

Experiments with Network Formation*

Dean Corbae
Department of Economics
University of Texas at Austin
E-mail: corbae@eco.utexas.edu

John Duffy
Department of Economics
University of Pittsburgh
E-mail: jduffy@pitt.edu

First Draft: October 2000
This Draft: September 2003

Abstract

Our paper examines groups of I subjects who repeatedly play a two-stage game. In the first stage, the network structure in which individuals interact is either exogenously imposed or endogenously determined through a noncooperative game. In the second stage, subjects play τ rounds of a game against all of their ‘neighbors’ in the network. The second-stage game is a multi-player generalization of a much studied 2×2 coordination game in which there are two pure strategy equilibria, one of which is payoff dominant. One interpretation of the environment is that of a regional banking system with an endogenously determined interbank network, where withdrawal decisions taken by depositors in one region affects payoffs faced by depositors in other regions. We focus on a specific case for $I = 4$ and show that when players are free to choose their network structure, the only network which is strictly immune to unilateral deviations when agents have beliefs that play in the second stage will be according to the ex-ante, payoff dominant perfect bayesian equilibrium strategies is a marriage network. Experimental findings largely confirm the predictions of our model.

Key Words: Networks, Contagion, Coordination, Stability, Experimental Economics.

JEL Codes C7, C9.

*We wish to thank Andreas Blume, Ken Hendricks, Stephen Morris, Jack Ochs, Thomas Palfrey, Dale Stahl, Max Stinchcombe, and Tom Wiseman for many helpful comments. We also thank seminar participants at the Econometric Society meetings, the SED meetings, the Workshop on Coordination, Incomplete Information and Iterated Dominance at Pompeu Fabra, Cal-Tech, Carnegie-Mellon, Cornell, New York University, and the Federal Reserve Banks of Cleveland and Philadelphia. Drew Saunders provided brilliant research assistance. Finally we thank the National Science Foundation for support under grant SES-0095234.

1 Introduction

In this paper we attempt to address the broad question of why agents choose certain trading networks and how trading strategies spread through network economies. The answers have potentially important implications for equilibrium selection and the propagation of shocks. Since there are multiple equilibria of our repeated two-stage network formation game, as well as the fact that it is often hard to get good micro data on network formation and contagion, we choose to examine these issues in the context of experiments.

Allen and Gale’s [1] article on financial contagion is an excellent example of the type of economic environment that we are interested in. It is a network version of the Diamond and Dybvig [7] banking model, which can be interpreted in the context of a coordination game. There are four regions composed of ex-ante identical agents who receive unobservable preference shocks. The fractions of patient and impatient agents vary across the four regions in two possible states of the world. In one state, regions 1 and 3 have a high fraction of impatient agents while regions 2 and 4 have a low fraction. In the second state, the regions experiencing high demand for liquidity are reversed. Thus, there is no difference in the aggregate demand for liquidity across the two states and it is clear that there is a set of potentially beneficial trades across regions to provide insurance against region specific liquidity shocks. The set of potential trades, however, depends on the network structure that links this four region economy. Allen and Gale call an economy where all four regions are linked to each other by an interbank deposit market a complete network.¹ They also consider incomplete network structures: one network where all regions are linked but only through deposit markets from/to two immediate neighbors and one network where pairs of regions provide insurance to each other but there is a decoupling of the interbank markets.² Allen and Gale provide the conditions in incomplete network where the presence of a zero probability aggregate excess liquidity shock in one region can generate a financial contagion where there are insufficient funds to ward off the bank run in that region which triggers a string of early liquidations (and financial collapse) in the other regions. The authors leave the issue of network (or more specifically, interbank deposit market) design for future research.

Here we attempt to address network design and the possibility of contagion in the context of a relatively simple model economy that yields a set of equilibrium predictions which we test by means of experiments. The basic experimental design consists of a group of I subjects who repeatedly play a two-stage game. In the first stage, a network structure is exogenously or endogenously through a noncooperative “proposal” game played by the I players. The resulting network specifies all the players with whom each individual player in the group will interact – each player’s ‘neighbors’. In the second stage, which we call the “investment” game, each agent plays τ rounds of a general coordination game against all of his neighbors, earning a payoff in each round according to his single action choice and the equal weighted average choices of all his neighbors. We report and analyze the results of a number of experimental sessions, where the main treatment variable consists of the network structure exogenously imposed on subjects in the initial game. In particular, we explore the stability of these network structures in subsequent games, when the network structure is endogenously determined by the subjects themselves. For instance, if subjects start by playing the investment game with all other

¹Consistent with the literature, we will call this “uniform matching (UM)”.

²Consistent with the literature, we call the first type of network “local interaction (LI)” and the second type “marriage (M)”. Allen and Gale’s labelling of these network structures as “incomplete” is drawn more from the literature on incomplete markets rather than graph theory. In particular, in graph theory, LI is a complete network because all agents are either directly or indirectly linked while M is an incomplete network since the two pairs are not linked.

subjects (i.e. what we call uniform matching), do the subsequent proposal games result in a uniform matching network?

The paper is organized as follows. After reviewing the literature in Section 2, we describe a fairly general economic environment (matching and productive technologies, as well as preferences and information structure) in Section 3 and define a perfect bayesian equilibrium for our environment in Section 4. Section 5 provides theoretical results on the equilibrium set of networks for a simple 4 agent case under a specific parameterization of the model. The interesting case is when one of the 4 agents is randomly chosen to play the safe action (corresponding to the risk dominant equilibrium in a 2 agent setting) for each of the τ rounds of the investment stage. While it is common knowledge that one of the agents' action sets will be restricted in this way, who receives the shock is private information. We describe the ex-ante, payoff dominant perfect bayesian equilibrium for all possible network structures. Under the assumption that subjects play according to those strategies in the second stage (i.e. we endow agents with such beliefs in the first stage), we show in Proposition 2 that the only network which is strictly immune to unilateral deviations is marriage. This provides us with a stark, testable hypothesis for the experiments which we take up in Section 6. Since there are multiple equilibria in both the proposal and investment games, we analyze the data as to whether subjects are playing according to the strategies considered in Proposition 2. Our experimental findings are fairly consistent with the predictions of the theory.

2 Literature

The literature on network economies is voluminous. Here we will list an incomplete set of papers that influenced our analysis. The theoretical literature on network economies can be split into those that: (i) take the network as given and study equilibrium selection in a coordination game (we will refer to this as the exogenous networks literature) and (ii) allow the network to be chosen endogenously. Papers that follow the first approach are Ellison [8], Kandori, Mailath, and Rob [15], Morris [18], and Young [21]. Papers that follow the second approach are Bala and Goyal [2], Jackson and Watts [13], Jackson and Wolinsky [14] and Tesfatsion [20]; Jackson [12] surveys this literature. By contrast, there are relatively few experiments examining the impact of exogenous network configurations, or experiments where the network structure is endogenously chosen by the subjects themselves.

Experiments that consider the impact of exogenous network structures include Keser, Ehrhart and Berninghaus [16] and Berninghaus, Erhart and Keser [3] who examine how subjects play coordination games in various network structures under different payoff weighting schemes. Another study by Charness, Corominas-Bosch, and Frechette [5] looks at how bargaining behavior changes as links in exogenously imposed buyer-seller networks are altered.

More closely related to this study, there are two prior experimental studies examining *endogenous partner selection*. Hauk and Nagel [11] examine behavior in repeated 2-player prisoner dilemma games where players are either forced to interact in fixed pairs or where individual players may form unilateral or mutually agreed upon links with another player prior to playing the 2-player repeated game. They find that cooperation levels are higher when players can chose their partners than when players are forced to play with another player. Callander and Plott [4] examined how groups of 6 subjects played a network formation game based on the environment of Bala and Goyal [2]. Subjects could link to any of the other five subjects at a fixed cost per direct link, but their payoffs were increasing in the sum of their direct and indirect links. For the parameters they considered, a "wheel" (LI) network is both the uniquely efficient and strict Nash equilibrium, though there are many other weak Nash

equilibria. They report that convergence to Nash equilibrium configurations is a frequent occurrence and appears to be guided more by Nash stability considerations than by considerations of efficiency or the focalness of particular network structures (e.g. 1-2-3-4-5-6-1).

We build on this prior work in several ways. First, we provide our own theory of endogenous network formation in 4-player groups, an environment that admits all of the network configurations that have appeared in the theoretical literature (i.e. uniform matching, local interaction, marriage, stars, etc.). In particular, we are able to characterize whether each of the various possible endogenous network configurations that are admissible in our environment are equilibria or not, thus delivering crisp predictions which we then test in the laboratory. Second, unlike Bala and Goyal and Callander and Plott, in our model, link formation is *two-sided*, that is, links have to be mutually agreed upon between two parties in order to be implemented. Unlike Jackson and Wolinsky, however, we implement two-sided link formation in a noncooperative game. Third, we are not simply interested in the question of which networks emerge when agents are free to propose network links; we also examine how players play a coordination game with their network neighbors, similar to the games studied by Keser et al. [16] and Berninghaus et al. [3], given the network structure they have implemented. Thus our study can be viewed as unifying the two different experimental literatures on networks. Finally, in our environment unlike any other that has been previously explored, one player in every group receives a “payoff shock” that limits the actions he can choose in the coordination game. Using this device, we are able to explore the issue of the *contagious spread* of actions as a function of network structure. As noted in the introduction, such contagious behavior is an important consideration in the design of financial market networks, as well as in other applications.

3 The Environment

The basic model is of a finite sequence of two-stage games. One can interpret the economy as in Allen and Gale with insurance arrangements or interbank markets that link regions, though there are numerous other possible interpretations. In the first stage, agents choose the network structure endogenously through a simultaneous set of proposals. This case nests the literature with exogenous network structures since it is always possible to restrict the proposal action space to effectively impose any feasible graph. In the second stage, agents play several rounds of a game with their neighbors. In each round, they make an investment decision with payoffs similar to an n -person coordination game.

Specifically, there is a finite set of players $\mathcal{I} = \{1, 2, \dots, I\}$. There are a finite number of discrete periods $t = 1, 2, \dots, T$. There are κ repetitions of two stage play. We call the first phase the *network proposal stage* and the second phase the *investment stage*, where the latter has τ rounds of play. Thus, $T = \kappa(1 + \tau)$. Let t_p denote the times at which network proposal are made and t_a denote the times at which investment decisions are made.³

3.1 Matching technology

Economic interactions are determined by a matching technology $\mu^{\overline{ij}}(g_t)$, which assigns a weight to the link \overline{ij} between agents i and j in network g_t . While we define networks in much the same way as Jackson and Wolinsky [14], we implement the network as a noncooperative dynamic game in contrast to their static, cooperative solution concept.⁴ At the network proposal stage, agent i chooses whether

³Thus $t_p = 1 + k(1 + \tau)$, where $k = 0, 1, 2, \dots, \kappa - 1$ and $t_a = 1 + k(1 + \tau) + \varsigma$, for $\varsigma = 1, 2, \dots, \tau$ and any k .

⁴See Jackson and Watts [13] for a dynamic version of network formation.

or not to link to each agent in the economy. In particular, agent i takes a *network proposal action*: $p_t^i = (p_{1,t}^i, \dots, p_{j,t}^i, \dots, p_{N,t}^i) \in P^i = \{0, 1\}^I$ where $t = t_p$. The action $p_{j,t}^i = 1$ denotes a proposal by agent i to link to agent j , while $p_{j,t}^i = 0$ denotes i 's choice not to link to j at time t . Let the *network proposal action profile* at time t be denoted $p_t = (p_t^1, \dots, p_t^I)$. A *link* at time t , denoted \overline{ij}_t , occurs iff $p_{j,t}^i p_{i,t}^j = 1$. Thus, unlike Bala and Goyal [2], links must be mutually agreed upon. A *network* is just the set of all agreed upon links, $g_t = \{\overline{ij}_t : \forall i, j \in \mathcal{I}\} \in \Gamma$, where Γ is the set of all possible networks. We assume that the network remains constant during the τ rounds of play of the investment stage.⁵ We treat the network as an endogenous state variable.

Network formation can be a costly activity. If there are costs of sending proposals, then

$$c^i(p_t^i) = \begin{cases} c & \text{for each } p_{j,t}^i = 1, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

while if the costs are for maintaining links, then

$$c^i(p_t) = \begin{cases} c & \text{for each } p_{j,t}^i p_{i,t}^j = 1, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

We define the *neighborhood* of agent i to be the set of all agents to whom he/she is linked and denote it $N^i(g_t) = \{j | \overline{ij}_t \in g_t, j \neq i\}$. Let the set of all possible neighborhoods of any given agent be denoted \mathcal{N} . The number of neighbors of agent i is simply the cardinality of $N^i(g_t)$ and is denoted $n^i(g_t)$. Given these definitions, the matching weights are given by:

$$\mu^{\overline{ij}}(g_t) = \begin{cases} \frac{1}{n^i(g_t)} & \text{if } j \in N^i(g_t), \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Thus, we assume that the interactions are *bounded* (i.e. that the weights add up to one) and *deterministic*.⁶ Examples of the types of interaction structures which fall within this class are the uniform matching rule (UM) of Kandori, Mailath, and Rob [15] or Young [21] and the local interaction rule (LI) of Ellison [8]. In these examples, $I - 1 \geq n^i(g_t) \geq 2$. The class of interaction structures we allow includes degenerate links (e.g. monogamous marriages ($n^i(g_t) = 1$) or polygamous marriages (a star)) which have also appeared in the literature (see, e.g. Jackson and Watts [13]).

To illustrate the types of interaction structures we study, let $I = 4$, which is the smallest number that admits both UM and LI as distinct, special cases. Our experiments will consider values of $I = 4$.

The set of all feasible graphs is shown in Figure 1.⁷ The only complete graph is UM , a 4-agent version of the *uniform matching* mode. In this case, each agent has $n^i(g^{UM}) = 3$ direct links to all other agents in the economy. All other cases are, in Allen and Gale's terminology, *incomplete*. The graph LI is the 4-agent version of a *local interaction* model. Each agent has $n^i(g^{LI}) = 2$ direct links (and 2 indirect links) to every agent in the economy. Graph M represents a *marriage* model where each agent has $n^i(g^M) = 1$ direct link and no indirect links. The final symmetric graph, A shown on the top row is the case of no links or *autarchy*. The second row of graphs in Figure 1 are hybrids of the symmetric graph forms that result from removing a single link from the symmetric graph shown

⁵That is, for all $\varsigma = 1, 2, \dots, \tau$ and any k , $g_{1+k(1+\tau)+\varsigma} = g_{1+k(1+\tau)}$.

⁶See Morris [18] for this taxonomy.

⁷In Figure 1, there are many other graphs that are *isomorphic* to the ones we present. For instance, in *LI* we choose only to illustrate the "square" form of LI rather than the "bow tie" or "hour glass" versions of LI.




Figure 1: Illustration of all symmetric and asymmetric graph forms when $I = 4$

just above, in the first row. The fourth graph, $UM - LI$ has agents 1 and 3 in LI neighborhoods while agents 2 and 4 are in UM neighborhoods. The fifth graph, $LI - M$, has agents 1 and 4 in LI neighborhoods and agents 2 and 3 in M neighborhoods. Other hybrid network possibilities are shown in the third row of Figure 1. The graph $UM - M$ is sometimes referred to in the literature as a *star* network or the case of a single *middleman*; as depicted in Figure 1, the center of the star or middleman is agent 3. The graphs $M - A$, $LI - A$ and $LI - M - A$ all entail some form of ostracism, where one or more players are unlinked.

3.2 Payoffs

In each of the τ rounds of the second stage investment game, players are endowed with one unit of indivisible capital. There are two types of investment (or storage) technologies. One is risky (R) and the other is safe (S). In the context of the banking example, the risky technology is the long-run technology which may be liquidated at a cost due to a bank run. The action set for each agent depends on an idiosyncratic shock $\omega_t^i \in \Omega^i = \{0, 1\}$ so that $a_t^i \in A^i(\omega_t^i)$ where $A^i(1) = \{S\}$ and $A^i(0) = \{R, S\}$. The shocks provide the experimenter with some control over the investment decisions that agents take in the economy at times $t = t_a$. For instance, the shock can be used to randomly assign one agent in the economy a tremble. In the context of the banking example, this allows us to enforce early liquidation by a strict subset of the population. In general we let $\pi^i(\omega_t^i; \omega_{t-1}^i) = \Pr(\omega_t^i | \omega_{t-1}^i)$ where $\omega_{t_p} = \emptyset$ and denote the shock profile by $\omega_t = (\omega_t^1, \dots, \omega_t^I) \in \Omega^I$. We will focus on two specific processes

that do not induce aggregate uncertainty despite a finite population. In the *transitory case*, in each round ($t = t_a$) of the second stage, one agent (say i) in \mathcal{I} receives $\omega_t^i = 1$ while all other agents receive $\omega_t^{-i} = 0$. Which agent in \mathcal{I} receives the “tremble” is drawn from a uniform distribution so that $\pi^i(\omega_t^i; \omega_{t-1}^i) = 1/I, \forall t_a$. This process can be thought of as a finite population approximation to an i.i.d. idiosyncratic technology shock where the population proportion receiving the adverse shock is $1/I$. In the second instance, which we call the *permanent case*, in the first round of the second stage, one agent (say i) in \mathcal{I} receives $\omega_t^i = 1$ while all other agents receive $\omega_t^{-i} = 0$ as before. However, in this case, agents maintain their type in all τ rounds of the second stage. That is, $\pi^i(\omega_t^i = 1; \emptyset) = 1/I$ and $\pi^i(\omega_t^i = 1; \omega_{t-1}^i = 1) = 1$ while for all other agents $\pi^{-i}(\omega_t^{-i} = 0; \omega_{t-1}^{-i} = 0) = 1$ for the remainder of the rounds $\varsigma = 2, \dots, \tau$.

Before agent i interacts with his investment partners (taking as given the network g_t from the first stage) and after he learns the state of his technology (ω_t^i), he takes an *investment action* which is implemented in all of player i 's interactions with other players in his neighborhood $j \in N^i(g_t)$. That is, actions cannot be made j contingent; otherwise, network structure would be inconsequential. Agent type contingencies of this variety were the object of the first stage network proposal game. Let the *investment action profile* at time t be denoted $a_t = (a_t^1, \dots, a_t^I)$.

We assume agent i 's payoffs for this investment activity, denoted $u^i(a_t^i, a_t^j), j \in N^i(g_t)$, are given by:

$$\begin{aligned} u^i(R, R) &= a & u^i(R, S) &= c \\ u^i(S, R) &= b & u^i(S, S) &= b \\ u^i(a^i, \emptyset) &= d & \text{if } N^i(g_t) &= \emptyset \end{aligned} \tag{4}$$

The last line of (4) simply says that if an agent is in autarky, he receives payoff d independent of his actions. If there were only 2 players, the payoffs are simply those of a pure coordination game with two pure strategy equilibria (*all-R* and *all-S*) and a mixed strategy. Furthermore, while *all-R* is payoff dominant (for players receiving the favorable payoff shock), if $b > (a + c)/2$ then *all-S* is the equilibrium with lowest *risk factor*.⁸ Agent i 's payoff in any round of the investment game from playing action a^i is then given as the weighted sum of all the payoffs associated with actions taken by his neighbors:

$$\sum_{j \in N^i} \mu^{ij} u^i(a^i, a^j). \tag{5}$$

We chose this weighted sum representation of payoffs rather than a simple aggregation since we did not want link formation to be simplistically increasing in the number of links, biasing comparisons across network structures towards large neighborhoods. For instance, under a simple aggregation rule where all neighbors are playing R , local interaction would be preferred to marriage simply because it yields $2a$ rather than a .

3.3 Information and the Timing of Events

Since each agent interacts with every other agent in his neighborhood, we assume that in any period $t > 1$, agents know the actions that have been played by their neighbors in all previous rounds, through $t - 1$. Consistent with the decentralized nature of agent interaction, we assume that players

⁸See Young [22] (p. 67) for a definition of risk factor. If the game were a perfectly symmetric 2×2 game, all-S would be more familiar as the *risk dominant* equilibrium.

do not know the actions that have been played by agents outside of their neighborhood. To be fully consistent with the decentralized nature of interactions, we would need to assume that agents do not know the network structure outside of their neighborhood (i.e. they only see $N^i(g_t)$ but not g_t). In our experiments, however, we provide agents with information about the actual network (i.e. g_t).

While the distribution of technology shocks is common knowledge, we assume the idiosyncratic shock is private information. That is, while agent i knows his own type ω^i he does not necessarily know the types of others ω^{-i} . This assumption is consistent with the information partition associated with a given network structure.⁹ In the context of the banking example, agents know the fraction of early consumers but not who are the early consumers. This assumption is also consistent with work on incomplete information games by Morris [17].

Figure 2 summarizes the timing of events.




Figure 2: The timing of events in each session

4 Equilibrium

In this section, we define strategies and an equilibrium of our dynamic game. We begin by defining an agent's history. In any given period, an agent is assumed to know his payoff shock (if $t = t_a$), his actions (both his proposals if $t = t_p$ and investments, as well as the investment actions of his neighbors, if $t = t_a$), as well as the realized neighborhood he is in. Thus we let $h_t^i = (\omega_t^i, p_{t-1}^i, N^i(g_{t-1}), a_{t-1}^i, a_{t-1}^{j \in N^i(g_t)}, h_{t-1}^i)$ denote agent i 's *history* in period t , where $h_1 = \emptyset$ and $\omega_{t_p}^i = \emptyset$.¹⁰ We assume that information which does not change across certain periods of time is simply updated with the identity function (e.g. for all $\varsigma = 1, 2, \dots, \tau$, $p_{1+k(1+\tau)+\varsigma}^i = p_{1+k(1+\tau)}^i$). Let the set of all possible histories for agent i at time t be denoted $H_t^i = \Omega \times P^i \times \mathcal{N} \times A^i \times \left(\times_{j \in N^i(g_t)} A^j \right) \times H_{t-1}^i$. In the case where we allow agents to know the network, then g_t is substituted for $N^i(g_t)$ in the definition of i 's history and Γ is substituted for \mathcal{N} in the definition of the set of all possible histories.

