there a significant deviation. In three cases, subjects in position A contributed in period 2 after observing two contributions in the previous period. The numbers are very small and should be attributed to the 'trembling hand.'

6.6.3 The one-link network

Finally, we consider the one-link network. Table 5 below summarizes the frequencies of contributions in different positions and information states in the one-link network. The number in parentheses in each cell represents the number of observations. Note that conditional on observing the subject in position B contribute, the contribution rates of subjects in position A are 0.500 and 0.583 in periods 2 and 3, respectively. It appears that subjects are randomizing, but the contribution rate of subjects in position C, 0.229, is much too low to make subjects indifferent between contributing and not contributing. Likewise, when subjects in position A do not observe the subject in position B contribute, it cannot be optimal for them to randomize in periods 2 and 3: the contribution rates of subjects in position C and subjects in position B in period 3 are too low.

[Table 5 here]

The data summarized in Table 5 are mixed, with several features that are difficult to reconcile with equilibrium behavior. By analogy with our findings in the line network, one might expect the salient equilibrium to be one in which B contributes first, A contributes after observing B contribute, and C never contributes. The bare facts appear inconsistent with this prediction. Overall, the isolated subjects in position C contribute on average 0.229 of the time. Similarly, subjects in position A contribute 0.178 of the time *without* having observed a contribution by the subject in position B. Even when they have observed a contribution by the subject in position B, the contribution rate of subjects in position A is only 0.534. One anomaly here appears to be the contribution behavior of subjects in position C. Since they can neither observe nor be observed, they have no ability to coordinate and yet they make a significant number of contributions. Since subjects learn the outcome of the game at the end of each round, subjects in position A may become aware of the contribution behavior of subjects in position C and decide to free ride to some extent. What cannot be ascertained from the information given in Table 5 is whether these anomalies are endemic or caused by a few subjects. To pursue this question, it is necessary to investigate behavior at an individual level, which is provided in Appendix III. Whatever the explanation, it is hard to argue that the average behavior of subjects in position A is optimal.

7 Conclusion

Our conclusion is that asymmetry in the network architecture is an important factor in creating the salience of certain strategies. Asymmetric networks give different roles to different subjects, making their behavior more predictable and aiding the coordination of their actions. These networks encourage strategic commitment in some positions, strategic delay in others, and passivity in still others (isolated subjects, who can neither observe nor be observed, are less likely to contribute). These features of network architecture make certain behaviors – and possibly certain equilibria – salient. The bottom line is that asymmetry gives rise to salience which, in turn, is an aid to predictability and coordination. These regularities lack a proper theoretical explanation, of course. For the time being we are forced to leave them as puzzles for theorists to ponder.

There is clearly a lot more to be done and the uses of our data set are far from exhausted. Perhaps most importantly, although we would like to focus more explicitly on the correspondence between theory and empirical behavior by estimating a theoretically grounded structural model, we have not yet found an appropriate way of doing so. In particular, we have tried adapting the model of Quantal Response Equilibrium (QRE) of McKelvey and Palfrey (1995, 1998), but it seems intractable because of the multiplicity of equilibria. In fact, because of the imperfection of information and the simultaneous moves, even in the case of three-person networks, finding the set of QRE is computationally intensive. We have not found an alternative, theoretically grounded, structural model that might be brought to the interpretation of the data. This is an important topic for future research. We also have not yet explored alternative "behavioral" approaches that might be brought to the interpretation of the data. Finally, the model and results that we have developed provide a foundation for future research and the techniques can be applied to other setups such as dynamic graphs where the set of neighbors changes over time. We can also use the same methodology to examine more complex network architectures. Obviously, different network architectures may lead to different outcomes. To determine which factors are important in explaining subject behavior in a variety of settings, it will also be necessary to investigate a larger class of games in the laboratory.

References

[1] Admati A. and M. Perry (1991) "Joint Projects Without Commitment." *Review of Economic Studies* **58**, pp. 259-276.

[2] Andreoni J. (1998) "Toward a Theory of Charitable Fund-Raising." *Journal of Political Economy* **106**, pp. 1186-1213.

[3] Bacharach, M. and M. Bernasconi (1997) "The Variable Frame Theory of Focal Points: An Experimental Study." *Games and Economic Behavior* **19**, pp. 1-45.

[4] Bagnoli M. and B. Lipman (1992) "Private Provision of Public Goods Can Be Efficient." *Public Choice* **74**, pp. 59-78.

[5] Bardsley, N., J. Mehta, C. Starmer and R. Sugden (2006) "The Nature of Salience Revisited: Cognitive Hierarchy Theory versus Team Reasoning." Unpublished.

[6] Blume, A. (2000) "Coordination and Learning with a Partial Language." *Journal of Economic Theory* **95**, pp. 1-36.

[7] Budescu, D., W. Au, and X.-P. Chen (1997) "Effects of Protocol of Play and Social Orientation on Behavior in Sequential Resource Dilemmas." *Organizational Behavior & Human Decision Processes* **69**, pp. 179-193.

[8] Camerer C. (2003) *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton University Press.

[9] Choi S., D. Gale and S. Kariv (2008) "Sequential Equilibrium in Monotone Games: Theory-Based Analysis of Experimental Data." *Journal of Economic Theory* **143** , pp. 302–330

[10] Cooper, R., D. DeJong, R. Forsythe and T. Ross (1990) "Selection Criteria in Coordination Games: Some Experimental Results." *American Economic Review* **80**, pp. 218-233.

[11] Cooper, R., D. DeJong, R. Forsythe and T. Ross (1993) "Forward Induction in the Battle-of-the-Sexes Games." *American Economic Review* **83**, pp. 1303-1316.

[12] Crawford, V. and H. Haller (1990) "Learning How to Cooperate: Optimal Play in Repeated Coordination Games." *Econometrica* **58**, pp. 571-595.

[13] Crawford, V. (1997) "Theory and Experiment in the Analysis of Strategic Interaction." In *Advances in Economics and Econometrics: Theory and Applications*, Seventh World Congress, vol. I, ed. D. Kreps and K. Wallis. Cambridge University Press.

[14] Crawford, V., U. Gneezy and Y. Rottenstreich (2008) "The Power of Focal Points is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures." *American Economic Review* **98**, pp. 1443–58

[15] Cubitt R. and R. Sugden (2003) "Common Knowledge, Salience and Convention: A reconstruction of David Lewis' Game Theory." *Economic and Philosophy* **19**, pp. 175-210.

[16] Devetag G. and A. Ortmann (2007) "When and why? A critical survey on coordination failure in the laboratory." *Experimental Economics* **10**, pp. 331-344.

[17] Duffy J., J. Ochs and L. Vesterlund (2007) "Giving Little by Little: Dynamic Voluntary Contribution Games." *Journal of Public Economics* **91**, pp. 1708-1730.

[18] Erev I. and Rapoport A. (1990) "Provision of step-level public goods: The sequential contribution mechanism." *Journal of Conflict Resolution* **34**, pp. 401–425.

[19] Gale D. (1995) "Dynamic Coordination Games." *Economic Theory* **5**, pp. 1-18.

[20] Gale D. (2001) "Monotone Games with Positive Spillovers." *Games and Economic Behavior* **37**, pp. 295-320.

[21] Goyal S. (2005) "Learning in Networks: A Survey." In *Group Formation in Economics; Networks, Clubs and Coalitions*, ed. G. Demange and M. Wooders. Cambridge University Press.

[22] Jackson M. (2005) "A Survey of Models of Network Formation: Stability and Efficiency." In *Group Formation in Economics: Networks, Clubs, and Coalitions,* ed. G. Demange and M. Wooders. Cambridge University Press.

[23] Jackson M. (2008) *Social and Economic Networks.* Princeton University Press.

[24] Kosfeld, M. (2004), "Economic Networks in the Laboratory: A Survey." *Review of Network Economics* **3**, pp. 20-41

[25] Lewis D. (1969) *Convention: A philosophical study.* Harvard University Press.

[26] Marx L. and S. Matthews (2000) "Dynamic Voluntary Contribution to a Public Project." *Review of Economics Studies* **67**, pp. 327-358.

[27] Mckelvey R. and T. Palfrey (1995) "Quantal Response Equilibria for Extensive Form Games." *Games and Economic Behavior* **10**, pp. 6-38.

[28] Mckelvey R. and T. Palfrey (1998) "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* **1**, pp. 9-41.

[29] Mehta J., C. Starmer and R. Sugden (1994) "The Nature of Salience: An Experimental Investigation of Pure Coordination Games." *American Economic Review* **84**, pp. 658-673.