A *behavior strategy* of agent i is a history contingent plan of proposal and investment actions which we denote $\sigma^i \in \Sigma^i = \left(\times_{h_t^i \in H_t^i} \Delta(P^i) \right) \times \left(\times_{h_t^i \in H_t^i} \Delta(A^i) \right)$ where $\Delta(X)$ denotes a probability distribution defined over the set X .¹¹ Let $\sigma^i(h_t^i) = (\rho^i(h_t^i), \alpha^i(h_t^i))$. Let the *behavior strategy profile* be denoted

⁹Of course, if i happens to be the agent who experiences the adverse state (say $\omega^i = 0$), then he knows all other agents are in the favorable, high payoff state ($\omega^{-i} = 1$). Otherwise he only knows that some other agent, quite possibly outside of his neighborhood, experiences the adverse shock.

¹⁰ p_t^i is included since knowing either $N^i(g_t)$ (or g_t) is not sufficient to infer $p_{i,t}^j$ (e.g. if i sends $p_{j,t}^i = 0$, he can't infer $p_{i,t}^j$ from g_t).

¹¹This follows the definition in Fudenberg and Tirole [9], pp. 84-5.

σ . Conditional on his own history, agent i forms *beliefs* about his neighbors' histories (including his neighbor's type) which we denote $\beta^i(h_t^j; h_t^i)$. Let the *belief profile* be denoted $\beta = (\beta^1, \dots, \beta^I)$.

Following these definitions, we are able to write expected payoffs for agent i for t in the investment stage, given h_t^i as:

$$\begin{aligned} v^i(\alpha^i, \alpha^{j \in N^i(g_t)}; h_t^i, \beta^i) &= \sum_{j \in N^i(g_t)} \sum_{h_t^j \in H_t^j} \sum_{a_t^i \in A^i, a_t^j \in A^j} \beta^i(h_t^j; h_t^i) \bar{\mu}^{ij}(g_t) u^i(a_t^i, a_t^j; \omega_i) \alpha^i(a_t^i) \alpha^j(a_t^j), \quad (6) \\ &+ \sum_{h_{t+1}^i \in H_{t+1}^i} q(h_{t+1}^i; h_t^i, a_t^i, a_t^{j \in N^i(g_t)}) v^i(\sigma^i, \sigma^{j \in N^i(g_{t+1})} \vee \mathcal{I}; h_{t+1}^i, \beta^i), \end{aligned}$$

where $q(h_{t+1}^i(\omega_{t+1}^i, p_t^i, g_t(p_t), \cdot, \cdot, h_t^i); h_t^i, a_t^i, a_t^{j \in N^i(g_t)})$ is a transition function for agent i giving the probability that next period's history is h_{t+1}^i conditional on the current history and actions and obviously includes the transition probabilities for the exogenous shock process ω_{t+1}^i .¹²

Given (h_t^i, β^i) , $\hat{\alpha}^i$ is a *best response* in the investment stage if

$$v^i(\hat{\alpha}^i, \alpha^{j \in N^i(g_t)}; h_t^i, \beta^i) \geq v^i(\alpha^i, \alpha^{j \in N^i(g_t)}; h_t^i, \beta^i), \forall \alpha^i.$$

Given (h_t^i, β^i) , let the set of all best responses to $\alpha^{j \in N^i(g_t)}$ be denoted $BR^i(\alpha^{j \in N^i(g_t)}; h_t^i, \beta^i)$.

Furthermore, expected payoffs for agent i for t in the proposal stage, given history h_t^i can be written

$$\begin{aligned} v^i(\rho^i, \rho^{-i}; h_t^i, \beta^i) &= - \left[\sum_{j(\neq i) \in \mathcal{N}} \sum_{h_t^j \in H_t^j} \sum_{p_{jt}^i, p_{it}^j \in \{0,1\}} \beta^i(h_t^j; h_t^i) c \rho^i(p_{jt}^i) \rho^j(p_{it}^j) \right] + \quad (7) \\ &\sum_{h_{t+1}^i \in H_{t+1}^i} q(h_{t+1}^i; h_t^i, \cdot, \cdot) v^i(\sigma^i, \sigma^j; h_{t+1}^i, \beta^i). \end{aligned}$$

Given (h_t^i, β^i) , $\tilde{\rho}^i$ is a *best response* if

$$v^i(\tilde{\rho}^i, \rho^{-i}; h_t^i, \beta^i) \geq v^i(\rho^i, \rho^{-i}; h_t^i, \beta^i), \forall \rho^i, h_t^i.$$

Given β^i , let the set of all best responses to ρ^{-i} be denoted $BR^i(\rho^i; \beta^i)$.

We are now ready to define an equilibrium.¹³

Definition 1 A *perfect Bayesian equilibrium* is a strategy-belief pair $(\hat{\sigma}, \beta)$ such that (i), given β^i , $\hat{\sigma}^i \in BR^i(\hat{\sigma}^{-i})$, $\forall i, t, h_t^i$ and (ii), wherever possible, posteriors satisfy $\beta^i(h^j; h^i) = \text{prob}(h^j | h^i)$.

5 Predictions

The next set of results is intended to provide us with insight into a set of predictions for our experimental study. The results are for a very simple $I = 4$ player game; there is one network proposal

¹²This definition of transition functions follows Fudenberg and Tirole [9], p. 503.

¹³This definition follows Fudenberg and Tirole [9], p. 333.

round followed by τ rounds of the investment game so that $T = \tau + 1$.¹⁴ We consider only two polar cases for the shock process discussed in Section 3.2: transitory and permanent shocks.

The logic of our analysis is to first characterize continuation equilibria of the investment game taking as given the network structure and then determine whether any agent would want to deviate from the given network in the proposal stage. Specifically, while there are many equilibria of the investment game (e.g. all- S), we characterize the ex-ante payoff dominant, pure strategy, symmetric Perfect Bayesian equilibrium (PBE) of the investment game, taking as given the network structure.¹⁵ In many cases, the equilibrium strategies are such that players who do not receive the shock play R in the first round. This strategy reveals which agent was shocked in one round in many network structures and we characterize mutual best responses in subsequent rounds after agents update their beliefs across the different network structures. We define a continuation equilibrium to be *ex-ante pareto efficient* if there is no other symmetric equilibrium of the investment game in a given network where some agent in that network can expect to receive a higher payoff in her network position before the occurrence of the technology shock and everyone else can expect to receive at least as much in their position. An ex-ante pareto efficient equilibrium is said to be *ex-ante payoff dominant for a given (M, LI, UM) type* if there is no other ex-ante pareto efficient equilibrium that gives that type a higher payoff. We say an equilibrium is *ex-ante payoff dominant* if it is ex-ante payoff dominant for all types.¹⁶ We then use these results to construct LI and M equilibria under the assumption that agents coordinate upon payoff dominant equilibria in each investment game continuation. Furthermore, we show that UM is not an equilibrium under this assumption. In particular, given our assumption in the environment that links are only formed by mutual agreement, we show that there are not deviations from the minimal set of proposals required to construct such networks which make any agent better off in the case of LI and M, while there is a deviation in the case of UM. There are, however, deviations from LI which make the agent as well off as he was in LI. In fact, we characterize the entire equilibrium set of networks.

Unless stated explicitly, we make the following parametric assumption.

Assumption 1 $a > \frac{2a+c}{3} > b > \frac{a+c}{2} > c$ and $b \geq d$.

The assumption that $a > b > c$ is standard in coordination games where coordinated risky (R) play yields a higher payoff than safe (S) play. The assumption that $\frac{2a+c}{3} > b$ ensures that in neighborhoods of three players, coordinated R play yields a higher payoff than S despite the fact that the shocked player is in one's neighborhood. The assumption that $b > \frac{a+c}{2}$ ensures that if a shocked player is in one's neighborhood of two players, R play is suboptimal. This assumption is necessary for contagion

¹⁴The results on existence of equilibrium can be extended to multiple rounds of proposals, simply by constructing an equilibrium where agents disregard history of the prior proposal and investment games. On the other hand, such equilibria are not necessarily ex-ante payoff dominant.

¹⁵By symmetry we mean that if two players have the same number of neighbors and experience the same history of actions, then they take the same actions. The sense in which we use the term symmetry for proposals is that since each agent starts with the null history, they send out the same number of proposals. So, for instance, symmetric proposal strategies for a 4-player LI game means that each player sends out two proposals $p_{i-1,0}^i = p_{i+1,0}^i = 1$ and $p_{i+2}^i = 0$, mod 4.

¹⁶The set of pareto efficient equilibria is non-empty for any given network. For each type, there is always a payoff dominant equilibrium for that type. It is not always the case that the payoff dominant equilibrium of a given type is also the payoff dominant equilibrium for another type. In fact there are only two cases of 11 we analyse in Lemmas 10 and 11 where this is the case.

to get started. It is also consistent with all- S play being risk dominant in a two player game.¹⁷ The assumption that $b \geq d$ ensures that participation is weakly optimal.

5.1 Transitory Shocks

The next result establishes that with transitory shocks, network structure does not matter (all proofs appear in the appendix). The intuition for the result is that with i.i.d. shocks, there is nothing learned about the shock in one period that can be used to infer who will be playing S next period. In that case, the optimal ex-post investment strategy is the one that maximizes ex-ante payoffs. Since all realizations are equally likely, network structure doesn't matter.

Proposition 1 *If the strategy in which each unshocked player chooses R in each period is a PBE for some network structure, then it is an equilibrium for any network structure when shocks are transitory.*

5.2 Permanent Shocks

Now we move onto the polar opposite case of permanent shocks. In the appendix, we analyze this case in two steps. Lemmas 1 to 11 establish certain results of the investment game under a given network structure. The only real issue for players in the investment game is to infer who has the shock and best respond accordingly. The results are useful to describe equilibrium play for the symmetric networks we will be examining (i.e. UM, LI, and M) as well as to define equilibrium strategies for asymmetric network structures that may result through a unilateral deviation from UM and LI. These latter results will in turn be used in establishing predictions for play in the proposal game that determines equilibrium networks in a series of Lemmas 12 to 20 in the Appendix. These lemmas are summarized in the text by Proposition 2. The idea is to endow agents with beliefs that play in a network which results from a unilateral deviation from a given network will follow the ex-ante payoff dominant perfect bayesian continuation equilibrium strategies discussed in the above lemmas in the subgame following that deviation.¹⁸ We illustrate the possible unilateral deviations and resulting networks in Figure 3, for all networks illustrated in Figure 1. In Figure 3, the arrows from a given network to a new network show the result of a single, unilateral deviation. In certain cases, e.g. UM-LI-M, a single deviation can result in several new and distinct network structures.

[Insert Figure 3 here.]

The first result is that an M network is a strict PBE in the sense that a unilateral deviation leaves the player strictly worse off.¹⁹ In particular, a unilateral deviation from sending a proposal to one's partner results in autarky, where payoff $d\tau$ is strictly less under Assumption 1 than the ex-ante payoff associated with M given by $\frac{2a+c}{3} + \frac{2a+b}{3}(\tau - 1)$. This ex-ante payoff is calculated using the ex-ante payoff dominant, perfect bayesian equilibrium strategy given by all unshocked agents playing R in the first round, players who have played R in each previous round continue to do so if their partner has

¹⁷A pair of strategies is risk dominant (Harsanyi and Selten [10]) if each strategy is a best response to a mixed strategy of the other player that weights all the player's pure strategies equally. In our case, since choice of R yields payoff $\frac{1}{2}a + \frac{1}{2}c$ while choice of S yields b , then we have S being the risk dominant action.

¹⁸There is always an issue about coordination of proposals, which we try to address in the experiments by actually making participants play τ rounds of the investment game under a given network structure before sending their proposals in a subsequent stage.

¹⁹This result is established in lemma 12 in the appendix.

played R in each previous round, and any player who has herself played S or whose partner has played S in some round plays S in each subsequent round.²⁰ Since agents are ex-ante more likely to be in an unshocked marriage and S play invokes an S response according to the subgame perfect strategy, in the first round it is optimal to play R until one knows whether one’s partner is the shocked player, in which case it is optimal to play S since $b > c$.

The second result establishes that LI is a weak PBE in the sense that there is no strictly profitable unilateral deviation that brings about LI-M or LI-M-A.²¹ However, the PBE is weak in the sense that a unilateral deviation leaves the player indifferent. To understand the result, suppose agent 2 deviates and chooses not to send a proposal to agent 3, while all other agents send two proposals associated with the original LI network. This deviation is illustrated in the first column, third and fourth rows of Figure 3. The equilibrium play in LI-M is identical to equilibrium play in LI since the M player, if he is unshocked, knows that one of the two LI players is linked to a shocked player after the first round, thereby altering his beliefs and best responding with S play in the subsequent rounds as dictated by the ex-ante payoff dominant PBE strategy in an LI network.²² This strategy (where unshocked agents in an LI network play R in the first round until the shocked agent is discovered is an equilibrium (though obviously not unique) follows from $b > \frac{a+c}{2}$. Using that strategy, the position of the shocked player can be inferred after one round. The agent who is diagonally across from the shocked agent anticipates that his unshocked neighbors will play S and hence plays S . Notice that this result would be very different if we were using a solution concept like naive best response. Thus, the “contagious” S play spreads very quickly in our application, but would take another round with naive players. Since equilibrium play is the same in LI-M and LI, ex-ante payoffs are identical so that the deviation is not strictly profitable. There is an important sense, however, in which LI is not stable which corresponds informally to an evolutionary stability type argument. That is, a best response to agent 2’s single proposal to agent 1 is for agent 1 to send a single proposal to agent 2. As above, agent 2 does no worse sending one proposal and both do better getting into a marriage. This type of proposal strategy would displace LI as an equilibrium.

The third result shows that a UM network is not stable in the sense that there is a strictly profitable unilateral deviation that brings about UM-LI.²³ This result is despite the fact that the ex-ante payoff dominant, pure strategy PBE in UM results in each unshocked agent playing R in every round so that this network is “contagion-proof”. To understand the result, suppose agent 1 deviates and chooses not to send a proposal to agent 3, while all other agents send proposals to all other agents. The resulting UM-LI network (see the first two rows of Figure 3) means that agent 1’s two neighbors (agents 2 and 4) “provide insurance” to agent 1 (continue to play R) in the event that agent 3 gets the shock. In that event, agent 1 receives payoff a while in the UM network he would receive $(2a+c)/3 < a$.²⁴ Thus, the resulting instability of the UM network is similar to a free-rider problem. That is, each agent has an incentive to enjoy the benefits of insurance against payoff shocks (the public good) provided by others while providing it insufficiently herself.

²⁰This result is established in lemma 6 in the appendix.

²¹This result is established in lemma 13 in the appendix.

²²These results are established in lemmas 3 and 8 in the appendix.

²³See lemma 15.

²⁴This payoff is consistent with the ex-ante payoff dominant, perfect Bayesian equilibrium strategy in UM where each unshocked agent plays R in the first round and thereafter plays R in each round in which at least three agents play R in the previous round, and plays S otherwise. That this is optimal follows since $\frac{2a+c}{3} > b$. See lemma 2 for play in UM and lemma 4 for play in UM-LI.

There are other related results that pertain to asymmetric networks that are variants of UM, LI, or M. For instance, a star (UM-M) network is not stable in the sense that the UM player could unilaterally deviate and send only one proposal, resulting in his own marriage. His ex-ante payoffs $\frac{2a+b+c}{4} + \frac{a+b}{2}\tau$ from being in a marriage are strictly higher than the expected payoffs by being the middleman $\frac{2a+b+c}{4}\tau$ since in the event that he is unshocked, he provides insurance against the shock with probability one each period.²⁵ That it is optimal for him to provide such insurance if he is unshocked follows since the ex-ante payoff dominant, perfect Bayesian equilibrium strategy is similar to that of the UM network discussed above.²⁶ We summarize all the results in the following proposition.

Proposition 2 *When shocks are permanent, T is sufficiently large, and we restrict play in the investment continuation game to satisfy ex-ante payoff dominance for at least one type, the set of weak PBE networks are LI, M, LI-A, M-A. Those which are strict PBE networks are M, M-A, and LI-A.*

Finally, we discuss the sensitivity of our results to the size $I = 4$ of the economy with permanent shocks to only one agent in the economy. We chose $I = 4$ since it was the smallest number of agents that allowed us to differentiate between symmetric M, LI, and UM networks. It also kept the number of asymmetric networks to a manageable level (for example, if $I = 6$, then there are 8 symmetric networks and Y asymmetric networks). It can be shown that under Assumption 1, the strategies that result in the ex-ante payoff dominate equilibrium in the UM and M networks with $I > 4$ are the same as those for $I = 4$ analyzed previously. In that case the ex-ante payoffs V_M to a given agent of being in an M network are

$$V_M(I) = \left(\frac{1}{I}\right)[b\tau] + \left(\frac{1}{I}\right)[c + b(\tau - 1)] + \left(\frac{I-2}{I}\right)[a\tau]$$

where the first term is if the agent is shocked, the second if his partner is shocked, and the third if someone else is shocked, while the ex-ante payoffs V_{UM} to a given agent of being in a UM network are

$$V_{UM}(I) = \left(\frac{1}{I}\right)[b\tau] + \left(\frac{I-1}{I}\right)\left[\frac{(I-2)a+c}{(I-1)}\right]\tau$$

where the first term is if the agent is shocked, the second if he is not. For any finite I , it is simple to see that an M network strictly dominates a UM network ex-ante (i.e. $V_M(I) - V_{UM}(I) \propto (b-c) > 0$). It can also be shown that the incentive to unilaterally deviate from UM holds because of the externality in the previous results and it is clear that deviating to A from M is suboptimal. The main difference from the previous results when $I = 4$ is in the LI network. In this case, the strategy that resulted in the ex-ante payoff dominant equilibrium in Lemma 3, generalizes as follows: each unshocked agent plays R in the first round, and then plays R either until one of his neighbors has played S , or until he can infer that one of his neighbors will play S in the current round, and he plays S thereafter. For I even with $I/2 > \tau$, the ex-ante payoffs V_{LI} to a given agent of being in an LI network following this strategy are

$$\begin{aligned} V_{LI}(I) = & \left(\frac{1}{I}\right)[b\tau] + \left(\frac{2}{I}\right)\left[\frac{a+c}{2} + b(\tau - 1)\right] + \left(\frac{2}{I}\right)\left[a + \frac{a+c}{2} + b(\tau - 2)\right] + \\ & \dots + \left(\frac{2}{I}\right)\left[a(\tau - 2) + \frac{a+c}{2} + b\right] + \left(\frac{2}{I}\right)\left[a(\tau - 1) + \frac{a+c}{2}\right] + \left(\frac{I-2\tau-1}{I}\right)a\tau, \end{aligned}$$

²⁵See lemma 16 for this result.

²⁶See lemma 9.

where the first term is if the agent is shocked, the second if the shocked agent is in his neighborhood, the third if the shocked agent is in his neighbor's neighborhood, etc.²⁷ It is straightforward to show that $V_{UM}(I) > V_{LI}(I)$ for all finite $I > 2\tau$.²⁸

It should also be noted that the above results are all ex-ante. Obviously, if a household is in a bad marriage, it would prefer ex-post to be in a UM network.

6 Experimental Design and Findings

Our experimental design focuses on the stability of network structures when players are free to choose links.²⁹ The aim of our design was to test the theory developed in the previous sections. In particular, we are interested in testing the finding summarized in Proposition 2: in the presence of permanent shocks, the only strict pure-strategy perfect bayesian equilibrium networks satisfying the ex-ante payoff dominance criterion in the investment continuation game are M, M-A, and LI-A networks. To reduce the number of treatments we considered to a manageable number, we have chosen to focus on the stability of the three symmetric networks, M, LI, and UM. Our main experimental treatment variable consists of the initial network configuration in which agents interact: M, LI, or UM. A secondary treatment variable consists of the number of two-stage games played (5 or 9). Following the first two-stage game, players were free to choose the players with whom they proposed to form links in the first stage of all two-stage games. Our main finding is that, consistent with the prediction of Proposition 2, only M networks appear to be stable.

6.1 The investment game

We work with a specific parameterization for payoffs in the second stage investment game, which satisfy Assumption 1. The payoff matrix we adopt for the benchmark, symmetric 2×2 case, as would apply in a M network, is given below:

		1 Neighbor	
(i, j)		R	S
R		60	0
S		35	35

Payoffs are shown only for the row player i . In the case of a symmetric M network, the other players's payoffs can be inferred from such a representation. We note that for this benchmark, symmetric,

²⁷When $\tau \geq I/2$, the expected payoff from the strategy described is

$$\begin{aligned} \tilde{V}_{LI}(I) = & \left(\frac{1}{I}\right) [b\tau] + \left(\frac{2}{I}\right) \left[\frac{a+c}{2} + b(\tau-1)\right] + \left(\frac{2}{I}\right) \left[a + \frac{a+c}{2} + b(\tau-2)\right] + \\ & \dots + \left(\frac{2}{I}\right) \left[\left(\frac{I}{2}-2\right)a + \frac{a+c}{2} + b\left(\tau - \left(\frac{I}{2}-1\right)\right)\right] \\ & + \left(\frac{1}{I}\right) \left[\left(\frac{I}{2}-1\right)a + b\left(\tau - \left(\frac{I}{2}-1\right)\right)\right]. \end{aligned}$$

²⁸Subtracting V_{LI} from V_{UM} gives $\frac{1}{I}(a-b)\sum_{i=1}^{\tau} 2(i-1) > 0$.

²⁹An earlier draft of this paper contained experimental findings from treatments where the network structure was *exogenously imposed* in all games played so that players only made choices in repetitions of the second-stage investment game. We have chosen to remove the discussion of these findings in the interest of brevity. The interested reader is referred to our earlier working paper, Corbae and Duffy [6].

2×2 case, there are two pure strategy Nash equilibria: all-R and all-S. It is easily verified that all-R is the payoff dominant equilibrium, while all-S is the risk dominant equilibrium.³⁰

In the case of asymmetric network configurations, players would need to know the network configuration (how many links each player in a four player group had) as well as the payoff tables that agents with various ($k = 0, 1, 2, 3$) links faced. Such information was indeed provided to subjects, as explained below. But first, we explain how the payoff table, as shown above for the 1-neighbor case, was represented in the case where a player had 2 or 3 links (neighbors).

If a player is in a network configuration with two neighbors, as in an LI network, the payoff matrix was represented to them as:

		2 Neighbors		
(i, j)		2R	1R1S	2S
R		60	30	0
S		35	35	35

where 2R means that 2 of the $j = 2$ neighbors chose R, 1R1S means that 1 of the two neighbors chose R and the other chose S, and 2S means that both of the 2 neighbors chose S. The payoffs for these outcomes are consistent with the calculation in (5).³¹

Analogously, a player with three links—the most possible in groups of 4 players—as in a UM network, would see the following payoff table:

		3 Neighbors			
(i, j)		3R	2R1S	1R2S	3S
R		60	40	20	0
S		35	35	35	35

where again, the different payoff amounts reflect the weighting scheme in (5). It is easily verified that our choices for the payoff parameters, $a = 60$, $b = 0$, and $c = 35$ are consistent with Assumption 1.

Finally, we had to choose a payoff that subjects would earn per round in the event that they had no links, i.e. the parameter $d = u^i(a^i, \emptyset)$. We chose to set $d = b = 35$, so that the payoff to a player with no links is the same that a player could earn by having one link and always playing action S. We settled on this choice, rather than setting $d < b$, because we did not want subjects to be concerned that they would be worse off if they failed to establish any links; such a fear might cause them to send out link proposals to more players than they desired to be linked with as insurance that they would be linked. We note further that our choice for d is consistent with Assumption 1.

The payoff parameter values for a , b , c , and d represent cents earned in U.S. currency per play of the second stage game (e.g., a player whose payoff was 35 for a round earned U.S. \$0.35 for that round). Subjects kept their payoffs from all rounds of all games played, and in addition were awarded a fixed, \$5 participation payment.

³⁰There is also a mixed strategy equilibrium to the symmetric, 2-player game where each player plays action R with probability $\frac{7}{12}$ and earns an expected payoff of 35.

³¹For example, if a player with two neighbors chose R, and his two neighbors' choices were R and S (i.e. 1R1S), the player's equal weighted average payoff (using the 2×2 payoff matrix parameters) was $\frac{1}{2}60 + \frac{1}{2}0 = 30$. We saw no reason to explain to subjects the equal weighted average scheme by which these payoff tables were constructed. Berninghaus et al. [3] presented payoffs to players in their network games in a similar manner.

6.2 Experimental procedures

The experiments were implemented using networked PCs in the University of Pittsburgh’s experimental economics laboratory. The subject pool consisted of inexperienced undergraduates, recruited from the population of undergraduates at the University of Pittsburgh.

Prior to the start of play, subjects were given written instructions that were also read aloud to ensure that the information in the instructions was public knowledge.³² These instructions explained the various choices available to subjects, how these choices determined payoffs, and how payoffs translated into monetary payments. Included in the instructions were all three payoff tables presented in section 6.1. In addition, these payoff tables were drawn on a blackboard visible to all participants. The payoff to a player without any links was also carefully explained, as was the process for link formation (as discussed below). Finally, the instructions carefully explained that following the link formation phase, one player in each four-player group would be randomly chosen to receive a payoff shock and would be forced to play action S in all rounds of the subsequent investment game (the case of permanent shocks). We carefully explained that the location within each economy of the shocked player would not be revealed, and that the player chosen to receive the shock was an independent and identically distributed draw made following the network formation stage but prior to the play of each τ -round investment game. Any questions that subjects had were answered in private before play of the games commenced.

Each experimental session involved exactly 12 subjects with no prior experience of our experimental design. At the start of each session, subjects were randomly divided up into three groups of 4 players, or “economies,” labeled A, B or C. They remained in the same 4-player economy for the duration of the session.³³ Within each economy, players were identified only by their ID number 1,2,3, or 4 which also remained fixed for the duration of the experimental session. They then played a sequence of either 5 or 9 two-stage games. The sequence of play followed the same timing convention illustrated in Figure 2. In the first stage of each two-stage game, the network structure was determined by the proposal game. In the second stage, subjects played $\tau = 5$ rounds of the investment game against all of their neighbors as determined in the first stage.

In the first stage of the very first, 2-stage game, an exogenous, symmetric network structure was always imposed. This was done by having players choose particular links -the experimenter verified that this was done correctly - according to instructions we gave them. Thus, in the first, two-stage game alone, it is as if agents’ proposal action sets were restricted. In particular, in each session, we required group A to choose links so as to implement an M network, group B was to choose links so as to implement an LI network and group C was to choose links so as to implement a UM network. As noted above, payoff tables for all three types of networks (where players have 1, 2 or 3 links) were provided to all subjects in all groups, as part of the written instructions.

We chose to exogenously impose a particular network in the first two-stage game so as to ensure that subjects had experience with different network structures as well as to help coordinate players’ beliefs in subsequent proposal stages. Indeed, in reading the instructions aloud to all three groups of subjects, we were able to explain all three of the payoff tables that players might subsequently face when

³²Copies of the instructions used in our experiment can be viewed or downloaded at <http://www.pitt.edu/~jduffy/networks/>

³³The spatial location of members of a particular economy in the computer laboratory was randomly determined; it was pointed out to subjects in the instructions that they could not ascertain whether subjects near them in the layout of the computer laboratory were members of their economy or members of some other economy, thus reducing possibilities of collusion.

network links were freely determined. In addition, starting each group out in an exogenously imposed network configuration provides a clean test of the theoretical predictions; if a particular network structure is stable, then we should see it repeatedly reemerge when players have the opportunity to choose their own links, and this observation will form the basis of our experimental hypotheses. Following the completion of the first two-stage game, agents were free to choose which of the other three players they wanted to propose to link to in the first, link-formation stage of all subsequent games.

A screenshot of the link formation first-stage decision screen is shown in Figure 4. In this screenshot, player number 1 of economy A is choosing to form links to players 2 and 4.

[Insert Figure 4 here.]

After all players from all three economies had submitted their link proposals, the computer program found all mutually agreeable links and implemented the resulting network. The resulting network for each 4-player economy was depicted using a graphic on each player's screens. Each individual's own links were shown in red and links within the same 4-player economy that did not involve that individual were shown in green. Illustrative screens for player numbers 1 and 3 in economy A are shown in Figures 5 and 6. The network structure shown in these screenshots is LI as can be seen by the graphic in the upper left corner. The payoff table for an LI network is also shown. Note that player's choices were X and Y, with X corresponding to action R and Y corresponding to action S. The payoff table lists only the individual's own payoffs; in the case of the symmetric LI network, it was public knowledge that all other players network faced the same payoff table so players could easily infer their neighbor's payoff incentives. If a player receives a permanent payoff shock for the game, this information is only revealed *after* the network has been implemented. For example, in Figure 6 we see that player ID 3 is the player in Economy A who has been shocked. The shocked player does not make a decision; the computer automatically chooses action X (corresponding to S in the preceding sections) for this player in every round of the game.

[Insert Figures 5-6 here.]

After the network structure was imposed, players entered the second stage where they chose actions in the coordination game shown on their screens (if unshocked). After each round of a game all players are informed of their own action, the actions of their network "neighbor(s)" (if any) and their payoff for the round. Thus for example, in Figure 5, we see that player 1 chose action X (R) in round 1, and her two neighbors 2 and 4 also chose action X. The action choice of player 3 is not revealed as this player was not a neighbor of player 1. In round 2, player 1 again chose action X but her two neighbors chose action Y, as both of them had player 3 as neighbors.³⁴ In round 3, player 1 chose Y (S) as did her two neighbors, and the same outcome arose again in round 4, etc. Thus players not only learned their own payoff outcome from each round; they also knew which of their neighbors chose which action in every round of a game. The aim of this design was to give players information on other players' behavior so that they might make better informed decisions in the first stage game when they were free to propose links to the other players in their group.

³⁴In theory, of course, player 1 should have chosen action Y (S) in round 2; the screenshots are just an illustration of what could happen.

Following the completion of play of the first two-stage game, players were given additional written and oral instruction. They were told that in the first round of all subsequent games, each player in each group would have the opportunity to choose the players with whom they would form links. Subjects were informed that these link proposals would be made simultaneously and without communication. They were instructed about the need for mutual agreement between players for the establishment of links and were also informed of the payoff they earned in every period in the event that they had no links. Finally, subjects were told that they would not learn whether they faced a payoff shock until *after* all players had submitted link proposals and the network structure for the next five rounds of play had been implemented. Players submitted their link proposals by checking the boxes next to the ID numbers of players in their group whom they wanted to form links with as illustrated in Figure 4. Players were instructed that they were free to choose 0, 1, 2, or 3 links at every opportunity they were given to form links.

After players submitted their link proposals, the computer program found all mutually agreeable links and implemented the resulting network. This network configuration was shown on subjects' screens just as in Figures 5 and 6; the player's own links were shown in red and other links within the four player group not involving the player were shown in green. Since players had the payoff tables for the case of 1, 2, and 3 links, and also knew the payoff for no links, the graphical depiction of the network configuration allowed them to determine the payoff tables that all players in their four player group were facing. Of course, their own payoff table was prominently featured on the screen as well. We carefully explained to subjects that once networks were endogenously constructed, the payoff tables of their neighbors might differ from their own, due to possible asymmetries in the number of links among the players in each group. They were told to refer to the graphic on their screen to determine how many links each player in their 4-player economy had, and to refer to the various payoff tables given in the instructions to understand the payoff incentives these other players were facing.

We have conducted a total of 8 experimental sessions. Each session had 12 players divided up into three groups that were initially in either a M, LI or UM network as described above. In 4 of these 8 sessions, the three groups of players played a total of 5 two-stage games each; while the network structure was exogenously imposed in the first stage of the first game, in the subsequent 4 two-stage games, the network structure was *endogenously* determined by players themselves. After conducting these first four sessions and reviewing the results, we were curious to discover whether giving players more experience with endogenous network formation would matter for our findings. We therefore conducted 4 more sessions that were identical in all respects with the first four except that the 12 players in each of these additional sessions played a total of 9 two-stage games. Again, in the first stage of the first two-stage game, a network structure, M, LI or UM, was exogenously imposed on one of the three groups, but in the 8 subsequent games, the network structure was endogenously determined by the players in each group.

6.2.1 Hypothesis

Our main hypothesis is based on Proposition 2

Hypothesis 1 *When given the opportunity to choose links they will act so as to continue implementing an M network when they start in an exogenously imposed M network, but their link choices will not result in the continued implementation of LI or UM networks when they exogenously start out in them.*

It is important to recognize that this hypothesis relies upon other testable hypotheses. In particular, that subjects play according to the ex-ante payoff dominant perfect bayesian equilibrium strategies in the continuation game following the implementation of the network via the proposal game. Our experimental evidence will focus on all of these.

6.2.2 Endogenous Network Findings

The entire experimental dataset collected for this study is illustrated graphically in a series of figures that are collected in Appendix 2. There are two main sets of figures. Figures M5, LI5 and UM5 show results for the first four sessions where subjects played a total of 5 two-stage games. As noted above, in each of these four sessions, one 4-player group started out in M, one in LI and one in UM, but here we have rearranged the results, so that the results for all groups starting in M are presented in Figure M5, all starting in LI are presented in Figure LI5, and all starting in UM are presented in Figure UM5. The other set of figures correspond to the four sessions where subjects played a total of 9 two-stage games. We have again grouped together the results according to the initial exogenous network that players started out in. Hence Figures M9a,b,c,d, show results for four 4-player groups that played 9 games starting out in M, and Figures LI9 a,b,c,d and UM a,b,c,d do the same for the groups starting out in LI and UM respectively.

The results for each game are represented by two graphics. In the first graphic, the link proposal choices of the individual players, identified by the numbers 1,2,3,and 4 are shown as arrows emanating from each subject to the other players in his/her group. Double tipped arrows indicate mutually agreed upon links. The second graphic shows the network that was actually implemented based on the link proposals of the individual group members. In this same graphic, the player receiving the payoff shock (forced to play action S in all 5 rounds) is circled. Next to each player number is shown the sequence of actions chosen by that player in all five rounds of the coordination game P2 as played against that player’s network neighbors (if s/he had any). Thus, for example, *RSSSS* means that the player chose action R in the first round and action S in the last four rounds.

Finally at the bottom of each figure we report the frequency of “best response” behavior by all 3 unshocked players who had at least one link in each game. These best response frequencies were calculated as follows. For each game we counted the total number of times that each unshocked and linked subject played a best response to the history of action choices he observed given his knowledge of the network structure and assuming he was playing according to the PBE strategies in Lemmas 2 to 11. We then divided this count by the total number of choices made by all unshocked and linked players. For example, In Figure M5, Group 2, game 1, player 1 chose action *S* in the first round counter to the PBE prescription, but then chose S four more times in accordance with his history of interaction with player 2 (the shocked player) and with the PBE strategy for a M network given in Lemma 6. Hence, in game 1, subject 1 chose the right action (played a best response) 4 times. Player 3 also started out playing S counter to the PBE prescription. Given that choice, player 3 should have expected that player 4 would resort to playing S in the four remaining rounds, and so player 3 should have continued playing action S in the remaining 4 rounds. Instead, player 3 played R in the next four rounds. Therefore, we conclude that in game 1, subject 2 played 0 best responses. Finally, player 4 started out playing R as prescribed by the PBE strategy. Once player 4 observed that player 3 played S in round 1, player 4 should have resorted to playing S in the remaining 4 rounds of the game. In fact, player 4 played S in rounds 2 and 4 and R in rounds 3 and 5. We conclude that player 4 played best responses in 3 of the five rounds of game 1. The total best response frequency for this group for

game 1 is the sum of the individual totals, $4 + 0 + 3 = 7$ divided by 15, the total number of action choices, or .467, and this is the frequency represented by the first bar in the bar chart for Group 2 as depicted in Figure M5. The other best response frequencies are calculated in a similar fashion.³⁵

Our discussion of our experimental results is divided up into three parts: (a) the stability of networks, (b) link proposals, and (c) players' behavior in the second stage coordination game.

6.2.3 Stability of networks

We first consider whether our data supports Hypothesis 1.

Finding 1 *M networks are frequently stable while LI and UM networks are always unstable.*

Support. Consider first the case where players initially start out in exogenous M networks as shown in Figure 5M. After some experimentation, players in this treatment settled on link choices that resulted in implementation of M networks. In groups 2 and 3 this had happened by game 3; in group 2 by game 4 and in group 5, the M network was not established until game 5. The most common scenario was that players eventually returned to implementing the same M network configuration that was exogenously imposed in the first game. Only one group – group 2 – managed to coordinate on a different network structure and this one was also a symmetric M network!

We obtain very similar results in the 9-game sessions where groups started out exogenously in a M network, as shown in Figures M9.a, M9.b, M9.c and M9.d. Counter to our expectations, the increase in the number of games from 5 to 9 appears to lead to a longer period of experimentation with non-M networks. Nevertheless, in all four of the 9-game sessions, we find that players end up in a M network, and this is sustained for at least two sequential games. In these four sessions, players always returned to the same M network configuration that was exogenously imposed on them in the first game.

On the other hand, as Figures LI5, LI9.a, LI9.b, LI9.c and LI9.d reveal, when players started out in exogenously imposed LI networks, they did not consistently choose links so as to continue to maintain a LI network; the result of the endogenous link choices was, with one exception, some kind of hybrid network where players had two or fewer links each. There does not appear to be any systematic pattern to the sequence of hybrid networks that players implemented through their link choices; no hybrid network appears to have been sustained for more than 3 games in a row. Efficiency, as measured by the frequency of best response behavior, is generally lower among these groups by comparison with the groups that started out in M networks.

The one interesting exception is Group 3 in the 9-game treatment, whose behavior is shown in Figure LI9.c. Beginning with the second game, Players 1 and 3 in this group formed a unique marriage link to one another and never deviated from it for the duration of the session. This behavior resulted in a perfect M network for the last 6 endogenous games of the session. This finding adds to our evidence that M is empirically stable.

In all four groups shown in Figure UM5, the initial, exogenously imposed UM network is *never* observed after the first 5-round game, when players are free to choose their own links. The resulting

³⁵Notice that we are not allowing “forgiving strategies” that depend only on the history of play in the previous round. Our equilibrium predictions do not make use of such forgiving strategies, and that is why we do not consider them in our analysis of best response behavior. It should be noted, however, that forgiving strategies do not improve subjects ex-ante payoffs (along the equilibrium path) relative to the strategies we consider. Whether or not the ex-ante payoff dominant strategies we consider are actually chosen by the subjects is thus a matter of empirical verification which we take up.

networks are again hybrid-types where most players have two or fewer links. The frequency of best response behavior is comparable to that of groups that started out in LI, i.e. below the best response frequency for groups starting out in M. These findings are again robust to an increase in the number of games from 5 to 9 as shown in Figures UM9.a, UM9.b, UM9.c and UM9.d; the UM network exogenously imposed in the first game is never observed again in any subsequent game. Instead, there is typically much experimentation with various different asymmetric network structures, with no particular network being sustained for many games. The best response frequencies are typically below 100%. The one exception is Group 1 in the 9-game treatment, whose behavior is shown in Figure UM9.a. This group implements an M network beginning with Game 5, and maintains this same M network configuration in all subsequent games of the session. Again, such behavior would appear to support our conclusion that M is empirically stable. We note also that an M network is achieved in the last two games played by group 4 in the 9-game treatment (see Figure UM9.d).

The unraveling of the UM network may seem somewhat surprising in light of the results we obtained in Lemma 2 for the exogenous UM network (i.e. that all unshocked players play R thus avoiding any contagion) and the fact that in ex-ante terms players are better off than in LI. Still, we know from Lemma 15 that UM is not a weak PBE, and the experimental findings are consistent with this result. The unraveling of the LI network may also seem surprising in light of the results we obtained in Lemma 13. However, if we use a strict PBE refinement, then we shouldn't necessarily expect to see LI networks as an equilibrium outcome.

In Table 1, we summarize the network results from the figures found in Appendix 2 by reporting the frequency of players who had 0, 1, 2, or 3 links in each game, ignoring the first game where network links were predetermined. In games 2, 3 and 4, we pooled the results from both the 4 short and the 4 long sessions, while for games 5–9, the results are from the 4 long sessions alone. Figure 7 illustrates the mean link frequencies reported in Table 1 for each game of a treatment (players beginning in M, LI or UM).

[Insert Figure 7 here.]

As a severe stability test, we use the nonparametric, Kolmogorov-Smirnoff goodness-of-fit test to ask whether the sample cumulative distribution function over links $X = 0, 1, 2, 3$, $F^S(X)$, differs from a theoretically predicted cumulative distribution function, $F^T(X)$, where $F(X)$ represents the proportion of observations that are less than or equal to X . In the case where players started out in an M network, our hypothesis is that each player should maintain their single link, so the theoretical cumulative distribution function would be $F^T(0) = 0$, $F^T(1) = F^T(2) = F^T(3) = 1.00$. In the case where players started out in a LI network, our hypothesis is that players will not maintain two links (i.e. LI is not a strict PBE). In this case, we will nevertheless specify the “theoretical” cumulative distribution function is $F^T(0) = 0 = F^T(1) = 0$, $F^T(2) = F^T(3) = 1.00$. Finally, in the case where players started out in a UM network, our theoretical prediction is that players will *not* endogenously choose to have 3 links (i.e. UM is not even a weak PBE). As in the preceding case, we will specify the cumulative distribution function as $F^T(0) = F^T(1) = F^T(2) = 0$, $F^T(3) = 1.00$. The null hypothesis, H_0 , is that there is no significant difference between $F^S(X)$ and $F^T(X)$. In the case where players started out in M, failure to reject H_0 can be taken as support for our theoretical predictions, while in the case where players started out in LI or UM, our theoretical prediction is that H_0 will be rejected. The results from applying the Kolmogorov-Smirnoff test to H_0 in each game are shown in the last columns of Table . When players start out in M, we can reject H_0 , that 100% of players have a single

link – in games 2,3,5, 6, and over games 2-5. However, in the other four games, and over games 6-9 we cannot reject H_0 ; that is the empirical distribution in these games is not significantly different from one where every player has a single link. Thus, we find some mixed support for our theoretical prediction that H_0 will not be rejected. When players start out in LI, we can always reject H_0 , that 100% of players have two links, which supports our hypothesis that LI is not stable. Finally, when players start out in UM, we can always reject H_0 , that 100% of players have three links. In this case, rejection of H_0 is also consistent with our hypothesis that UM is not stable.