[30] Ngan C. S. and Au W. T. (2008) "Effect of Information Structure in a Step-Level Public-Good Dilemma Under a Real-Time Protocol." In *New Issues and Paradigms in Research on Social Dilemmas*, ed. A. Biel et al. Springer.

[31] Rapoport, A. (1997). "Order of Play in Strategically Equivalent Games in Extensive Form." *International Journal of Game Theory* **26**, pp. 113-136.

[32] Schelling, T. (1960) *The strategy of conflict*. Harvard University Press.

[33] Straub, P. (1995) "Risk Dominance and Coordination Failures in Static Games." *Quarterly Review of Economics and Finance* **35**, pp. 339-363.

[34] Sugden, R. (1995) "A Theory of Focal Points." *Economic Journal* **105**, pp. 533-550.

[35] Van Huyck, J., R. Battalio and R. Beil (1990) "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review* **80**, pp. 234-248.

[36] Van Huyck J., R. Battalio and R. Beil (1991) "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games." *Quarterly Journal of Economics* **106**, pp. 885-911.

[37] Weber, R., C. Camerer, and M. Knez (2004) "Timing and Virtual Observability in Ultimatum Bargaining and 'Weak Link' Coordination Games." *Experimental Economics* **7**, pp. 25-48.

Table 1. The total number of contributions and provision rate by network

Network	Total contributions				Under provision	Provision
	0	1	2	3		
Empty	0.059	0.370	0.489	0.081	0.430	0.570
One-link	0.141	0.252	0.437	0.170	0.393	0.607
Line	0.033	0.128	0.744	0.094	0.161	0.839
Star-out	0.224	0.097	0.624	0.055	0.321	0.679
Star-in	0.087	0.213	0.620	0.080	0.300	0.700
Pair	0.120	0.287	0.487	0.107	0.407	0.593

Wilcoxon (Mann-Whitney) rank-sum test - under (white) / over (gray)

	Empty	One-link	Line	Star-out	Star-in	Pair
Empty	--	0.028	0.690	0.353	0.964	0.469
One-link	0.537	--	0.046	0.001	0.021	0.119
Line	0.000	0.000	--	0.161	0.645	0.713
Star-out	0.053	0.199	0.001	--	0.367	0.088
Star-in	0.023	0.101	0.003	0.685	--	0.428
Pair	0.695	0.809	0.000	0.116	0.054	--
	--	0.537	0.000	0.053	0.023	0.695

Table 2. The evolution of contributions over time by uninformed and informed types

A. Uninformed

Network	Position	Period 1	Period 2	Period 3	Contribution rate
One-link	*B*	0.570 (135)	0.345 (60)	0.158 (42)	0.763
Line	*C*	0.900 (135)	0.167 (14)	0.200 (11)	0.933
Star-out	*B* , *C*	0.318 (135)	0.187 (92)	0.044 (75)	0.470
Star-in	*A*	0.620 (180)	0.439 (68)	0.094 (38)	0.807

B. Informed

Network	Position	Period 1	Period 2	Period 3	Contribution rate
One-link	*A*	0.104 (135)	0.306 (121)	0.429 (84)	0.644
Line	*A*	0.006 (180)	0.034 (179)	0.069 (173)	0.106
	B	0.172 (180)	0.584 (149)	0.597 (62)	0.861
Star-out	*A*	0.006 (165)	0.110 (164)	0.514 (146)	0.570
Star-in	*B* , *C*	0.157 (300)	0.170 (253)	0.205 (210)	0.443
Pair	*A* , *B*	0.300 (300)	0.352 (210)	0.353 (136)	0.707

C. Isolated

Network	Position	Period 1	Period 2	Period 3	Contribution rate
Empty	*A* , *B* , *C*	0.402 (405)	0.091 (242)	0.136 (220)	0.531
One-link	*C*	0.163 (135)	0.0354 (112)	0.0459 (107)	0.230
Pair	*C*	0.100 (150)	0.022 (135)	0.053 (132)	0.167

( ) - # of obs.

Table 3. The frequencies of contributions at different states in the line network

		A		*B*		*C*
1	n_i	--		--		--
	Freq.	0.006 (180)		0.172 (180)		0.900 (180)
2	n_i	0	1	0	1	--
	Freq.	0.007 (148)	0.161 (31)	0.077 (13)	0.632 (136)	0.167 (18)
3	n_i	0	1	0	1	--
	Freq.	0.115 (61)	0.045 (112)	0.182 (11)	0.686 (51)	0.200 (15)

( ) - # of obs.