Finding 2 *Regardless of whether players start out in exogenously imposed M, LI or UM networks, in the subsequent endogenously chosen networks, most players have just one link.*

Support. As seen in Table 1, over all endogenous network games 2-9, and across the three treatments, the mean frequency of players with just one link is the largest and the magnitude of this frequency is highest in groups that started out in M and lowest in groups that started in UM.

Notice further that among players who started out in M, there is never any instance of a player having three links, and very few instances of players with two links. The relative frequencies of links appears to be more similar between the treatments where players started out in LI or UM networks. Confirming these findings, a nonparametric chi-square test reveals a significant difference in the relative frequencies with which players have 0, 1, 2, or 3 links over all games (2-9) between the treatment where players started out in M and 1) the treatment where players started out in LI ($p < .001$); 2) the treatment where players started out in UM ($p < .001$). Furthermore, it turns out that there is a significant difference in the relative frequencies of links between the treatment where players started out in LI and the treatment where players started out in UM ($p < .02$). These findings suggest that initial conditions with respect to the number of links players started out with, were important in subsequent games where players were free to choose links.

6.2.4 Link proposals

So far we have focused on the network links that players were able to form by mutual consent. We now take a step back and examine individual link proposals. We note that our theory is silent about the manner in which players form links, focusing instead on the stability of the resulting network structures. Still, it is of some interest to consider subjects' link proposals.

Finding 3 *Players nearly always proposed to link to at least one other player.*

Support. A careful study of the figures found in Appendix 2 reveals that when players were given the opportunity to choose links, 95% of all players – 91 out of 96 players – chose to make at least one link proposal in every game played.³⁶ Even the five exceptions to this rule *did make at least one link proposal in at least one game where players were free to propose links*. We conclude from this finding that our decision to set the autarchic payoff parameter, $d = b$ did not cause players to avoid proposing links. Indeed, it also serves to highlight the difficult coordination problem that players faced in the link formation stage. There were several instances where all players in a group submitted link proposals but there were no mutually agreed upon links, resulting in an A network.

³⁶The five exceptions to this rule are: player number 4 in Group 1 of the 9-game M treatment, Figure M9.a, games 2-4; player number 3 in Group 2 of the 9-game M treatment, Figure M9.b, game 5 only; player number 3 in Group 2 of the 9-game LI treatment, Figure LI9.b, games 2-6; player number 1 in Group 4 of the 9-game LI treatment, Figure LI9.d, game 5 only; and player number 2 in Group 3 of the 9-game UM treatment, Figure UM9.c, games 4-9.

We next examine the number of link proposals that players made over time and across the three treatments. In Table we report the mean frequencies with which players proposed 0, 1, 2, or 3 links in each game of a treatment, as well as over the first 5, the last 4 and all games played.

Finding 4 *In sessions where players started out in M networks, proposing a single link is the most common action in the first-stage game. The number of links players proposed in treatments where they started out in LI or UM networks is more varied and is often indistinguishable from a uniform random distribution over 1, 2, or 3 links.*

Support. Table 2 shows that the mean frequency of a single link proposal increases steadily over games in treatments where players started out in M networks. In treatments where subjects started out in LI networks, players appear to move away from proposing two links in favor of proposing one or three links as they gain experience. In treatments where players started out in UM networks, there also appears to be roughly equal numbers of subjects proposing 1, 2 or 3 links over all games played.

More precisely, let us exclude the case of zero link proposals, and ask whether the remaining distribution of proposals for 1, 2, or 3 links differs from a uniform random distribution, i.e. one in which the frequency of proposals for 1, 2 or 3 links is one-third each. We first rebalanced the proposed link frequencies, plf_i , $i=1,2,3$ as reported in Table 2 removing cases of 0 links, so that after rebalancing, $\sum_i = 1^3 plf_i = 1$. We then conducted a Pearson chi squared test of the null hypothesis that $plf_i = 1/3$ for all $i = 1,2,3$. The results of this test are reported in the last column of Table 2. We see that in the case where players started out in an M network, we can reject the null hypothesis that link proposals were randomly determined in every game, and over various groupings of games. In the case where players start out in LI networks, we cannot reject the null hypothesis for any individual game, however, over games 6-9 or all games 2-9 we can reject the null hypothesis. The reason for this outcome is that the Pearson chi-squared statistic is sensitive to the number of observations. In the case where players started out in UM networks, with a single exception - game 6, we are not able to reject the null hypothesis. The exception in game 6 is largely due to the anomalous behavior of a single group - group 4 of the UM treatment (see Figure UM9.d); each member of this group proposed a single link in Game 6 and the resulting network was the autarkic one (A).

Finally, we examine the ratio of the number of network links a player had to the number of links they proposed to form. As links are formed only by mutual consent, this ratio must always lie in the interval $[0, 1]$. While our theory is silent on link *proposals*, focusing instead on the resulting, mutually agreed upon *links*, it remains of interest to consider the ratio of links to proposals an important indicator of the *consistency* of players' beliefs (or desires) for a particular type of network; one could argue that a ratio of unity among all players of a group is consistent with the notion that players had achieved an equilibrium.

Finding 5 *The ratio of links to proposals is typically greater than 50%. When players start out in M or LI networks, this ratio increases as players gain experience.*

Support for this finding is shown in Figure 8, which plots the mean ratio of links to proposals over all players for all games of a session, including the first. Of course, in the first game, players were told which link proposals to make, so the links-to-proposal ratio in this first game is 1.0. A ratio of .50 would imply that each link required two proposals, on average, while a ratio of 1.0 would seem to be consistent with the notion of an equilibrium, where player's planned links equalled their actual links.

We see in the figure that for all treatments, the ratio of links-to-proposals generally lies between .5 and 1.0, especially toward the end of a session. Indeed, the ratio sharply increases in the 9-game M sessions from an average of .47 over games 2-5 to an average of .82 over games 6-9. Similarly for the 9-game LI sessions, the ratio increases from an average of .66 over games 2-5 to an average of .79 over games 6-9. In the case of the 9-game UM sessions, the ratio actually declines a little, from an average of .76 over games 2-5 to an average of .74 over games 6-9. Similar results are found for the 5-game sessions; there is an increase in the average links-to-proposals ratio from games 2-3 to games 4-5 in both the M and LI sessions, but not in the UM sessions. This finding tells us that player's beliefs or desires for a particular type of network are becoming more consistent with experience in the M and LI treatments, but there appears to be little or no increase in consistency in the UM treatments.

[Insert Figure 8 here.]

6.2.5 Behavior in the second-stage coordination game

Finally, we turn our attention to the choices players made in the second-stage coordination game. Given a network configuration, our theory prescribes how play should evolve in every round. In particular, regardless of the network structure, all of our equilibria have unshocked players with one or more links choosing action R in the first round of every new-second stage coordination game. It is therefore of interest to consider what actions players chose in this first round.

Finding 6 *In the treatment where players start out in M or LI networks, most (more than 50%) unshocked (and linked) players choose action R in the first round of each second-stage game, and the frequency of first round R choices is increasing over time. The same finding does not hold in the treatment where players start out in UM networks.*

Support. Figure 9 shows the frequency of play of action R in the first round of each second-stage coordination game using data from all sessions of a treatment. In the treatment where players started out in M networks, this frequency is typically in excess of 90% and is slightly increasing as players gain experience; in the first 5 games, the mean frequency is 93%, while over the last 4 games it is 95%. Similarly, for the treatment where players started out in LI networks, the mean frequency of play of R in the first round is 68% over the first five games and 76% over the last 4 games. By contrast, in the treatment where players started out in UM networks, the mean frequency of play of R in the first round falls from 68% over the first five games to 47% over the last four games.

[Insert Figure 9 here.]

The play of action R in the first round is an important indicator of whether players are ex ante payoff maximizers, since in this first round they cannot possibly know whether their neighbor is the lone, shocked player. It seems that our assumption of ex ante payoff maximization comes closest to being realized in the treatment where players started out in M networks. If we disaggregate the results presented in Figure 9 by group we find that for 6 of the 8 groups that started out in M networks, the frequency of play of action R in the first round exceeded 90% (over all games), and for all 8 groups that started out in M networks, this frequency always exceeded 50%. A binomial test confirms that we may reject the null hypothesis that players were equally likely to choose action R or S in the first

round of all games in favor of the alternative that they were more likely to choose action R ($p=.004$). For the treatment where players started off in LI networks, we find that for 7 of the 8 groups, the frequency of play of action R in the first round exceed 50%. Again, using a binomial test of the null hypothesis that players were equally likely to choose action R or action S, we can reject this hypothesis in favor of the alternative that they were more likely to choose action R ($p=.035$). For the treatment where players started out in UM networks, we find that for 5 out of the 8 groups, the frequency of play of action R in the first round exceeded 50%. In this case, the binomial test does not allow us to reject the null hypothesis that players were equally likely to choose action R or action S in the first rounds of this treatment. Furthermore, for the UM treatment, we observed a slight decrease over time, in the frequency of first round R choices, as illustrated in Figure 9. One possible explanation for these findings is that for many of the groups that started out in UM networks, the subsequent endogenously chosen networks were ones where most players were directly or indirectly linked with all of the other three players – perhaps a lasting legacy (hysteresis?) of the initial imposition of a UM network configuration. As a consequence, we surmise that players became wary of playing action R in any round, including the first, given the knowledge that they were frequently indirectly linked to the shocked player.

Consider next the frequency of best response play over all 5 or 9 rounds of the second stage game.

Finding 7 *The frequency with which unshocked players play best responses over all rounds does not differ across treatments. Further, we cannot reject the hypothesis that subjects play best responses at least 75% of the time across all treatments.*

Support. The bar charts at the bottom of Figures M5, LI5, UM5 and M9, LI9 and UM9 show best response frequencies over all rounds of each game for each group. Notice that with a single exception, group 4 in Figure M9.d, no group played according to the prescribed PBE strategy in every round of every second-stage game. While a perusal of these bar charts together with Finding 6, might lead one to the conclusion that players who started out in M networks played more frequently in accordance with our PBE predictions than players who started out in LI or UM networks, the evidence does not warrant such a strong conclusion. Consider Figure 10, which consists of two charts showing mean best response frequencies over all rounds for each group in the 5-game (top chart) and 9-game (bottom chart) sessions. Recall that in each experimental session (consisting of either 5 or 9 games), one 4-player group started out in a M network, one in a LI network and one in a UM network. It therefore seems reasonable to collect best response frequencies over all games played according to group/session numbers as is done in Figure 10. Consider first the four 5-game sessions. We see that in all four sessions, the best response frequencies over all rounds were always slightly higher for groups starting out in M than for groups starting out in LI. This difference, according to a nonparametric, robust rank order test is significant ($p=.10$). However, one cannot make the same claim in comparing the four groups starting out in M with the four starting out in UM, as one of the four UM groups, number 3, had the same overall best response frequency as their M group counterpart. Similarly, one cannot claim that groups that started out in UM played best responses over all games more frequently than groups that started out in LI, as one of the four LI groups, number 2, has a higher best response frequency than their UM counterpart. As for the four sessions where players played 9 games, we cannot reject the null hypothesis of no significant difference in best response frequencies between any two treatments.

[Insert Figure 10 here.]

We conclude that action choices were not that dissimilar over all rounds of the second stage coordination game. This finding is not difficult to reconcile with Finding 6; while players were more likely to play according to the PBE strategy in the first round of the M treatment than in the LI or UM treatments, such a finding does not imply that players will be less likely to play according to the PBE strategy in subsequent rounds. For instance, an unshocked player who plays S in all 5 rounds is only violating the PBE prescription in round 1; in the remaining four rounds he is playing a best response given the history of play and any network structure.

To more accurately gauge the overall frequency of best response behavior, we considered the frequency of best response play by all subjects, in all treatments and tested the null hypothesis that the frequency of best response play was at least .75, (which appears to be a reasonable guess, given the bar charts shown in Figure 10). A one-sided binominal test reveals that we are unable to reject this null hypothesis ($p = .913$). Further restricting the best response data to treatments where players started out in M, we are unable to reject the null hypothesis that the frequency of best response play was at least .90 ($p = .947$). We take this as evidence that the great majority of players were playing best responses in all of our treatments.

7 Conclusions

Consistent with the theory we have developed, in the environment we study we find evidence that M networks are stable (in the sense of replicating themselves), while both UM and LI networks break down. This is consistent with network choice on the basis of ex-ante payoff dominance where payoffs in a given network can be ranked $M > UM > LI$. Note that this ranking is nonlinear in the sense that payoffs in incomplete networks both dominate and are dominated by expected payoffs in complete networks. Note also that there are interesting ex-post issues involved, since the payoff in UM dominates the payoff in a shocked marriage. Finally, the instability of the complete UM network is due to a free rider problem. Thus, for these two reasons, there may be some justification for outside intervention (e.g. government intervention in banking networks)

There are several interesting extensions. One would consider variations in the shock processes. For instance, we could consider a process where there is no shock with probability q and permanent shock with probability $1 - q$. Would this result in risk dominant contagious behavior to kick in more frequently? Another interesting extension would examine whether we would see renegotiation of network proposals. For instance, would an unshocked M player in a bad marriage provide side payments to form an LI network with the other two unshocked partners?

Appendix 1: Proofs

Proposition 1. *Suppose shocks are transitory. If the strategy in which each player chooses R in each period he is not shocked is a PBE for some network structure, then it is an equilibrium for any network structure. In particular, whether or not these strategies constitute an equilibrium depends on the game parameters, but not on the structure of the network.*

Proof. With i.i.d. shocks, history is not informative about action sets. Thus, it suffices to show that the strategies described are an equilibrium of the stage game. Suppose that these strategies are an equilibrium of the stage game for some network structure. This means that it is a best response for a player who does not receive the payoff shock to play R given that other players who do not receive

the shock will play R. Consider such a player with n links. If she is not linked to the shocked player, her payoff from playing R is a ; if she is linked to that player, her payoff will be

$$\frac{n-1}{n}a + \frac{1}{n}c.$$

This player's subjective assessment that she is linked to the shocked player is $n/3$. Thus her expected payoff from playing R is

$$\left(1 - \frac{n}{3}\right)a + \frac{n}{3}\left(\frac{n-1}{n}a + \frac{1}{n}c\right) = \frac{2a+c}{3},$$

which is independent of the number of links the player has. Thus, whether or not these strategies constitute an equilibrium depends only on $\frac{2a+c}{3} \geq b$. ■

We now move onto the case with permanent shocks, where the crucial part of the analysis is to try to infer neighbors' type. To this end, we introduce the following notation. Let ω^i denote player i 's type, and denote the action taken by player i in round t by $a_t^i \in A^i(\omega^i) \subset \{R, S\}$. Write $a^{i,t-1} := (a_1^i, a_2^i, \dots, a_{t-1}^i)$ for the history of player i 's actions prior to round t . It will be convenient to use identical notation for the functions describing a generic behavioral strategy for the player. For example, if player i has two neighbors, say j and k , in some network, then the action prescribed by a pure strategy for player i in round t is a function of his own type and the actions taken by his neighbors prior to round t ; that is, $a_t^i = a_t^i(\omega^i, a^{j,t-1}, a^{k,t-1})$.³⁷

The next lemma is useful in establishing results for the case of neighborhoods of one or two players. It follows from $b > \frac{a+c}{2}$ that if an agent knows that one of his neighbors has the shock (and must play S), he should play S in subsequent rounds if his neighborhood has one or two neighbors.

Lemma 1 *Suppose shocks are permanent. Following any history in which some player with two or fewer neighbors knows that one of his neighbors has the shock, any strategy which calls for that player to play R in a subsequent round is strictly dominated by a strategy that is otherwise identical, except that the player plays S for the remainder of the game from this history.*

Proof. If an agent plays R in a state in which he knows one of his neighbors has the shock, he gets at most $\frac{a+c}{2}$ for the round. By committing instead to play S in each round subsequent to the discovery that a neighbor has the shock, this player guarantees himself a sure payoff of b in each round. Since $b > \frac{a+c}{2}$, this change represents a strict improvement over a strategy which calls for play of R in this case, and is otherwise identical. ■

The next lemma shows that despite the fact that the shocked agent is certainly in one's neighborhood in a UM network, it is possible to support R play by unshocked players since $\frac{2a+c}{3} > b$. With this strategy, there is no "contagious" S play in a UM network.

Lemma 2 *When the network is given by UM, there is an ex-ante payoff dominant, pure strategy PBE in which each unshocked agent plays R in every round.*

³⁷Allowing for dependence on the player's own past actions is redundant under pure strategies; this can be seen as follows. Suppose that we write $a_t^i = f_t(\omega^i, a^{i,t-1}, a^{j,t-1}, a^{k,t-1})$ in the example above. Then clearly we could write

$$\begin{aligned} a_t^i &= f_t(\omega^i, f_{t-1}(\omega^i, a^{i,t-2}, a^{j,t-2}, a^{k,t-2}), a^{j,t-1}, a^{k,t-1}) \\ &=: h_t(\omega^i, a^{i,t-2}, a^{j,t-1}, a^{k,t-1}). \end{aligned}$$

By recursion, it is clear that we can always reduce the function to the form presented in the text.

Proof. Consider the following strategy. Each unshocked agent plays R in the first round and thereafter plays R in each round in which at least three agents play R in the previous round, and plays S otherwise. First note that play of S by each player constitutes an equilibrium continuation following any deviation. To see that the stated strategies are a PBE, it suffices to note that the payoff of a player who has not received the shock is $\frac{2a+c}{3}$ in each period, which is greater than the one period payoff from playing S . Moreover, by deviating a player ensures that she will receive this lesser amount in every subsequent period as well. Thus, the prescribed strategies are mutual best responses. The ex-ante payoff is the average of the payoff received by the three players who do not receive the shock and the shocked player; that is $\frac{3}{4} \cdot \frac{2a+c}{3}\tau + \frac{1}{4}b\tau = \frac{2a+b+c}{4}\tau$.

To see that the equilibrium constructed is ex-ante payoff dominant, it suffices to show that the ex-post payoffs of each player in each round is as large as it can be. That this is so for the shocked player is true by definition. Now consider an unshocked player. The four possible payoffs for any given period are $\frac{2a+c}{3} > b > \frac{a+2c}{3} > c$ if 2 unshocked neighbors play R (the equilibrium path), if the agent chooses S , if 1 unshocked neighbors plays R , and if no unshocked neighbors play R , respectively. Playing the above strategy thus yields the highest in each period. ■

The next result, which follows from $b > \frac{a+c}{2}$, shows that a strategy where unshocked agents in an LI network play R in the first round until the shocked agent is discovered is an equilibrium (though obviously not unique). Using that strategy, the position of the shocked player can be inferred after one round. The agent who is diagonally across from the shocked agent anticipates that his unshocked neighbors will play S by lemma 1 and hence plays S . Notice that this result would be very different if we were using a solution concept like naive best response. Thus, the “contagious” S play spreads very quickly in our application, but would take another round with naive players.

Lemma 3 *When the network is given by LI, there is an ex-ante payoff dominant, pure strategy PBE in which all unshocked agents play R in the first round, then all agents play S in subsequent rounds.*

Proof. Consider the following strategy. Each unshocked agent plays R in the first round and plays S in subsequent rounds. Since play of S is an equilibrium of the one period game, play of S in rounds 2 through τ is an equilibrium continuation following any mode of play for date 1; therefore, to show that the play described is an equilibrium, it suffices to show that play of R in the first round is a mutual best response. In the first round, conditional on not receiving the shock, an agent’s prior that one of her neighbors has the shock is $\frac{2}{3}$ and that neither has it is $\frac{1}{3}$. Thus, the expected payoff for playing R in the first round is given by $\frac{2}{3} \left(\frac{a+c}{2} \right) + \frac{1}{3} (a) = \frac{2a+c}{3}$ while the payoff to playing S in the first round is b , which is less by Assumption 1. Thus, playing R is a mutual best response in round 1 when each continuation is independent of first round play. This shows that the stated strategy constitutes a PBE. Payoffs for an agent who is shocked are given by $b\tau$, for an agent whose neighborhood contains the shocked player is $\frac{a+c}{2} + b(\tau - 1)$, and for an agent whose neighborhood does not contain the shocked player is $a + b(\tau - 1)$. Thus, ex-ante payoffs are given by $\frac{2a+b+c}{4} + b(\tau - 1)$.

Next we show that play of R by three players can occur at most once in any pure equilibrium in an LI network, and that such a round can only be followed by play of S in every subsequent round. For concreteness, suppose that players are linked as $1 - 2 - 3 - 4 - 1$ in Figure 1 and suppose that after some history of shocks and actions $(\omega^i, a^{i,t-1})_{i=1}^4$, players’ strategies call for play of R at t by players $1 - 3$ and S by player 4. Then state ω is known to at least two players after such an occurrence: player 4, who has the shock; and player 2 who played R and is situated between two other players who did likewise. (Player 2 knows her own type and those of two others, and so she can infer the

type of player 4 whose play she cannot observe directly.) On the other hand, players 1 and 3 see play of one R and one S but cannot see whether each other is playing S and so cannot necessarily infer whether 4 is shocked or just playing S . Thus while players 2 and 4 are informed, players 1 and 3 are uninformed. However, the uninformed player 1 knows that agent 2 is informed. To see this, it is common knowledge that player 2 knows that if $\omega^3 = 0$, player 2's action history were $a^{2,t-1}$, and player 4's action history were $\{S, S, \dots, S\}$, then player 3 would play R by the above supposition (i.e. $a_t^3(0, a^{2,t-1}, \{S, S, \dots, S\}) = R$). Therefore, if there was another history with $\tilde{\omega}^3$ and $\tilde{a}^{4,t-1}$ but the same $a^{2,t-1}$ with $a_t^3(\tilde{\omega}^3, a^{2,t-1}, \tilde{a}^{4,t-1}) = S$, then player 1 knows that if agent 2 had observed play of S by agent 3, player 2 could infer player 3's type. That is, either $\tilde{a}^{4,t-1} \neq \{S, S, \dots, S\}$ in conjunction with $\omega^1 = \omega^2 = 0$ or $\tilde{\omega}^3 \neq 0$ implies $\tilde{\omega}^3 = 1$.