Table 4. The frequencies of contributions at different states in the star-out network

		A			*B,C*
1	n_i	--			--
	Freq.	0.006 (165)			0.318 (330)
2	n_i	0	1	2	--
	Freq.	0.027 (73)	0.195 (77)	0.071 (14)	0.187 (225)
3	n_i	0	1	2	--
	Freq.	0.089 (45)	0.922 (77)	0.000 (24)	0.044 (183)

( ) - # of obs.

Table 5. The frequencies of contributions at different states in the one-link network

		A		*B*	*C*
1	n_i	--		--	--
	Freq.	0.104 (135)		0.570 (135)	0.163 (135)
2	n_i	0	1	--	--
	Freq.	0.039 (51)	0.500 (70)	0.345 (58)	0.035 (113)
3	n_i	0	1	--	--
	Freq.	0.222 (36)	0.583 (48)	0.158 (38)	0.046 (109)

( ) - # of obs.

Figure 1: The networks

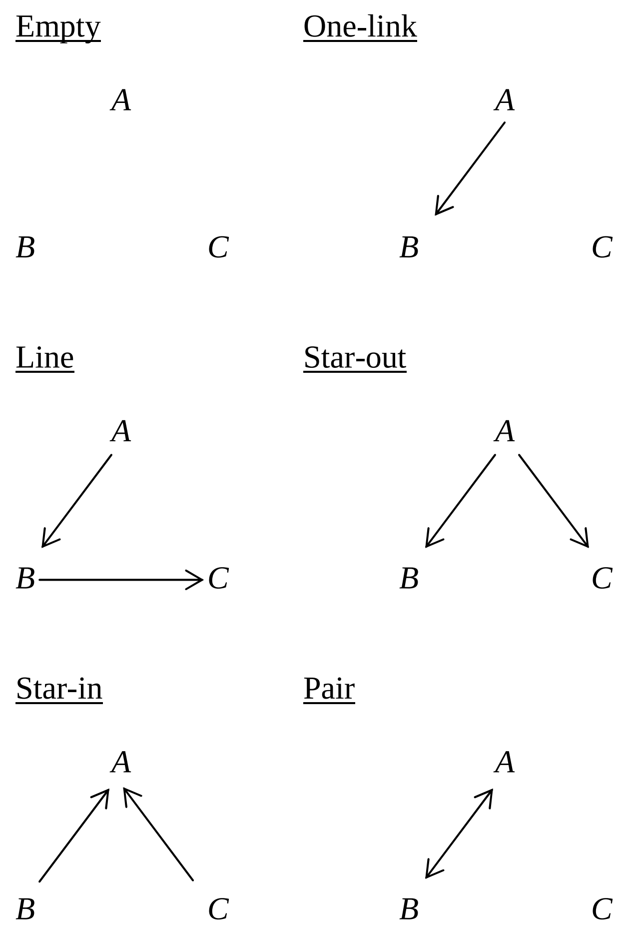

Figure 2. The frequencies of contributions across time for selected positions

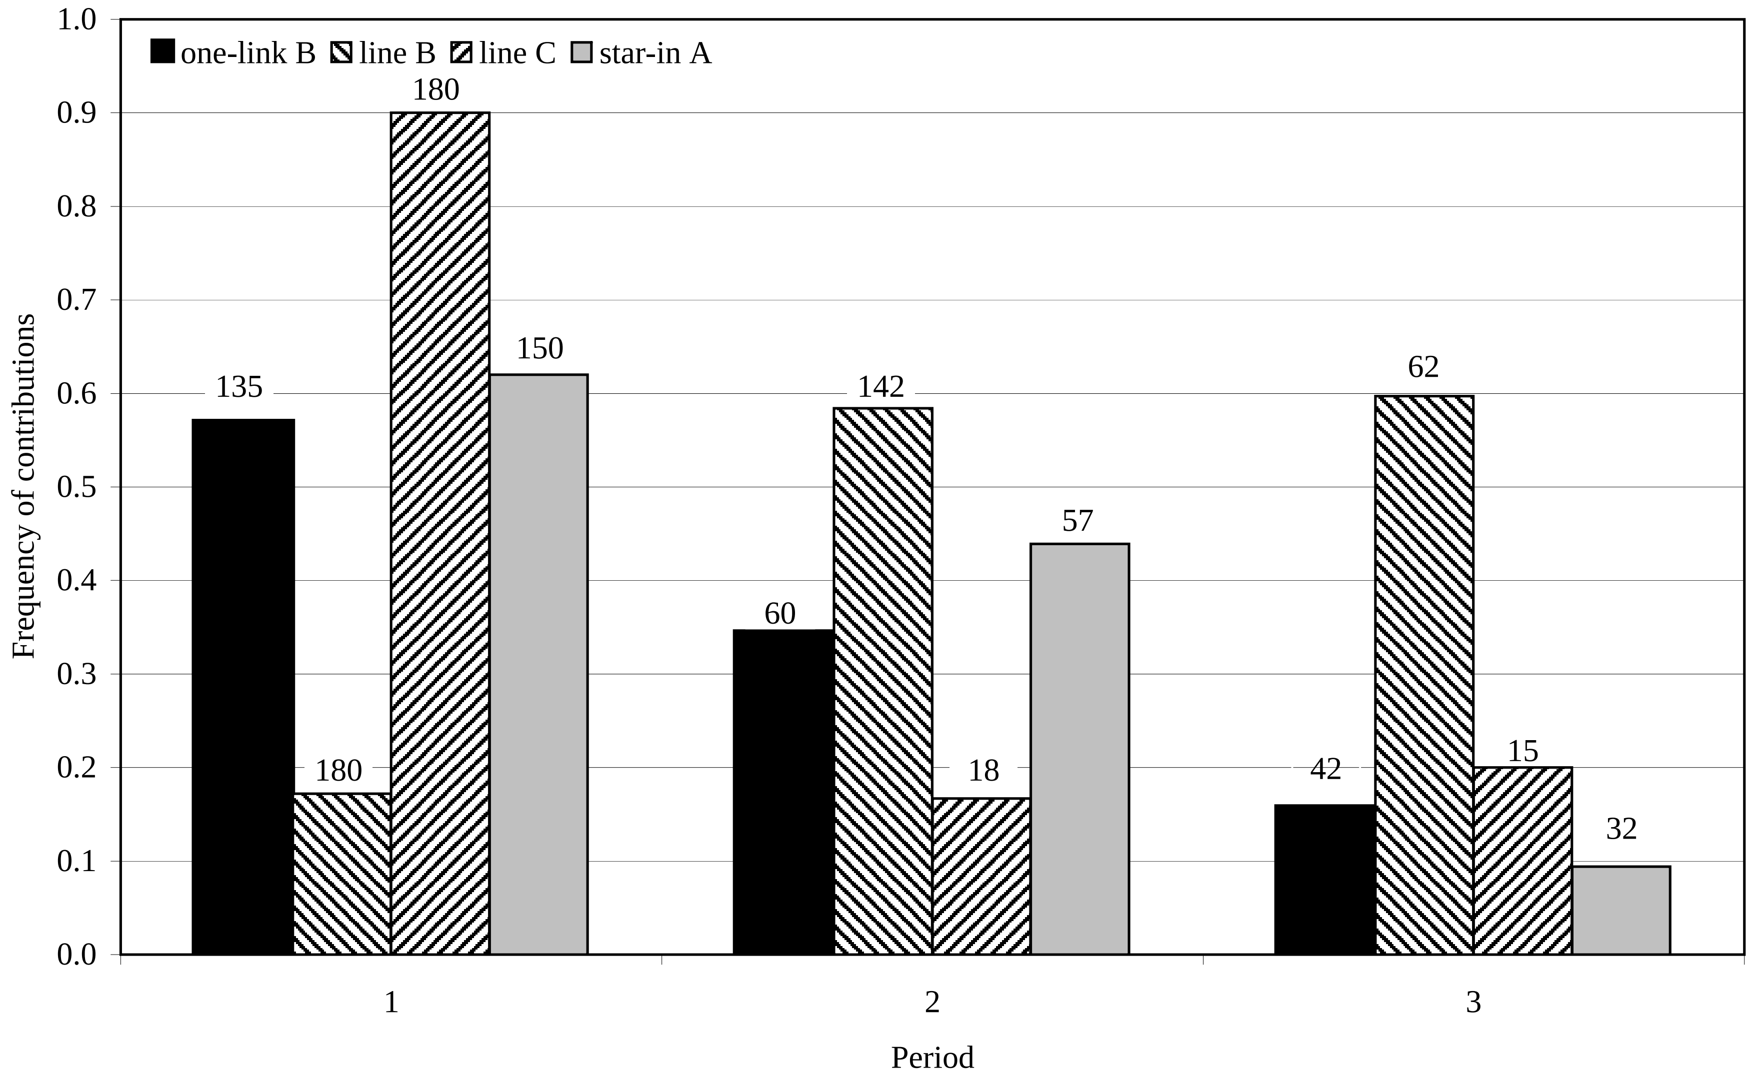

Figure 3. The frequencies of contributions at payoff-relevant states for selected positions