To finish the proof that play of R by three players can occur at most once and that such a round can only be followed by play of S in every subsequent round, first note that subsequent play of R by an LI agent who knows that at least one neighbor has the permanent shock is strictly dominated by Lemma 1. Next consider the problem faced by a neighbor of the informed players considered above. Player 1, say, knows that one of two things is true: either her neighbor who played S (player 4) actually has the shock, or player 3 has the shock and player 2 knows it. In either case, player 1 has at least one neighbor who will play S in each subsequent round. Therefore, she must play S herself in any equilibrium continuation. It follows that each player must do likewise. Since any situation in which three players play R in equilibrium is isomorphic to that described above, this result holds generally, as claimed.

To see that this equilibrium is ex-ante payoff dominant among symmetric equilibria, first note that the all- S equilibrium, which yields ex-ante payoff b in each round is dominated since $\frac{2a+b+c}{4} > b$ by Assumption 1. By symmetry, the first time anyone plays R , all three unshocked agents must play R , which we know from above must be followed by all- S . Since we have shown that three R s can occur at most once, no pure symmetric PBE can deliver a higher ex ante payoff than the one constructed, Q.E.D. ■

The next result establishes that a network where two agents have two neighbors while two agents have three neighbors (i.e. a mixture of UM and LI) has properties similar to that of either a UM or LI network depending on who gets the shock.

Lemma 4 *When the network is given by UM-LI, there is an ex-ante payoff dominant, pure strategy PBE in which all unshocked agents play R in the first round, all agents play S in subsequent rounds if a UM agent is shocked and all agents play R in subsequent rounds if an LI agent is shocked.*

Proof. Consider the following strategy. All unshocked agents play R in the first round, an LI agent continues to play R if both his neighbors have played R in all previous rounds, a UM agent continues to play R if his UM and one LI neighbor play R in all previous rounds, and agents play S otherwise. To show that the strategies described constitute an equilibrium, first notice that play of S in each round by each player is clearly an equilibrium continuation for any history, since a player gets $c < b$ by a unilateral deviation. Next, a deviation by any unshocked player in round 1 gets him at most $b\tau$ which is strictly dominated by expected payoffs to each player under the equilibrium: an unshocked UM player expects to receive $\frac{1}{3} \left[\frac{2a+c}{3} + b(\tau-1) \right] + \frac{2}{3} \left[\frac{2a+c}{3} \right] \tau$ and an unshocked LI player expects $\frac{2}{3} \left[\frac{a+c}{2} + b(\tau-1) \right] + \frac{1}{3} a\tau$, both of which are greater than $b\tau$. Finally, we must consider deviations from prescribed play after the first round. We have already established there are no profitable deviations if the UM player is shocked (and all agents play S as prescribed). In the case where it has been revealed

that an LI player has the shock, note that this player constitutes only $\frac{1}{3}$ of the neighbors of each of the UM players, so that they each get $\frac{2a+c}{3} > b$ by playing R when each of the other players does so. Furthermore, the other LI player gets a in this situation, so that play of R is a mutual best response following a shock to an LI player. Thus we conclude that the stated strategies constitute a PBE by the one-shot deviation property.

To see that the strategy is ex-ante payoff dominant, we must establish that the UM and LI players each receive the highest expected payoff among symmetric equilibria. There are four possible cases. First, that the UM players get the highest possible payoff when an LI agent is shocked follows from the proof of Lemma 2. Second, to show that an unshocked UM player gets the highest possible equilibrium payoff when the other UM neighbor is shocked we consider two cases: one where the UM agent always plays S and one where he plays R at least once. In the second case, we know that the location of the shock is revealed to everyone after the first play of R by a UM player in a symmetric equilibrium which results in S play for the remainder of the game by the LI agents by Lemma 1 and therefore the best he can get afterwards is b in each period under symmetric strategies. In the first case he receives $b\tau$ while in the second case he receives at most $\frac{2a+c}{3} + b(\tau - 1)$, which is what he gets in the proposed equilibrium. Since in both the first and second cases we have shown the UM agent receives the highest payoff ex-post, we know he receives it ex-ante. Third, to show that an unshocked LI player gets the highest possible equilibrium payoff when the other LI agent is shocked follows since he receives $a\tau$. Fourth, we must consider the case of an LI player when a UM agent is shocked. Ex-post, in this continuation, the only way for the LI player to do better is to play S in every round. Any symmetric equilibrium with this property must have the LI agent play S in the first round. In that case, the highest ex-ante payoff is $b + \frac{2}{3}b(\tau - 1) + \frac{1}{3}a(\tau - 1)$. But this is dominated by the expected payoff under the above mentioned strategy. Since in the third case we showed the LI agent receives the highest payoff ex-post and in the fourth case that any symmetric equilibrium with a higher ex-post payoff actually has a lower ex-ante payoff, we know he receives the highest payoffs ex-ante relative to any symmetric equilibrium. ■

The next result establishes that a network where three neighbors are linked to each other while one is in autarky (i.e. a mixture of LI and A) has similar properties to that of a UM network provided the agent in autarky has the shock while it has properties similar to an LI network otherwise.

Lemma 5 *When the network is given by LI-A, there is an ex-ante payoff dominant, pure strategy PBE in which all unshocked linked agents play R in the first round, and each plays R in the subsequent rounds if the autarkic player received the shock, and play S otherwise.*

Proof. Consider the following strategy for linked agents. All unshocked agents play R in the first round, each plays R in the subsequent rounds if the others did, and plays S otherwise. An unshocked autarkic player's strategy can be any arbitrary sequence of actions. It is easy to see that the strategies described constitute an equilibrium. That the equilibrium is payoff dominant among symmetric equilibria is easy to see in light of Lemma 1. ■

The next set of results shows that R play in all periods by a pair of unshocked players in a marriage is a PBE which yields the highest possible payoff a in each round. Since agents are ex-ante more likely to be in an unshocked marriage and S play may invoke an S response, in the first round it is optimal to play R until one knows whether one's partner is the shocked player, in which case it is optimal to play S since $b > c$.

Lemma 6 *When the network is given by M, there is an ex-ante payoff dominant, pure strategy PBE in which all unshocked agents play R in the first round, partners in the unshocked marriage play R in each subsequent round, and partners in the other marriage play S in each subsequent round.*

Proof. Consider the following strategy. All unshocked agents play R in the first round, players who have played R in each previous round continue to do so if their partner has played R in each previous round, and any player who has herself played S or whose partner has played S in some round plays S in each subsequent round. First note that play of S by each of a pair of linked players is a mutual best response for the two, so that the play prescribed following play of S by either of them constitutes a PBE for the continuation. Now, in equilibrium, an unshocked agent plays R in the first round, and continues to play R if her partner has not been shocked (that is, if her partner is observed to play R in round 1); the payoff to this player in this case is $a\tau$. If her partner has been shocked, she will observe S by that player, and she will then play S herself for the remainder of the game; the payoff to this player in this case is $c + b(\tau - 1)$. Since a player who has not been shocked perceives the probability that her partner has the shock as $\frac{1}{3}$, her expected payoff from playing according to the equilibrium is

$$\frac{2}{3}a\tau + \frac{1}{3}[c + b(\tau - 1)] = \frac{2a + c}{3} + \frac{2a + b}{3}(\tau - 1) > b\tau.$$

Thus, (by the one deviation property) playing R in round 1 is a best response when others are expected to play according to the equilibrium. Moreover, it is easy to see that continuing to play R in each period is also a best response in subsequent rounds when one's partner is expected to do so. In the PBE described, the shocked player gets $b\tau$, the two partners who have not been shocked receive $a\tau$, and the other player gets the payoff $c + b(\tau - 1)$. Ex-ante payoffs are then $\frac{2a+b+c}{4} + \frac{a+b}{2}(\tau - 1)$.

To see that the PBE is ex ante payoff dominant, it is sufficient to note that the ex-ante payoff obtained in the equilibrium is the largest expected payoff achievable by any single player in the M network without knowing the type of the player to whom she is matched. That is, a larger payoff can be obtained by a player only by knowing the type of her neighbor ex ante, but this is impossible. Since each player achieves this maximum feasible expected payoff, there can be no other PBE which gives a higher payoff to the players. ■

A related result for a marriage follows directly for the M-A network.

Lemma 7 *When the network is given by M-A, there is an ex-ante payoff dominant, pure strategy PBE in which each unshocked agent in a marriage plays R in the first round, partners in an unshocked marriage play R in each subsequent round, and partners in a shocked marriage play S in each subsequent round.*

Proof. Consider the following strategy for linked agents: each unshocked agent plays R in the first round, continues to play R if her partner has done so in the past, and plays S thereafter. The actions of unshocked agents in autarky may be chosen arbitrarily. The fact that this is a payoff-dominant equilibrium follows from Lemma 6. ■

The next result establishes that a network where two agents have two neighbors while two agents have one neighbor (i.e. a mixture of LI and M) has similar properties to that of an LI network since the shocked player will be in some LI player's 2 person neighborhood.

Lemma 8 *When the network is given by LI-M, there is an ex-ante payoff dominant, pure strategy PBE in which all unshocked agents play R in the first round and then all agents play S in subsequent rounds.*

Proof. Consider the following strategy. All unshocked agents play R in the first round, then all agents play S in the subsequent rounds. Since play of S is an equilibrium of the one period game, play of S in rounds 2 through τ is an equilibrium continuation following any mode of play for date 1; therefore, to show that the play described is an equilibrium, it suffices to show that play of R in the first round is a mutual best response. Now conditional on not receiving the shock, an LI player's prior that one of her neighbors has the shock is $\frac{2}{3}$ and that neither has it is $\frac{1}{3}$. Thus, the expected payoff of such a player for playing R in the first round is given by $\frac{2}{3} \left(\frac{a+c}{2} \right) + \frac{1}{3} (a) = \frac{2a+c}{3}$ while the payoff to playing S in the first round is b , which is less by Assumption 1. Furthermore, conditional on not receiving the shock, an M player's prior that her neighbor has the shock is $\frac{1}{3}$ and that she does not have it is $\frac{2}{3}$. Thus the expected payoff to such a player from playing R in the first round is $\frac{2}{3}a + \frac{1}{3}c$, while the payoff to playing S is b . Thus play of R by each non-shocked player is a mutual best response, and the stated strategies constituted a PBE, as claimed. In equilibrium, the expected payoff of an LI player is $\frac{2}{3} \left(\frac{a+c}{2} \right) + \frac{1}{3}a + b(\tau - 1)$ and that of an M player is $\frac{1}{3}c + \frac{2}{3}a + b(\tau - 1)$.

To see that the strategy is ex-ante payoff dominant for both the LI and M players, we first need three results, which we establish here. For concreteness, suppose the LI-M network is as depicted in Figure 1. Result 1 is that in the LI-M network in play of pure strategies, at each history which occurs with positive probability, each LI player knows whether the other LI player knows the type of the M player to whom the second is linked. To see this, we can write the behavioral strategy of an M player, say player 2, in round t as a function of only that player's type ω^2 and the history of her neighbor's, say player 1's, actions $a^{1,t-1}$; that is $a_t^2 = a_t^2(\omega^2, a^{1,t-1})$. Next suppose that we have that $a_t^2(\omega^2, a^{1,t-1}) \neq a_t^2(\tilde{\omega}^2, a^{1,t-1})$ for $\omega^2 \neq \tilde{\omega}^2$, so that player 2's type is revealed through her actions to player 1 after this history. Then it can be seen that, since the other LI player, say player 4, can observe the actions of player 1 (and players know the strategies of the other players), player 4 knows when $a_t^2(\omega^2, a^{1,t-1})$ is different from $a_t^2(\tilde{\omega}^2, a^{1,t-1})$; i.e., she knows when player 2 has revealed her type to player 1, proving the first result. Result 2 is that in the LI-M network, if an LI player knows that the other LI player knows the type of the M player to whom the second is linked, then the first LI player must play S in each subsequent round in equilibrium. To see this, note that if the first LI player, say agent 1, has the shock herself, then the result is trivial. Suppose agent 1 does not have the shock. In this case, the other LI player, agent 4, will play R in a subsequent round only if her M neighbor (player 3) was revealed not to have the shock. But if this is the case, then agent 4 plays R only if she knows that agent 1's M neighbor (player 2) has the shock. Thus in every case in equilibrium, at least one of agent 1's neighbors must subsequently play S in each round. The second result then follows from lemma 1. Result 3 is that whenever three agents play R , at least one of the LI players learns the type of her M neighbor, the proof of which is obvious.

Taken together, these three results imply that R play by three players can occur at most once in any pure strategy PBE: that is, by result 3, play of three R s causes at least one LI player to know the type of her M neighbor; by result 1, the other LI player knows that the first one knows this; then by result 2 and iterated elimination of dominated strategies, both of the LI players must play S in each round subsequent to play of R by three players.

Furthermore, these results imply that if no one has played R previously, then any strategy that calls for unshocked M agents to play R must be followed by S play by all agents thereafter. To see this, it is obvious that the LI agents learn the types of their M neighbors; by result 1, each LI player knows that the other LI player knows the type of the M player to which he is linked; by result 2, those LI agents play S thereafter; by lemma 1 all agents play S thereafter. Thus, there is no symmetric equilibrium in which only the two M agents play R in the first round in which R is played because

the M agents receive at most $c + b(\tau - 1)$, which is dominated by all- S .

Finally, we need to show that if no one has played R previously, then there is no symmetric equilibrium in which only the unshocked LI agents play R . To see this, first note that an LI agent who discovers one of his neighbors is shocked, plays S thereafter by lemma 1. Therefore, an unshocked M agent assigns probability $\frac{1}{2}$ to the event that his LI partner plays S forever after. Furthermore, by results 1 and 2, the M agent can only receive a payoff of a once. Thus, his maximum expected payoff from playing R in the period after his LI partner first plays R in this case is at most $\frac{a+c}{2} + b(\tau - 1)$, which is less than the payoff he receives from playing S in all periods. It follows from this reasoning that an M agent will never play R unless he sees his LI partner play R a second time. This would signal that neither his partner nor his partner's partner (the other LI agent) has the shock; at best this signal affords the LI agent one round with payoff a . The expected payoff for an unshocked LI player in this case is at most

$$\frac{1}{3} \left[3 \left(\frac{a+c}{2} \right) + b(\tau - 3) \right] + \frac{1}{3} (c + b(\tau - 1)) + \frac{1}{3} \left[2 \left(\frac{a+c}{2} \right) + a + b(\tau - 3) \right]$$

which can be shown to be less than what he receives by playing all- S which is $b\tau$. Thus, no such signaling strategy can be part of a symmetric equilibrium. Further, it is possible to show that more than two rounds of R play by LI agents to signal their type results in a lower expected payoff than $b\tau$.

To see that this equilibrium is ex-ante payoff dominant, note that the all- S equilibrium, which yields ex-ante payoff b is dominated since both the LI and M agents expect payoff $\frac{2a+b+c}{4} + b(\tau - 1)$ by Assumption 1. Since there is no symmetric equilibrium where only one type of LI or M agent plays R the first time anyone plays R and we have shown that play of R by both types gives at most a payoff equal to our proposed equilibrium, the result follows. ■

The next result shows that a star network where one agent has three neighbors while three agents have one neighbor (i.e. a mixture of UM and M) has similar properties to that of a UM network provided the middleman is not shocked and properties of a shocked marriage otherwise.

Lemma 9 *When the network is given by UM- M , there is an ex-ante payoff dominant, pure strategy PBE in which all unshocked agents play R in the first round, then they play R in the subsequent rounds if the UM player is not shocked and play S otherwise.*

Proof. Consider the following strategies. All unshocked agents play R in the first round, then they play R in the subsequent rounds if the UM player did, and play S otherwise. It is easy to see that these strategies constitute an equilibrium. To see that it is a payoff dominant equilibrium, note that we can restrict attention to strategies in which some player plays R if she is able in the first round. This is because given any equilibrium strategies which call for play of S by all players in round 1 independent of the shock, and play of R in some states in round 2, we can construct a payoff equivalent equilibrium strategy in which first play of R may occur in round 1.

First, consider an unshocked M agent. We claim that there is no action profile for the agents in which this agent's play in round 1 is independent of the shocks of the other agents that is as good for this player. To see this, note that, conditional on the location of the shock, the proposed strategies offer the highest possible payoff in each round after the first for this agent. Moreover, given her information in round 1, the expected first round payoff of this player is as good as it can be. Therefore, there can be no other strategy profile that is feasible under the information structure which gives a better payoff for this agent.

Now consider the UM agent. It is easy to see that this player's *ex post* payoff is maximal conditional on the location of the shock in each equilibrium outcome. Therefore there can be no other strategy profile which does better for this agent. ■

The LI-M-A network presents the first case where there is no ex-ante payoff dominant equilibrium for all types. There is one equilibrium that yields the highest ex-ante equilibrium payoff for the *LI* player and another equilibrium that yields the highest ex-ante equilibrium payoff for the *M* types. Each equilibrium makes at least one type of agent strictly worse off relative to the other equilibrium. Like the M network, the shocked agent may lie outside one's direct or indirect neighborhood (in this case it may reside with the *A* agent) and this leads to an inference problem which raises the possibility of signaling. In particular, strategies are such that in the equilibrium where the *LI* agent is best off, information about who is shocked is resolved only after the second round so that there are states of the world where the *M* agent could be made better off (i.e. when one of the *M* agents is playing *R* in round 2 while the network is heading to all-*S* because he doesn't know the other *M* agent has the shock). This information is made available in the equilibrium where the *M* agents are best off; in particular, in that equilibrium the *LI* agent provides a costly signal in the second round in the event that both *M* agents don't have the shock.

Lemma 10 *When the network is given by LI-M-A, if T is large, the ex-ante payoff dominant equilibrium for the LI type is strictly worse for the M players than the ex-ante payoff dominant equilibrium for the M types (and vice versa). Provided $T \geq 3$, play in the ex-ante payoff dominant equilibrium for the LI type is given by: in round 1, each unshocked linked agent plays R; thereafter, all linked agents play R provided the shock was not revealed to have been to one of them, and play S otherwise. Provided T is sufficiently large, the ex-ante payoff dominant (signaling) equilibrium for the M type is given by: in round 1, each unshocked linked agent plays R; in round 2, the M agents play S and the LI agent plays R if both of her neighbors did so in round 1 (i.e. he provides a costly signal as to whether one of the M players is shocked which is revealed by round 1 play) and plays S otherwise; from round 3 on, the three linked agents play R if the shock is outside of their network (i.e. if the A agent has the shock) and each plays S otherwise.*

Proof. The proof proceeds by showing that, if T is large enough, one strategy is payoff dominant for the M agents, while a distinct strategy is payoff dominant for the LI agent. Consider first the following LI-optimal strategy profile. In each round, each linked agent plays *R* if she has not previously played *S* and has not witnessed another agent play *S*, and she plays *S* otherwise. It is straightforward to show that these strategies constitute an equilibrium whenever $T \geq 3$, and that they give a highest payoff to the LI agent.

Now consider the following M-optimal strategy profile. In round 1, each unshocked linked agent plays *R*. In round 2, the M agents play *S* and the LI agent plays *R* if both of her neighbors did so in round 1; otherwise she plays *S*. From round 3 on, the three linked agents play *R* if each agent's neighbors has signaled that they do not have the shock, and each plays *S* otherwise. We will show that this strategy gives the best possible payoff for the M players, and that it is an equilibrium if T is large. (Note that it is symmetric in type.) To see this, we begin by stating several obvious facts:

(1) It is clear that a payoff dominant equilibrium for the M players must have the LI agent playing *R* in some states.

(2) It is clear that a (symmetric) payoff dominant equilibrium for the M players will have the M agents play *R* in some states. Thus, they will reveal their types to the LI player (simultaneously, by symmetry) in some period.

(3) Lemma 1 shows that, after learning the types of the M players, the LI player will subsequently play R only if both of the M players are revealed unshocked.

(4) In the round immediately after the M agents have first revealed their types to the LI agents, the subjective probability assigned by an unshocked M player to the possibility that the LI player will play S in every subsequent round is at least $\frac{1}{2}$. (It is $\frac{1}{2}$ if the LI player has already revealed to the M players that she is unshocked herself, and it is $\frac{2}{3}$ otherwise.) It follows that an unshocked M agent's subjective expected continuation payoff in the period immediately following the first play of R by unshocked M agents (at round t) is no greater than

$$x_T(t) := \frac{b}{2}(T-t) + \frac{1}{2}[b + a(T-t-1)]. \quad (8)$$

(This calculation reflects play of S by the M player in the current round followed by earning the maximum possible payoff in each subsequent period conditional on the location of the shock.) In particular, any equilibrium continuation in which the M player plays R in the current round must earn a worse (subjective expected) continuation payoff.

(5) The subjective expected period payoff of an unshocked M agent in the period in which the LI player first reveals herself is no greater than $(2a + c)/3$. (This follows from the fact that no information about the other players can be gleaned without this information; i.e., the type of the LI agent must be the first information that the unshocked M agent learns about the types of the other players.) This payoff can obviously be obtained only if the unshocked M player simultaneously plays R .

(6) In the period t that the M players first reveal their types, the best continuation payoff they can expect is

$$a + x_T(t),$$

but then only if the LI player has revealed previously, so that the M players are sure of her type.

It can be seen that there are 3 possible itineraries: the M agents reveal first, the LI agent reveals first, or all linked players reveal simultaneously. (Without loss of generality, we assume that one of these events occurs in round 1.) In the first case, the payoff of an unshocked M agent is at most

$$c + x_T(1).$$

In the second case, the unshocked M agent gets at most

$$\frac{1}{3}bT + \frac{2}{3}[b + a + x_T(2)].$$

In the third case, the unshocked M agent gets at most

$$\frac{2a + c}{3} + x_T(1).$$

Of these three payoff bounds, it is straightforward to show that the last is the greatest. Moreover, this payoff is attained in by the proposed strategies; thus, the proposed strategies offer the highest possible payoff for the M agents among all strategies which respect the result of Lemma 1.

It remains to show that this is acceptable to an unshocked LI agent. The unshocked LI agent's payoff under this scheme is

$$\frac{2a + c}{3} + \frac{2}{3}b(T-1) + \frac{1}{3}[c + a(T-2)].$$

It is clear that this is better than any possible deviation if T is large enough. Thus we have shown that the proposed strategies give the best feasible payoff to the M agents, and that they constitute an equilibrium if T is large enough.

Thus, there is no equilibrium which is payoff dominant for all of the players. ■

The last (asymmetric) network we consider is UM-LI-M, which provides another example where there is no equilibrium which is ex-ante payoff dominant for all types. There is one equilibrium that yields the highest ex-ante equilibrium payoff for the UM and LI players and another equilibrium that yields the highest ex-ante equilibrium payoff for the M player. Each equilibrium makes at least one type of agent strictly worse off relative to the other equilibrium. This network has properties similar to that of LI if either the UM or LI agents are shocked and properties similar to that of UM if the M agent is shocked. Unlike the previous case (LI-M-A), in the equilibrium which is best for the UM and LI players, there is no inference problem after the first round (i.e. after the first round, all agents either directly observe an S play by a neighbor or can infer that someone out of their neighborhood has received the shock). However, if T is sufficiently large, there is another equilibrium where information about shocks is actually revealed slowly which results in higher ex-ante payoffs for the M agent. Each equilibrium makes at least one type of agent strictly worse off relative to the other equilibrium.

Lemma 11 *When the network is given by UM-LI-M, if T is large, the ex-ante payoff dominant equilibrium for the UM and LI types is strictly worse for the M players than the ex-ante payoff dominant equilibrium for the M types (and vice versa). Play in the ex-ante payoff dominant equilibrium for the UM and LI types is given by: in round 1, all unshocked agents play R ; if the M agent is shocked, all unshocked agents continue to play R and play S otherwise. Provided T is sufficiently large, the ex-ante payoff dominant equilibrium for the M types is given by: in round 1, an unshocked UM agent plays R and other players play S ; in round 2, if the UM agent previously played R , then the UM player and an unshocked M player play R while the other players play S , and everyone plays S otherwise; in the third round, if the UM player has played R in the first two rounds, then all unshocked agents play R and play S otherwise; after the third round, agents continue to play R if the M agent played R in the second round, and they play S otherwise.*

Proof. Consider the following strategies for the $UM - LI$ -optimal equilibrium. In round 1, all unshocked agents play R . If the M agent is revealed to have the shock, then all of the other agents continue to play R as long as each of them has always played R in the past. In all other continuations, all agents play S . It is straightforward to show that this is an equilibrium and that it is the unique ex-ante payoff dominant equilibrium for both the UM and LI agents.

If T is sufficiently large, then the following strategies constitute the M -optimal pure-strategy symmetric equilibrium for the M player. In round 1, the UM agent plays R if he is able and other agents play S . If the UM agent played S in round 1, all agents play S subsequently. If the UM agent played R , then in round 2, the UM agent and the M agent (if she's able) play R , and the other agents play S . If the UM agent has played R in the first two rounds, then all agents who are able play R in round 3. After round 3, agents continue to play R if they have observed two neighbors play R all rounds after round 2, and they play S otherwise.

To see that this is the best possible equilibrium for the M player, consider each of the four possible states. When either the M player himself or the UM player gets the shock, then this equilibrium delivers the highest possible payoff bT to this agent. I claim that the UM player will play R at most once in any equilibrium after learning that the M player is unshocked. To see this, note that the only period payoff better than b for this player comes in a round when 2 of his neighbors play R . By

symmetry, the LI players must learn whenever this occurs. Thus if the state is such that one of the LI players has the shock, Helpful Lemma implies that neither of them will ever play R again. Thus the UM player will not play R after the LI players learn that one of them has the shock, proving the claim. Thus it can be seen that the M player can not earn a higher payoff in any equilibrium. ■

Now we move onto predicting which networks we expect to arise in the above game. The idea is to endow agents with beliefs that play in a network that results from a unilateral deviation from a given network will follow the ex-ante payoff dominant perfect bayesian equilibrium strategies discussed in the above lemmas in the subgame following that deviation.

The first result, that an M network is an equilibrium in the sense that it is (strictly) immune to unilateral deviations, follows simply from lemmas 6 and 7 (trivially). In particular, a unilateral deviation from sending a proposal to one's partner results in autarky, where payoff $d\tau$ is strictly less under Assumption 1 than the ex-ante payoff associated with M given by $\frac{2a+c}{3} + \frac{2a+b}{3}(\tau - 1)$.

Lemma 12 *An M network is a (strict) equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. Consider any proposal strategies which result in a given M network. (For simplicity, we will assume that players make only those proposals which are necessary and sufficient for such a network to obtain, though this is not necessary for this case.) Now the only unilateral deviations for the proposal stage which can affect a player's payoff are to propose zero links. This deviation is clearly worse for the player, since she receives a payoff of d in any continuation. This shows that these link proposals are an equilibrium with the prescribed mode of continuation play. ■

Corollary 1 *An M-A network is a (strict) equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. Since the deviation we are considering is for an M type agent that lands him in autarky, the same reasoning as in lemma 12 holds in an M-A network. ■

The next result uses lemmas 3 and 8 to establish that LI is an equilibrium. It shows that an LI network in which all agents send two proposals is stable in the sense that there is no strictly profitable unilateral deviation that brings about LI-M. To understand the result, suppose agent 2 deviates and chooses not to send a proposal to agent 3, while all other agents send two proposals associated with the original LI network. As we saw in lemma 8 the equilibrium play in LI-M is identical to equilibrium play in LI since the M player, if he is unshocked, knows that one of the two LI players is linked to a shocked player after the first round, thereby altering his beliefs and best responding with S play in the subsequent rounds as in lemma 3. Since equilibrium play is the same, ex-ante payoffs are identical so that the deviation is not strictly profitable. There is an important sense, however, in which LI is not stable which corresponds informally to an evolutionary stability type argument. That is, a best response to agent 2's single proposal to agent 1 is for agent 1 to send a single proposal to agent 2. As above, agent 2 does no worse sending one proposal and both do better getting into a marriage. This type of proposal strategy would displace LI as an equilibrium.

Lemma 13 *An LI network is a (weak) equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. Consider the strategies for the proposal stage in which players make the proposals which are necessary and sufficient for the formation of a given LI network; i.e., no player makes a proposal which is not reciprocated. Now the only unilateral deviations for the proposal stage which can materially affect a player's payoff are to propose zero or only one link instead of the prescribed two. The affect of proposing zero links is that the deviating player becomes an A player in a LI-M-A network. This deviation is clearly worse for the player, since she receives a payoff of d in any continuation. By proposing only one link, the deviating player becomes an M player in a LI-M network. Using the result above about the equilibrium followed in the investment stage of the LI-M network, it can be seen that such a player earns a payoff identical to what he would earn by playing according to the prescribed equilibrium. Thus, no player can have a profitable deviation. ■

The next result uses lemmas 5 and 10. It shows that an LI-A network where three agents send out two proposals to each other is stable in the sense that there is no profitable unilateral deviation that brings about LI-M-A. To understand the result, suppose the LI-A network is as depicted in Figure 1, where agent 3 is the sole A player, and the other 3 players are LI players. Suppose that agent 2 deviates and chooses not to send a proposal to agent 4, while agents 1 and 4 continue to send two proposals that would have resulted in LI-A. To show that 2's deviation is not profitable, it is sufficient to assign beliefs to the deviator that play in LI-M-A will be payoff dominant for the M type.³⁸ The deviation does nothing to insulate agent 2 from the shock and in the case that the shocked player is A, actually leads to a lower payoff. We also note that this equilibrium is stable in a sense that LI is not; an LI player in LI-A would do worse to accept a proposed link by the A player. This is in contrast to what is the case for the LI network, where a player does better to accept a proposed link by a player to whom he is not linked.

Lemma 14 *An LI-A network is an equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. Considering the minimal profile of proposals necessary, we need only consider the outcome where an LI player drops one or more of his links. By doing so unilaterally, this player can become either an M player in an LI-M-A network, or an A player in an M-A network. In the payoff dominant equilibrium for the LI-A network, an LI player gets

$$\frac{a\tau}{4} + \frac{b\tau}{4} + \frac{1}{2} \left[\frac{a+c}{2} + b(\tau-1) \right].$$

In the equilibrium that is payoff dominant for the M type in the LI-M-A network, an M player gets

$$\frac{1}{4} [b + a(\tau-1)] + \frac{b\tau}{4} + \frac{1}{4} [c + b(\tau-1)] + \frac{1}{4} [a + b(\tau-1)].$$

Subtracting the latter expression from the former gives $a-b$, so that an LI player in the LI-A network cannot improve his payoff by deviating unilaterally to effect an LI-M-A network under our assumptions. Since it is easy to see that the player does worse as an A player in a M-A network, LI-A can be seen to be an equilibrium network, as claimed. ■

The next result uses lemmas 2 and 4. It shows that a UM network in which all agents send three proposals is not stable in the sense that there is a profitable unilateral deviation that brings about

³⁸If this is dominated, then payoffs for the M player associated with the payoff dominant equilibrium for LI types would even be lower.

UM-LI. To understand the result, suppose agent 1 deviates and chooses not to send a proposal to agent 3, while all other agents send proposals to all other agents. The resulting UM-LI network (as shown in Figure 1) means that agent 1's two neighbors (agents 2 and 4) "provide insurance" to agent 1 (continue to play R) in the event that agent 3 gets the shock. In that event, agent 1 receives payoff a while in the UM network he would receive $(2a + c)/3 < a$. Thus, the resulting instability of the UM network is similar to a free-rider problem. That is, each agent has an incentive to enjoy the benefits of insurance against payoff shocks (the public good) provided by others while providing it insufficiently herself.

Lemma 15 *A UM network is not an equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. By deviating unilaterally from the prescription that all players propose links to all of the other players, a given player may find herself as an LI player in a UM-LI network. Given the lemma above guiding play in the UM-LI network, the LI player earns

$$\frac{a}{2} + \frac{b}{4} + \frac{c}{4} + \left(\frac{a}{4} + \frac{3b}{4}\right)(\tau - 1).$$

This is greater than

$$\left(\frac{a}{2} + \frac{b}{4} + \frac{c}{4}\right)\tau,$$

the expected payoff for a player in a UM network.³⁹ Since there is a profitable deviation from any strategies which result in a UM network when play continues as we've assumed, such a network cannot occur in equilibrium. ■

A star network is considered in many papers. Under Assumption 1, the UM player could unilaterally deviate and send only one proposal, resulting in his own marriage. His ex-ante payoffs $\frac{2a+b+c}{4} + \frac{a+b}{2}\tau$ established in lemma 6 from being in a marriage are strictly higher than the expected payoffs by being the middleman $\frac{2a+b+c}{4}\tau$ since in the event that he is unshocked, he provides insurance against the shock with probability one each period.

Lemma 16 *A star (M-UM) network is not an equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. By deviating from a proposal strategy which would cause her to be the UM player in an M-UM network, a player can become an M player in an M-A network. By lemma 7, it is easy to show that this player then earns an ex ante expected payoff of $\frac{2a+b+c}{4} + \frac{a+b}{2}\tau$ in the payoff dominant continuation. This is greater than $\frac{2a+b+c}{4}\tau$, the ex ante expected payoff of the UM player in an M-UM network in any payoff dominant continuation. ■

³⁹To see this, note that Assumption 1 implies that

$$\frac{a}{4} + \frac{3b}{4} = \frac{a}{4} + \frac{b}{2} + \frac{b}{4} > \frac{a}{4} + \frac{1}{2} \left(\frac{a+c}{2}\right) + \frac{b}{4} = \frac{a}{2} + \frac{b}{4} + \frac{c}{4}.$$

The following result shows that UM-LI network is not an equilibrium since a UM player can deviate unilaterally to become an M player in a UM-M network. Provided the middleman in the resulting UM-M network does not receive the shock, this deviation generates the benefits of an unshocked marriage to the deviator and yields a strict improvement over having to provide insurance in the UM-LI network as above.

Lemma 17 *A UM-LI network is not an equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. A UM player in a UM-LI network can deviate unilaterally to become an M player in a UM-M network. Computing the payoff of the M player in the payoff dominant equilibrium described in the UM-M Lemma above, it can be seen that the player gets

$$\frac{2a}{4} + \frac{b}{4} + \frac{c}{4} + \left(\frac{a}{2} + \frac{b}{2}\right)(\tau - 1).$$

This payoff is higher than

$$\frac{2a}{4} + \frac{b}{4} + \frac{c}{4} + \left(\frac{a}{3} + \frac{b}{2} + \frac{c}{6}\right)(\tau - 1),$$

which can be seen from Lemma 4 to be the payoff of the UM player in a UM-LI network. ■

The next set of results establish that an LI-M or LI-M-A network is not an equilibrium since an LI player can deviate unilaterally and receive the strictly higher ex-ante payoff of an M or M-A network respectively.

Lemma 18 *An LI-M network is not an equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant among all pure PBE continuations in the investment stage.*

Proof. An LI player in an LI-M network can deviate unilaterally to become an M player in an M network. The latter player earns

$$\frac{a}{2} + \frac{b}{4} + \frac{c}{4} + \left(\frac{a}{2} + \frac{b}{2}\right)(\tau - 1),$$

whereas it can be computed from Lemma ?? that the payoff of the LI player in the LI-M network is

$$\frac{a}{2} + \frac{b}{4} + \frac{c}{4} + b(\tau - 1),$$

which is lower. ■

Lemma 19 *An LI-M-A network is not an equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant for at least one type among all pure PBE continuations in the investment stage.*

Proof. By deviating unilaterally from this network, the LI player can become an M player in an M-A network. From lemma 7, it is easily seen that the latter player does better than the former player. ■

The next result is based upon the same idea but now it is the UM type in a UM-LI-M network who can deviate by sending only one proposal to be linked with one neighbor resulting in strictly higher ex-ante payoffs associated with an M network.

Lemma 20 *A UM-LI-M network is not an equilibrium outcome under the assumption that players always play according to a PBE continuation which is ex-ante payoff dominant for at least one type among all pure PBE continuations in the investment stage.*

Proof. By deviating from this network, the UM player can become an M player in an M network. From lemma 6, it is easy to see that the latter player receives a higher payoff than the former in either continuation game ■

Appendix 2: Data

In this appendix, we graphically depict on the following pages, *all* data collected in our experimental study. An explanation of how to read the following figures is provided in section 6.2.2.

Figure M5: Endogenous Network Treatment, 4 Groups Initially in M Play 5 Games

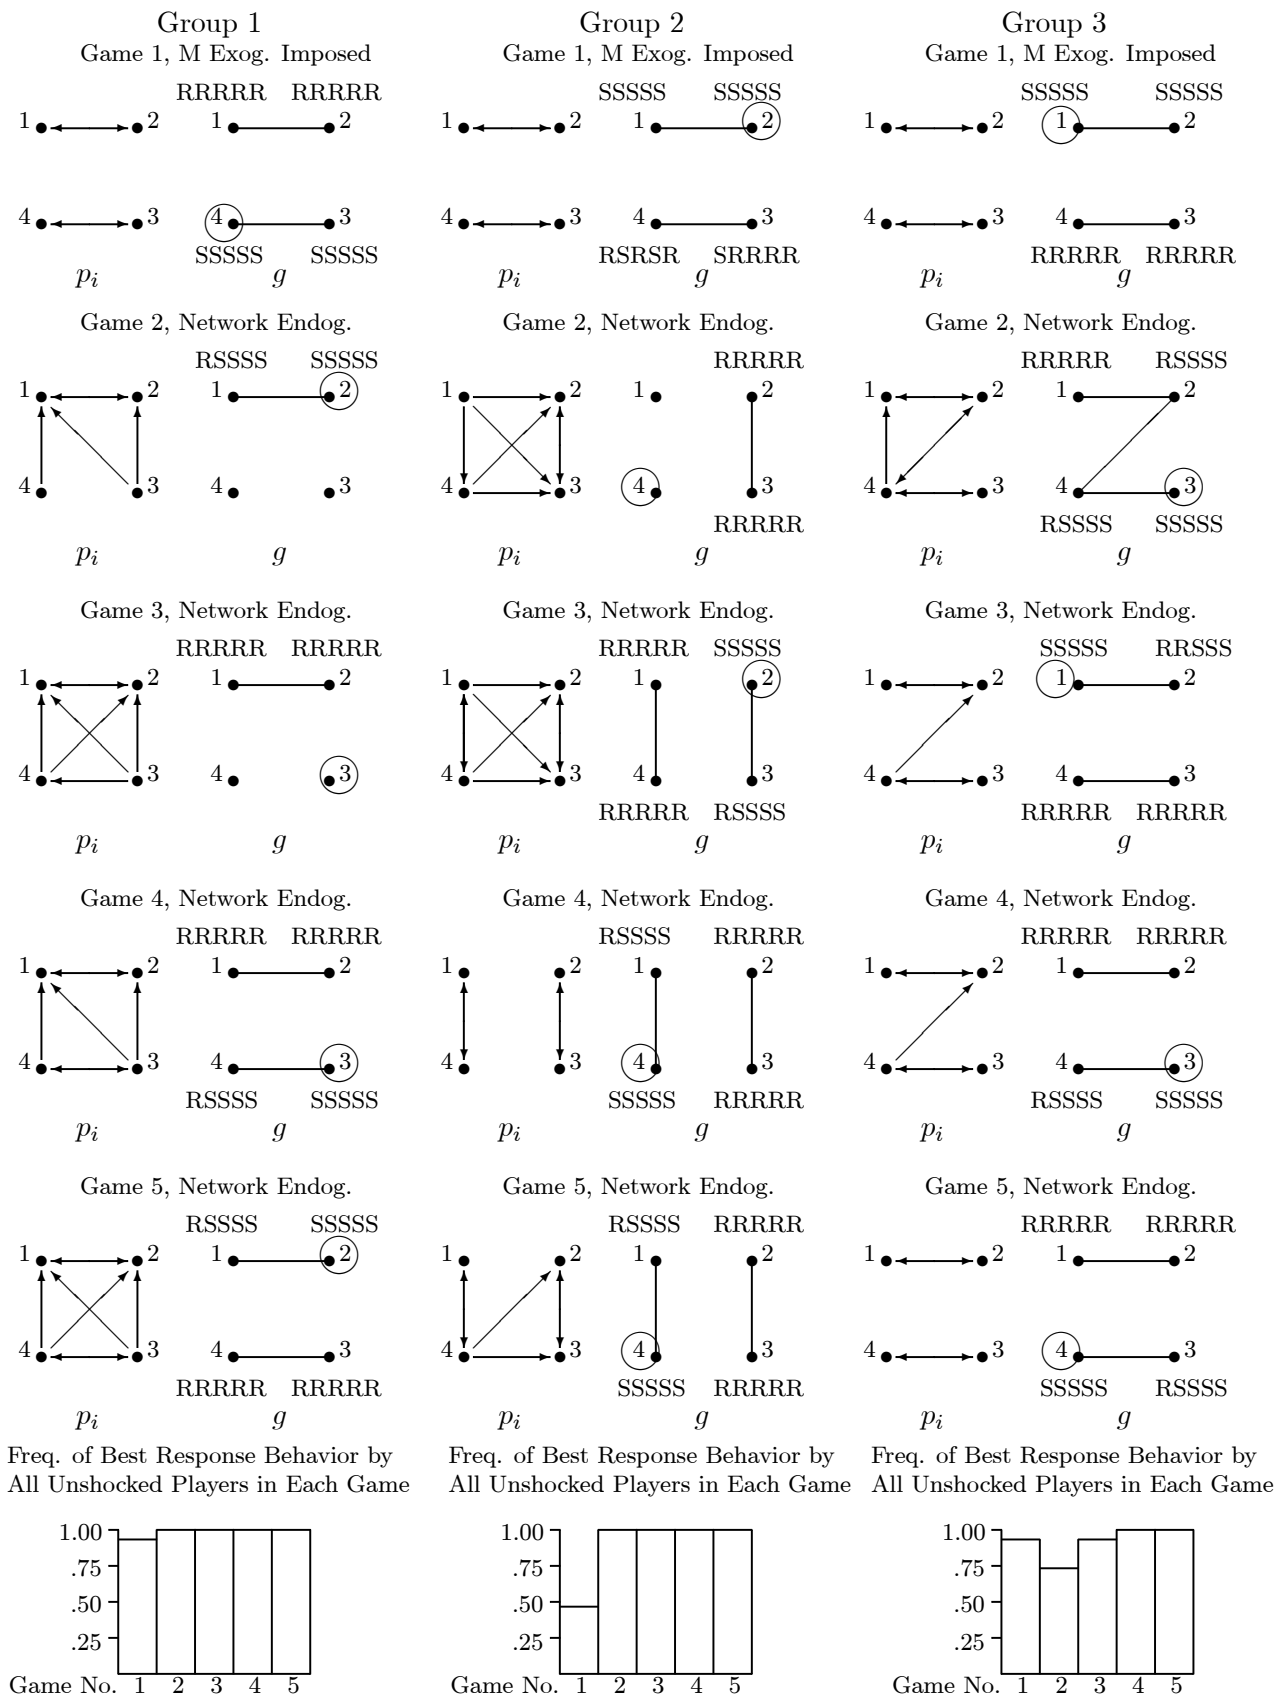
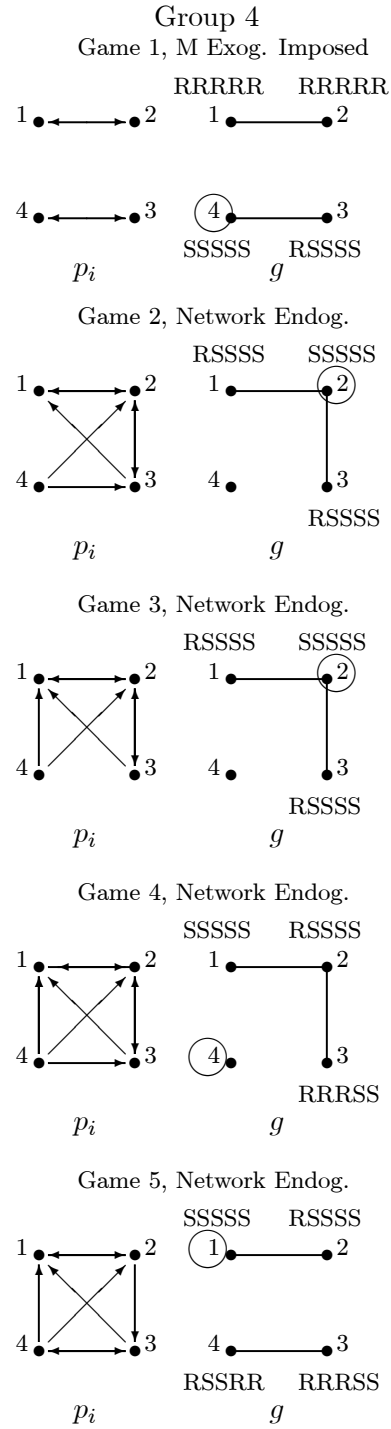