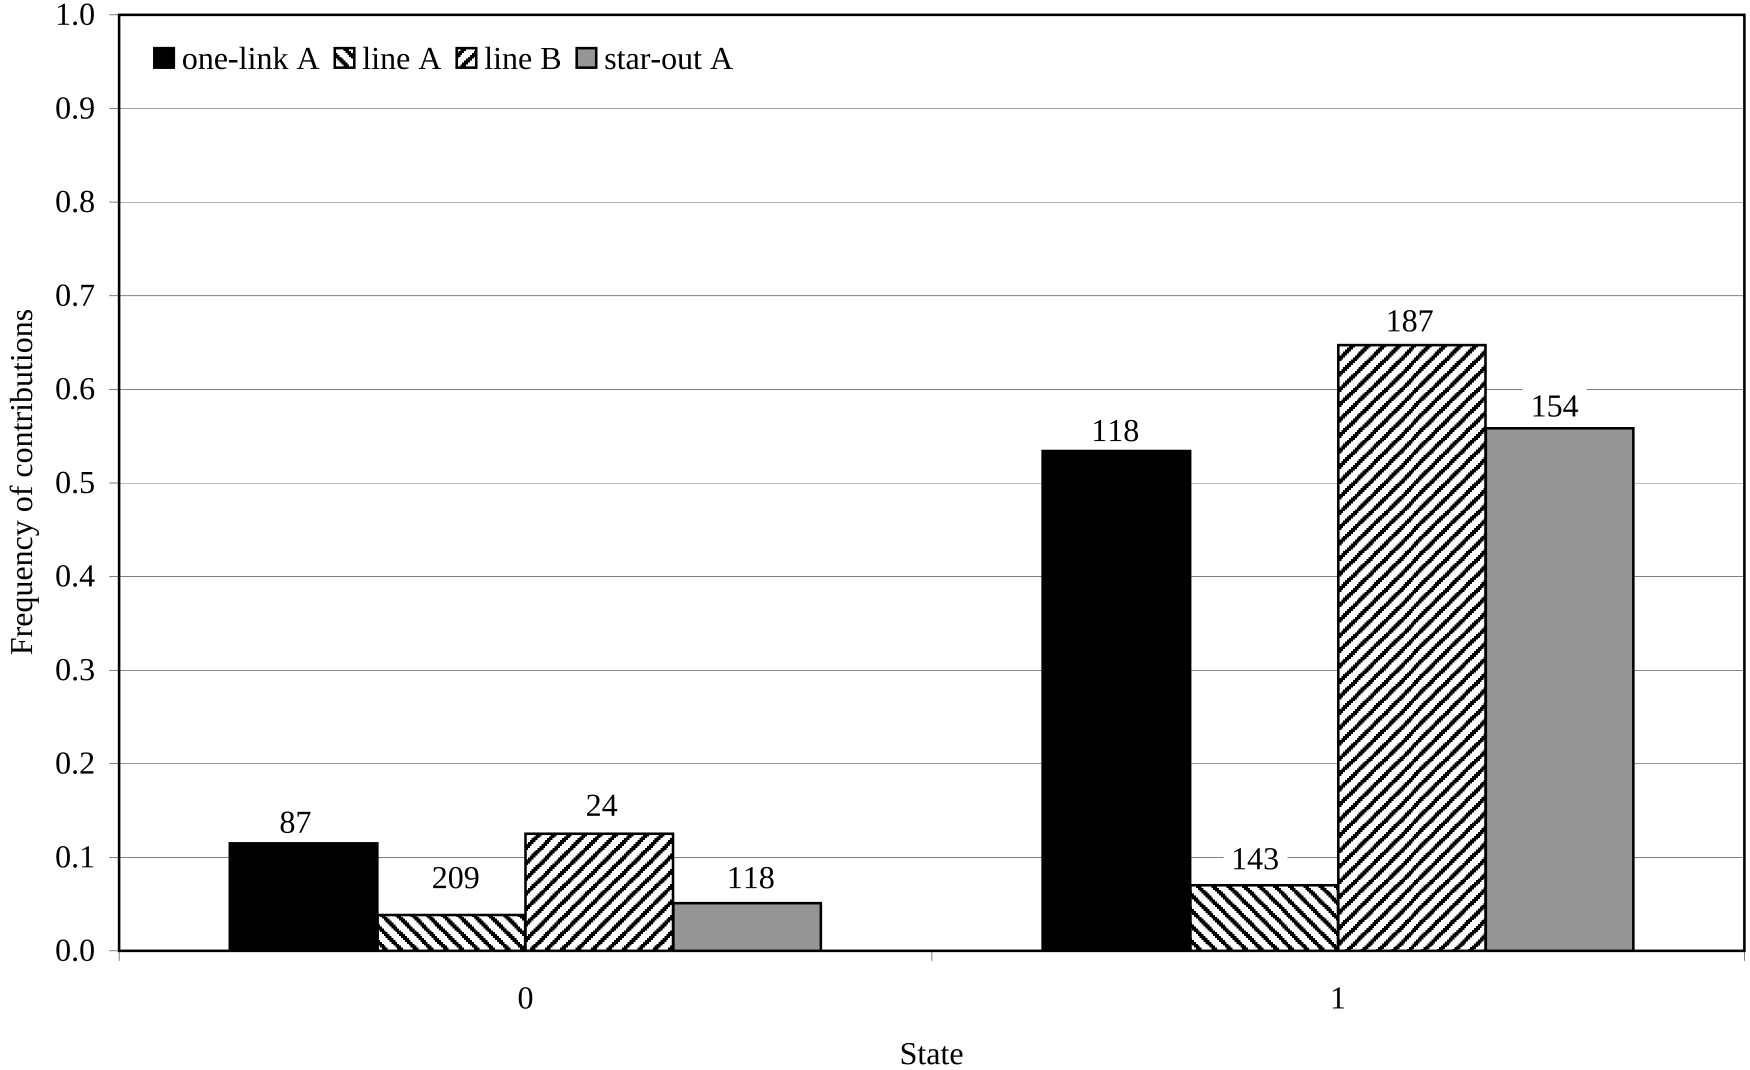

Figure 4. The total contributions across time in the star-out network by subjects in positions *B* and *C*