Figure M5: Continued, Endogenous Network Treatment, 4 Groups Initially in M Play 5 Games



Freq. of Best Response Behavior by All Unshocked Players in Each Game


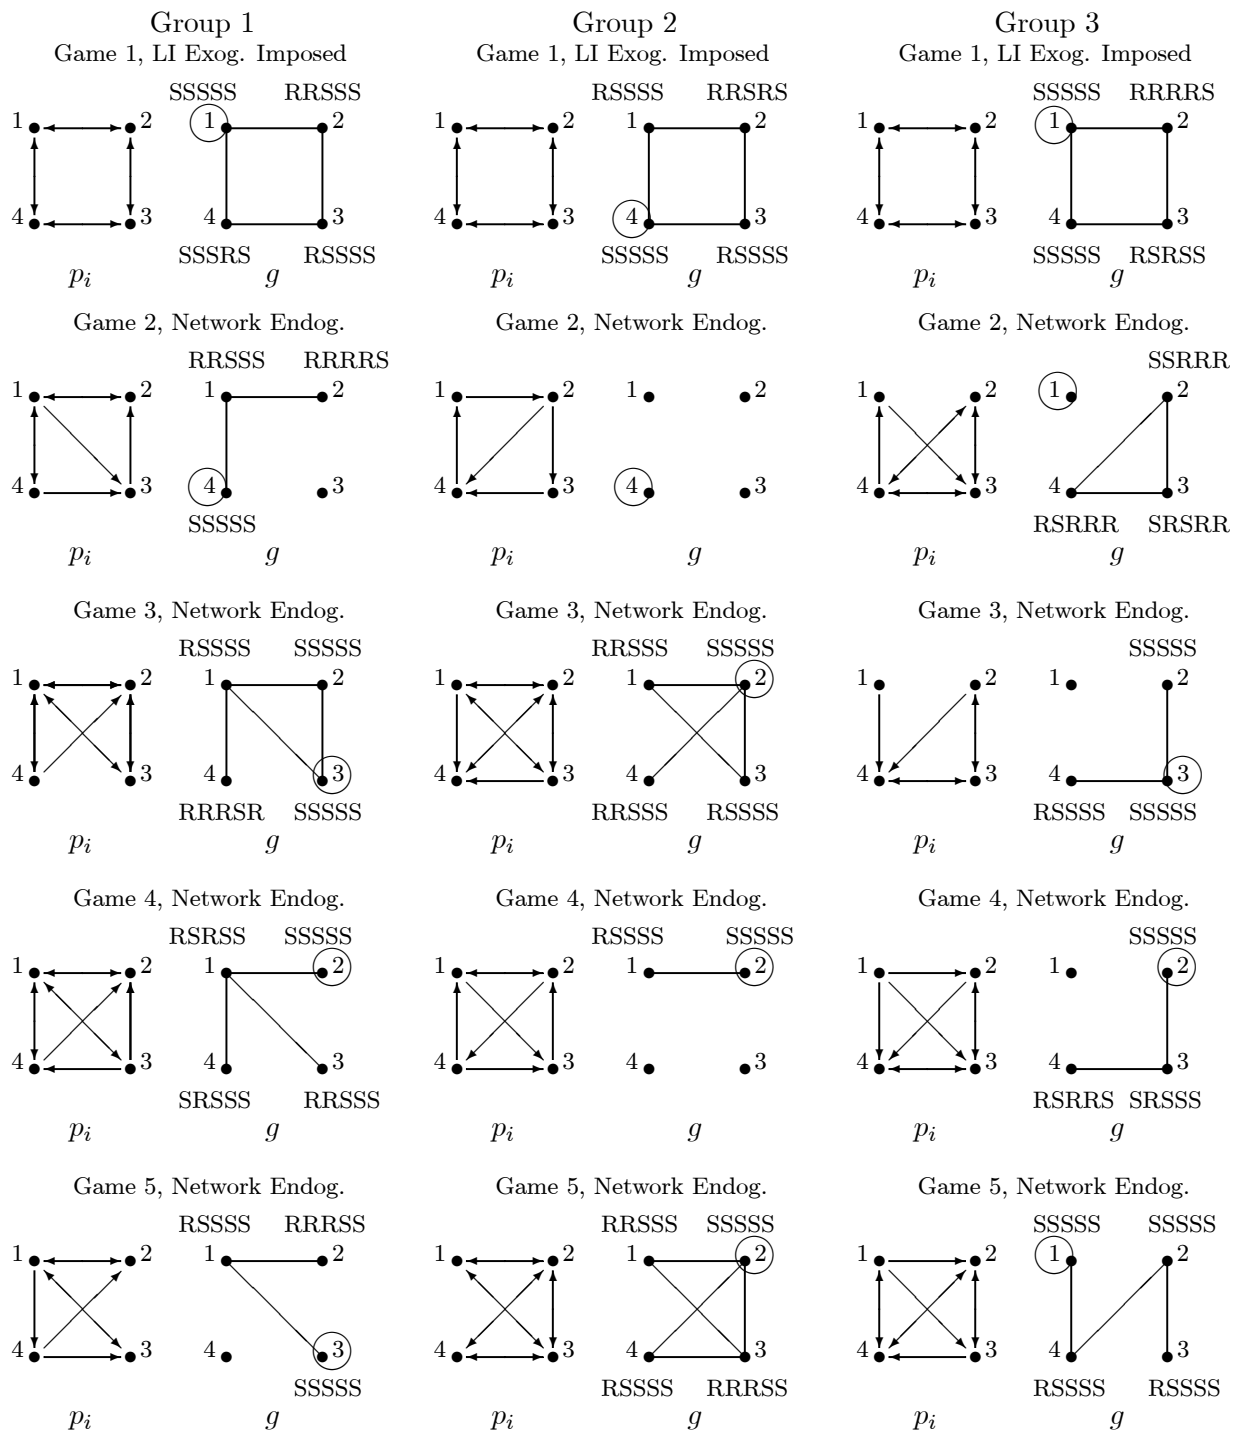




Figure LI5: Endogenous Network Treatment, 4 Groups Initially in LI Play 5 Games



Freq. of Best Response Behavior by All Unshocked Players in Each Game



Freq. of Best Response Behavior by All Unshocked Players in Each Game



Freq. of Best Response Behavior by All Unshocked Players in Each Game









Figure LI5: Continued, Endogenous Network Treatment, 4 Groups Initially in M Play 5 Games



Freq. of Best Response Behavior by
All Unshocked Players in Each Game


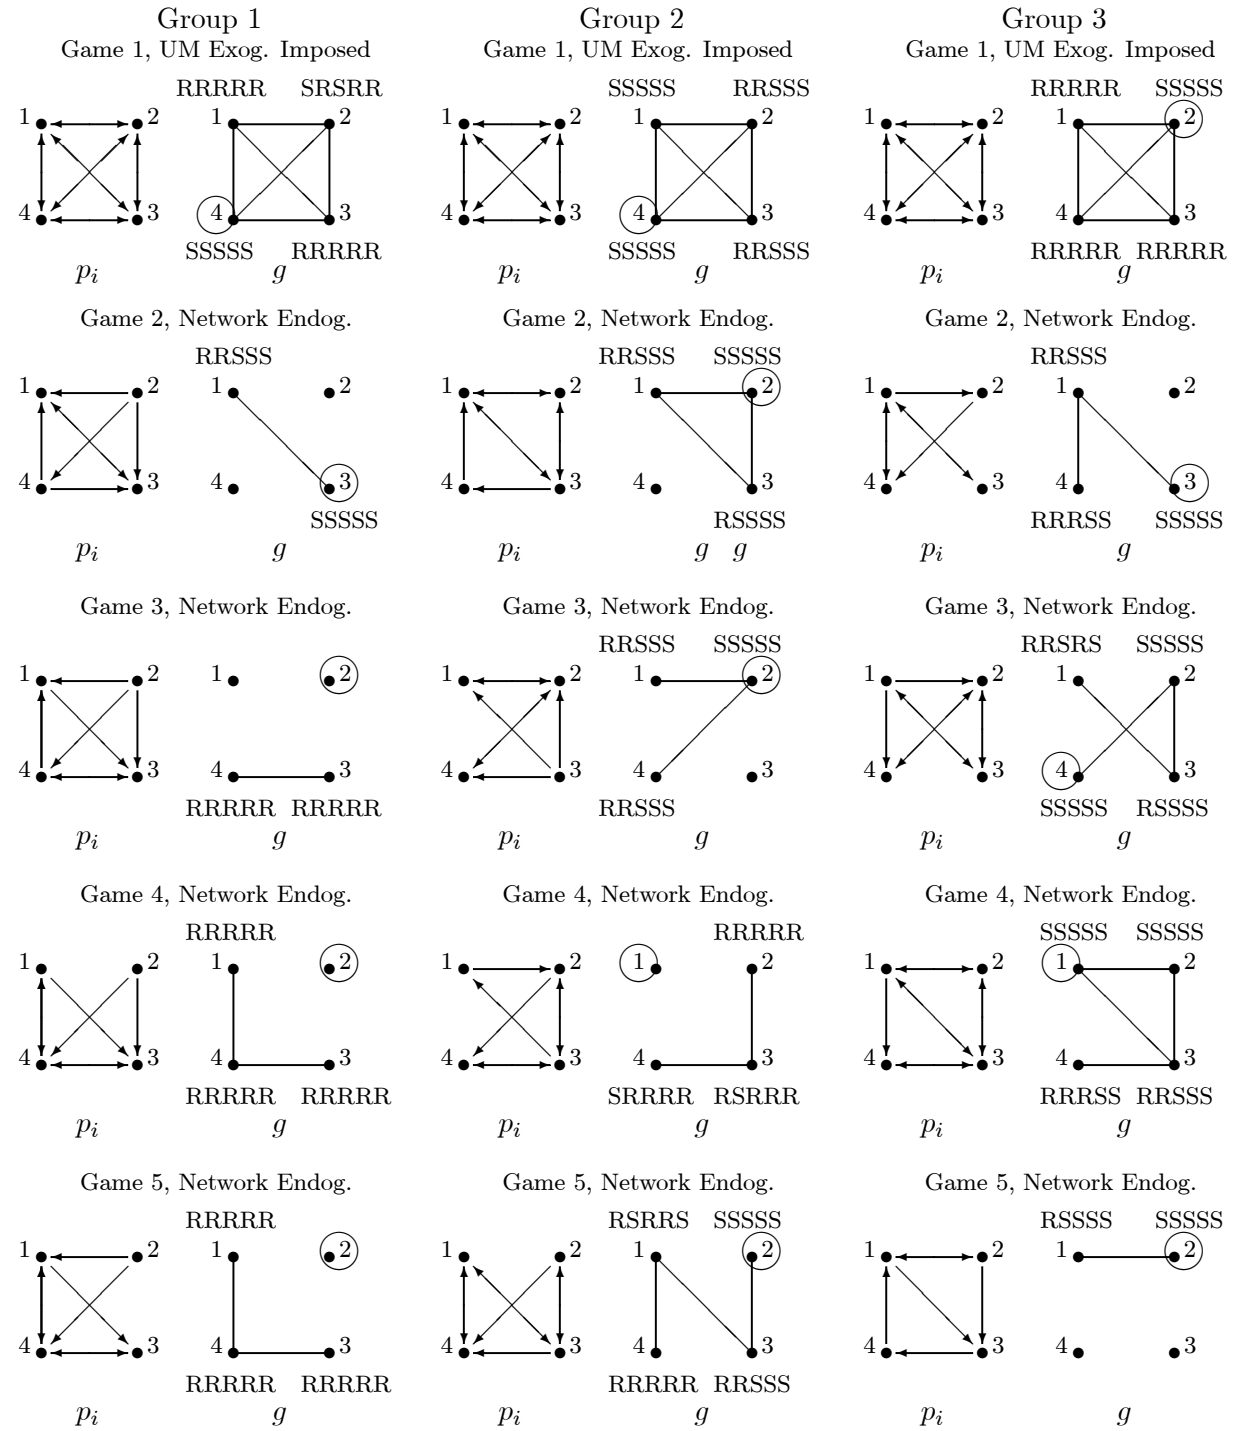




Figure UM5: Endogenous Network Treatment, 4 Groups Initially in UM Play 5 Games



Freq. of Best Response Behavior by All Unshocked Players in Each Game



Freq. of Best Response Behavior by All Unshocked Players in Each Game



Freq. of Best Response Behavior by All Unshocked Players in Each Game









Figure UM5: Continued, Endogenous Network Treatment, 4 Groups Initially in UM Play 5 Games