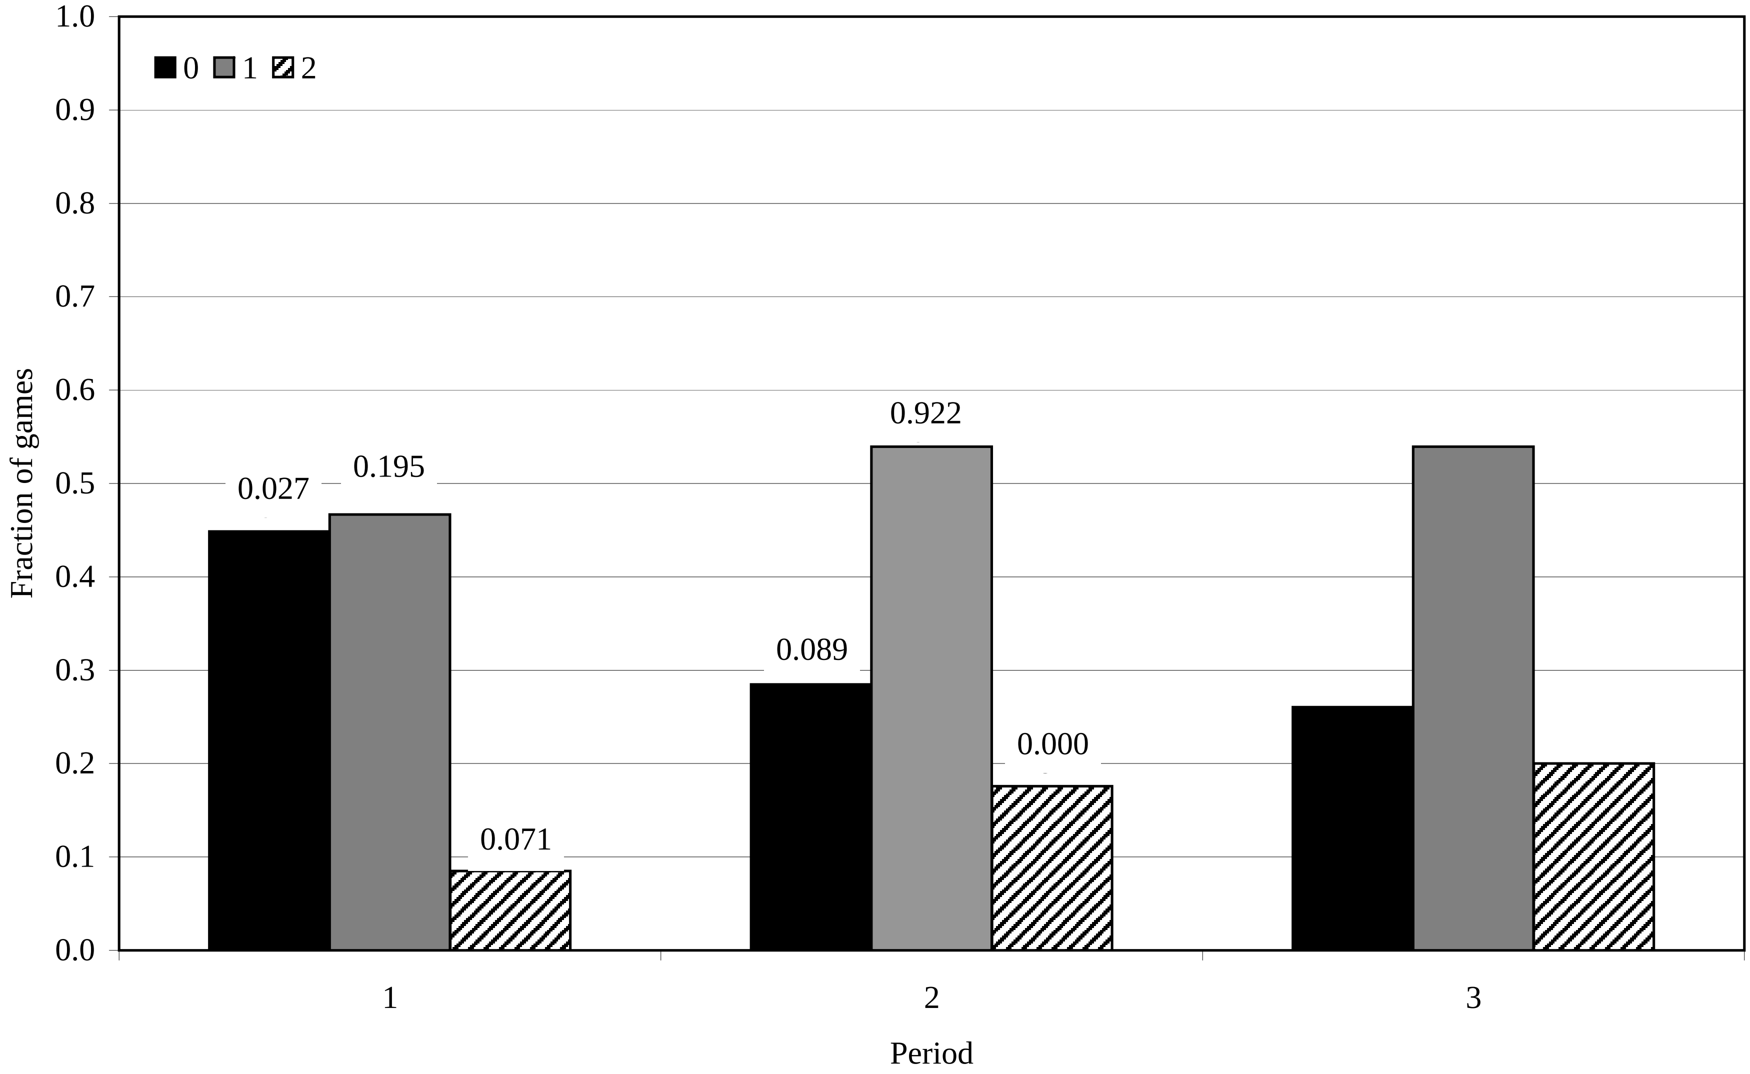

Figure 5. Efficiency in the star-in network conditional on the timing of contribution of position-*A* subjects