Freq. of Best Response Behavior by All Unshocked Players in Each Game




Figure M9.a Endogenous Network Treatment, Group 1 Initially in M Plays 9 Games


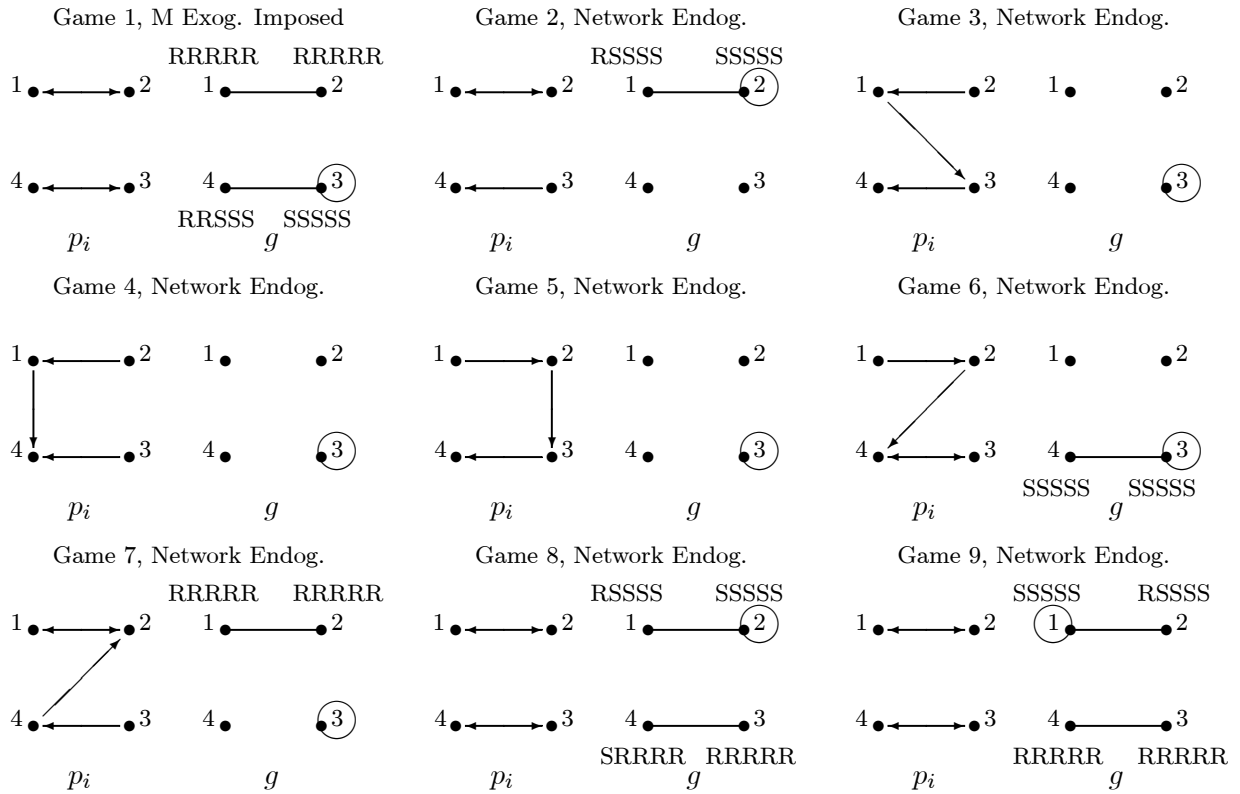


Figure M9.b Endogenous Network Treatment, Group 2 Initially in M Plays 9 Games


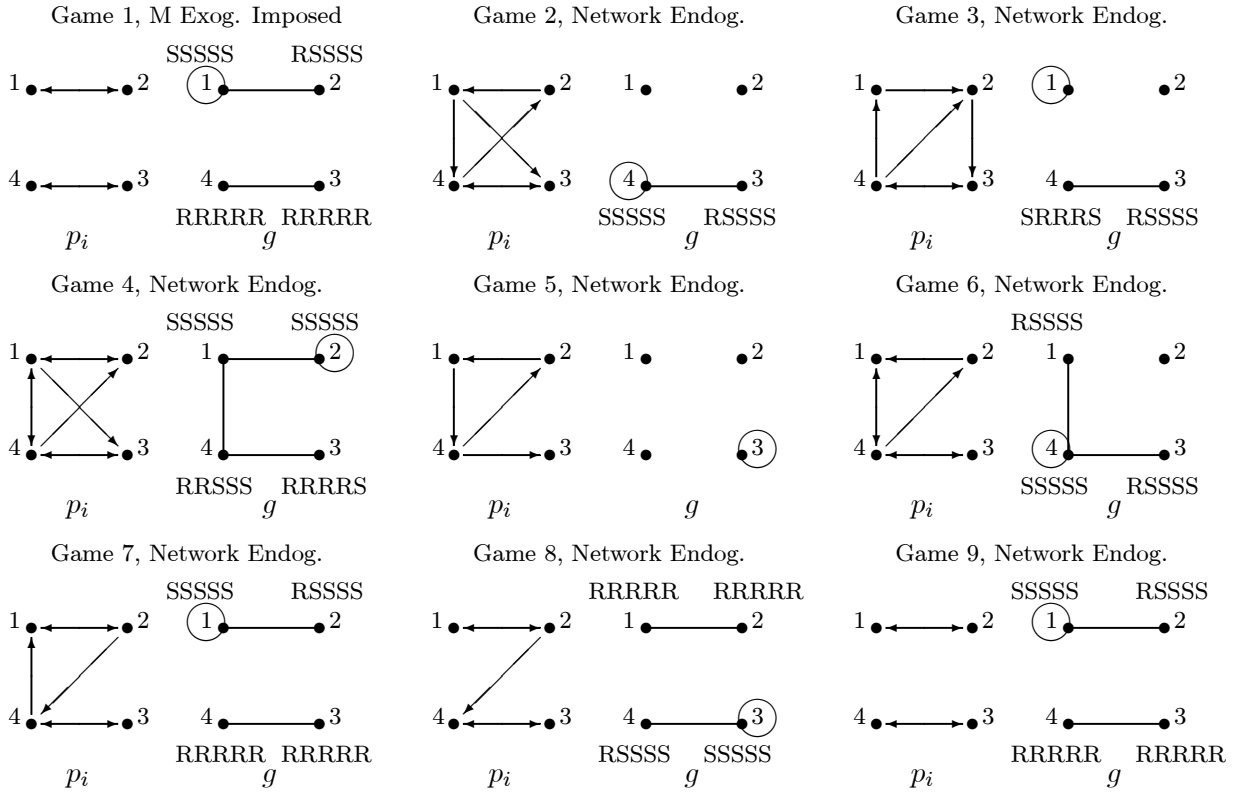


Figure M9.c Endogenous Network Treatment, Group 3 Initially in M Plays 9 Games





Figure M9.d Endogenous Network Treatment, Group 4 Initially in M Plays 9 Games


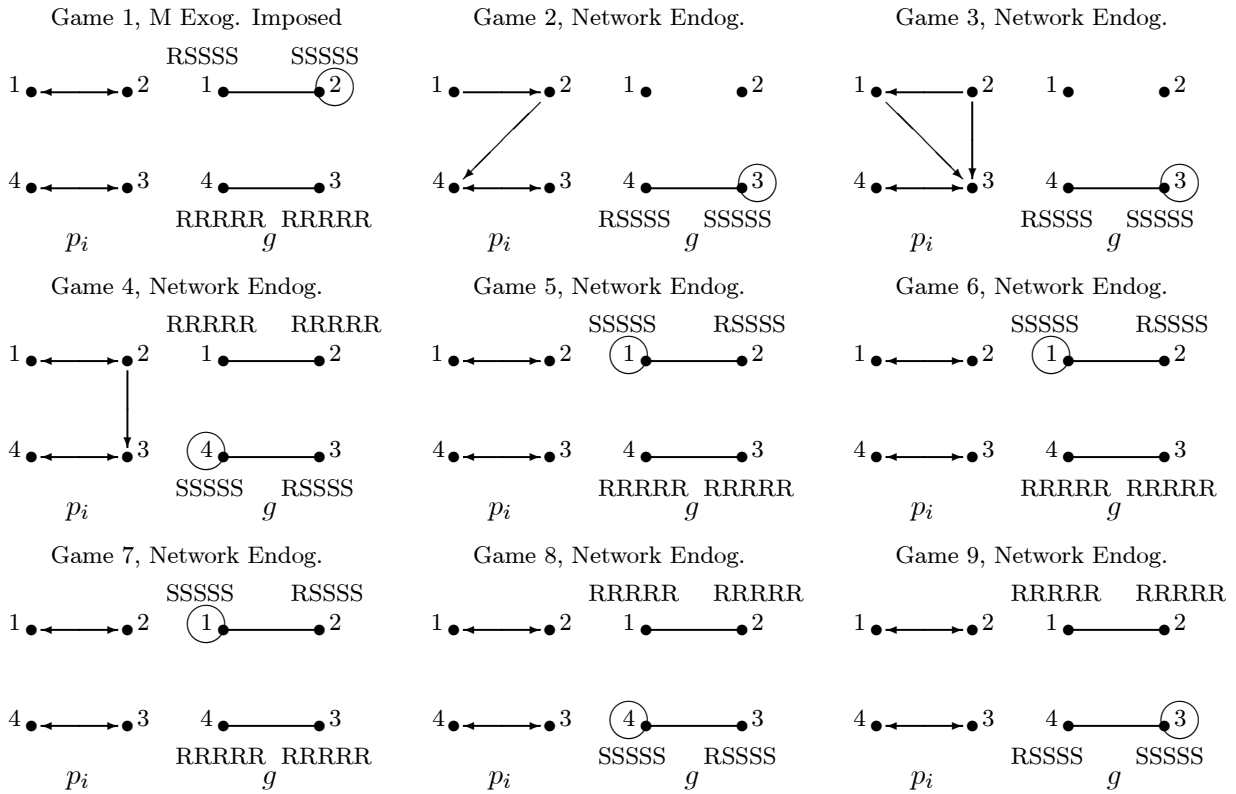


Figure LI9.a Endogenous Network Treatment, Group 1 Initially in LI Plays 9 Games






Figure LI9.b Endogenous Network Treatment, Group 2 Initially in LI Plays 9 Games



Freq. of Best Response Behavior by All Unshocked Players in Each Game




Figure LI9.c Endogenous Network Treatment, Group 3 Initially in LI Plays 9 Games





Figure LI9.d Endogenous Network Treatment, Group 4 Initially in LI Plays 9 Games


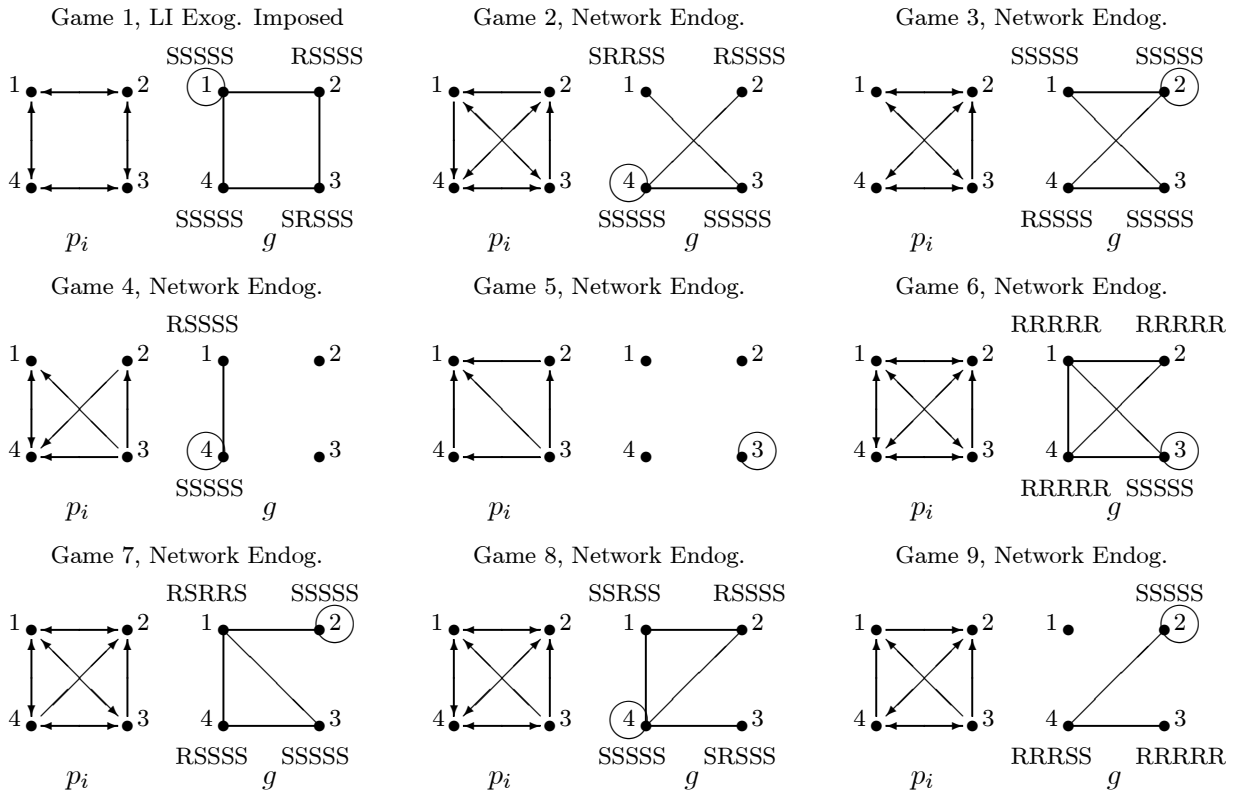


Figure UM9.a Endogenous Network Treatment, Group 1 Initially in UM Plays 9 Games


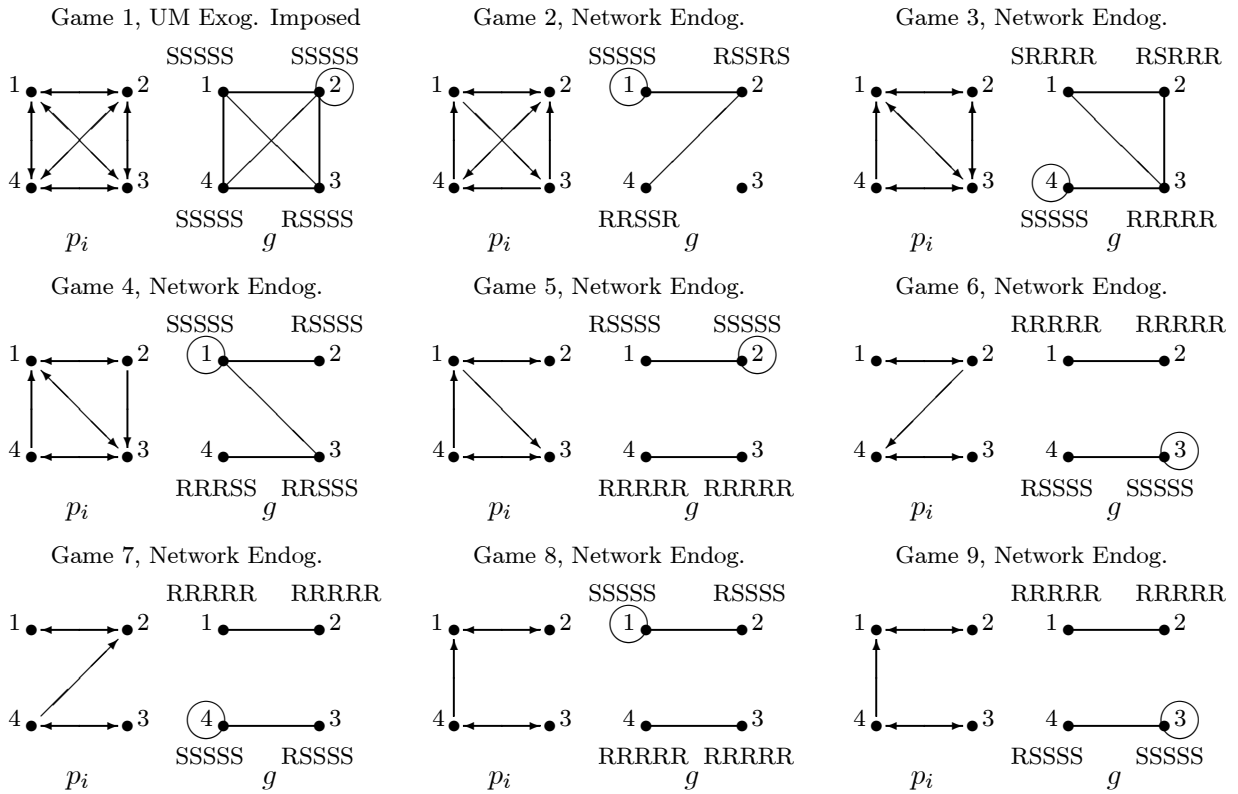


Figure UM9.b Endogenous Network Treatment, Group 2 Initially in UM Plays 9 Games






Figure UM9.c Endogenous Network Treatment, Group 3 Initially in UM Plays 9 Games



Freq. of Best Response Behavior by All Unshocked Players in Each Game






Figure UM9.d Endogenous Network Treatment, Group 4 Initially in UM Plays 9 Games



Freq. of Best Response Behavior by All Unshocked Players in Each Game



References

- [1] Allen, F. and D. Gale (2000): “Financial Contagion,” *Journal of Political Economy*, 108, 1-33
- [2] Bala, V. and S. Goyal (2000): “A Noncooperative Model of Network Formation,” *Econometrica*, 68, 1181-1229.
- [3] Berninghaus, S.K., K-M. Ehrhart and C. Keser (2002): “Conventions and Local Interaction Structures: Experimental Evidence,” *Games and Economic Behavior*, 39, 177-205.
- [4] Callander, S., and C.R. Plott (2002): “Networks: An Experimental Study,” Social Science Working Paper No. 1156, California Institute of Technology.
- [5] Charness, G., M. Corominas-Bosch and G.R. Frechette (2001): “Bargaining and Network Structure: An Experiment,” working paper, U.C. Santa Barbara, Universitat Pompeu Fabra and The Ohio State University.
- [6] Corbae, D. and J. Duffy (2002): “Experiments with Network Economies,” working paper, July, available at <http://www.pitt.edu/~jduffy/networks/>
- [7] Diamond, D. and P. Dybvig (1983): “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy* 91, 401-19.
- [8] Ellison, G. (1993): “Learning, Local Interaction, and Coordination,” *Econometrica* 61, 1047-1071.
- [9] Fudenberg, D. and J. Tirole (1991): *Game Theory*, Cambridge, MA: MIT Press.
- [10] Harsanyi, J. and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, Cambridge, MA: MIT Press.
- [11] Hauk, E. and R. Nagel (2001): “Choice of Partners in Multiple Two-Person Prisoner’s Dilemma Games,” *Journal of Conflict Resolution* 45, 770–793.
- [12] Jackson, M. (2003): “A Survey of Models of Network Formation: Stability and Efficiency,” working paper, California Institute of Technology.
- [13] Jackson, M. and A. Watts (2001): “The Evolution of Social and Economic Networks,” forthcoming *Journal of Economic Theory*.
- [14] Jackson, M. and A. Wolinsky (1996): “A Strategic Model of Social and Economic Networks,” *Journal of Economic Theory*, 71,44-74.
- [15] Kandori, M., G. Mailath, and R. Rob (1993): “Learning, Mutation, and Long-Run Equilibria in Games,” *Econometrica*, 61,29-56.
- [16] Keser, C., K-M. Ehrhart, and S.K. Berninghaus (1998): “Coordination and Local Interaction: Experimental Evidence,” *Economics Letters*, 58, 269-275.
- [17] Morris, S. (1997): “Interaction Games: A Unified Analysis of Incomplete Information, Local Interaction, and Random Matching,” Santa Fe Institute Working Paper #97-08-072 E.

- [18] Morris, S. (2000): “Contagion,” *Review of Economic Studies*, 67,57-78.
- [19] Morris, S. and Shin, H. S. (2002): “Global Games: Theory and Applications” working paper, Yale University.
- [20] Tesfatsion, L. (1997): “A Trade Network Game with Endogenous Partner Selection,” in H. Amman et al. (eds.), *Computational Approaches to Economic Problems*, (Boston: Kluwer), 249–269.
- [21] Young, H.P. (1993): “The Evolution of Conventions,” *Econometrica*, 61, 57-84.
- [22] Young, H.P. (1998): *Individual Strategy and Social Structure*, Princeton: Princeton University Press.




Figure 3: Illustration of all unilateral deviations when $I = 4$

Your Player ID is: 1. Your Player Economy is: A.

Which players do you want to propose a link to?
(Proposing a link is free)

Player 2 Player 3 Player 4

History of action choices

Game	Round	My Choice	Neighbor Choices	My Payoff

Figure 4: Illustration of the Link Formation Screen

Figure 5: Illustration of Unshocked Player's Decision Screen

Your Player ID is: 3. Your Player Economy is: A.

1 2
4 3

PAYOFF
Your Neighbors

	2X	1X 1Y	2Y
Your Choice	X 60.00	30.00	0.00
Y	35.00	35.00	35.00

Your Choice

You have recieved a payoff shock. You must chose action Y.
The program has already made this choice for you.

Results

Game	Round	My Choice	Neighbor Choices	My Payoff
1	4	Y	2Y4Y	35.00
1	3	Y	2Y4Y	35.00
1	2	Y	2Y4Y	35.00
1	1	Y	2X4X	35.00

Figure 6: Illustration of Shocked Player's Screen

Link Frequencies: Groups Beginning in M

Game No.	Frequency of Players With				No. of Obs.	K-S Test Reject H_0 ?
	0 Links	1 Link	2 Links	3 Links		
2	0.41	0.50	0.09	0.00	32	Y ($p < .01$)
3	0.41	0.56	0.03	0.00	32	Y ($p < .01$)
4	0.16	0.75	0.09	0.00	32	N
5	0.25	0.69	0.06	0.00	32	Y ($p < .05$)
Mean (2-5)	0.30	0.63	0.07	0.00	128	Y ($p < .01$)
6	0.31	0.63	0.06	0.00	16	Y ($p < .10$)
7	0.13	0.88	0.00	0.00	16	N
8	0.00	1.00	0.00	0.00	16	N
9	0.00	1.00	0.00	0.00	16	N
Mean (6-9)	0.11	0.88	0.02	0.00	64	N
Mean (2-9)	0.21	0.75	0.04	0.00	192	

Link Frequencies: Groups Beginning in LI

Game No.	Frequency of Players With				No. of Obs.	K-S Test Reject H_0 ?
	0 Links	1 Link	2 Links	3 Links		
2	0.25	0.50	0.22	0.03	32	Y ($p < .01$)
3	0.13	0.38	0.44	0.06	32	Y ($p < .01$)
4	0.25	0.59	0.13	0.03	32	Y ($p < .01$)
5	0.19	0.47	0.28	0.06	32	Y ($p < .01$)
Mean (2-5)	0.20	0.48	0.27	0.05	128	Y ($p < .01$)
6	0.06	0.56	0.25	0.13	16	Y ($p < .01$)
7	0.06	0.50	0.38	0.06	16	Y ($p < .01$)
8	0.06	0.50	0.38	0.06	16	Y ($p < .01$)
9	0.13	0.50	0.25	0.13	16	Y ($p < .01$)
Mean (6-9)	0.08	0.52	0.31	0.09	64	Y ($p < .01$)
Mean (2-9)	0.14	0.50	0.29	0.07	192	

Link Frequencies: Groups Beginning in UM

Game No.	Frequency of Players With				No. of Obs.	K-S Test Reject H_0 ?
	0 Links	1 Link	2 Links	3 Links		
2	0.22	0.22	0.47	0.09	32	Y ($p < .01$)
3	0.19	0.34	0.38	0.09	32	Y ($p < .01$)
4	0.16	0.41	0.38	0.06	32	Y ($p < .01$)
5	0.16	0.56	0.22	0.06	32	Y ($p < .01$)
Mean (2-5)	0.18	0.38	0.36	0.08	128	Y ($p < .01$)
6	0.31	0.38	0.31	0.00	16	Y ($p < .01$)
7	0.13	0.56	0.25	0.06	16	Y ($p < .01$)
8	0.06	0.50	0.31	0.13	16	Y ($p < .01$)
9	0.06	0.50	0.31	0.13	16	Y ($p < .01$)
Mean (6-9)	0.14	0.48	0.30	0.08	64	Y ($p < .01$)
Mean (2-9)	0.16	0.43	0.33	0.08	192	

Table 1: Link Frequencies by Game and Across Treatments




Figure 7: Link Frequencies by Treatment: Groups Starting Out in M, LI or UM

Proposal Frequencies: Groups Beginning in M

Game No.	Frequency of Players Proposing				No. of Obs.	χ^2 Test Reject H_0 ?
	0 Links	1 Link	2 Links	3 Links		
2	0.03	0.63	0.31	0.03	32	Y ($p < .001$)
3	0.03	0.63	0.22	0.13	32	Y ($p < .001$)
4	0.03	0.63	0.22	0.13	32	Y ($p < .001$)
5	0.06	0.66	0.16	0.13	32	Y ($p < .001$)
Mean (2-5)	0.04	0.63	0.23	0.10	128	Y ($p < .001$)
6	0.00	0.81	0.13	0.06	16	Y ($p < .001$)
7	0.00	0.81	0.19	0.00	16	Y ($p < .001$)
8	0.00	0.81	0.13	0.06	16	Y ($p < .001$)
9	0.00	0.94	0.00	0.06	16	Y ($p < .001$)
Mean (6-9)	0.00	0.84	0.11	0.05	64	Y ($p < .001$)
Mean (2-9)	0.02	0.74	0.17	0.07	192	Y ($p < .001$)

Proposal Frequencies: Groups Beginning in LI

Game No.	Frequency of Players Proposing				No. of Obs.	χ^2 Test Reject H_0 ?
	0 Links	1 Link	2 Links	3 Links		
2	0.03	0.38	0.38	0.22	32	N
3	0.03	0.34	0.34	0.28	32	N
4	0.03	0.44	0.28	0.25	32	N
5	0.06	0.38	0.28	0.28	32	N
Mean (2-5)	0.04	0.38	0.32	0.26	128	N
6	0.06	0.31	0.19	0.44	16	N
7	0.00	0.44	0.13	0.44	16	N
8	0.00	0.44	0.25	0.31	16	N
9	0.00	0.50	0.19	0.31	16	N
Mean (6-9)	0.02	0.42	0.19	0.38	64	Y ($p < .05$)
Mean (2-9)	0.03	0.40	0.25	0.32	192	Y ($p < .10$)

Proposal Frequencies: Groups Beginning in UM

Game No.	Frequency of Players Proposing				No. of Obs.	χ^2 Test Reject H_0 ?
	0 Links	1 Link	2 Links	3 Links		
2	0.03	0.22	0.34	0.41	32	N
3	0.00	0.31	0.34	0.34	32	N
4	0.03	0.31	0.41	0.25	32	N
5	0.03	0.31	0.38	0.28	32	N
Mean (2-5)	0.02	0.29	0.37	0.32	128	N
6	0.06	0.56	0.13	0.25	16	Y ($p < .10$)
7	0.06	0.38	0.25	0.31	16	N
8	0.06	0.31	0.25	0.38	16	N
9	0.06	0.31	0.31	0.31	16	N
Mean (6-9)	0.06	0.39	0.23	0.31	64	N
Mean (2-9)	0.04	0.34	0.30	0.32	192	N

Table 2: Link Proposals: Frequencies by Game and Across Treatments

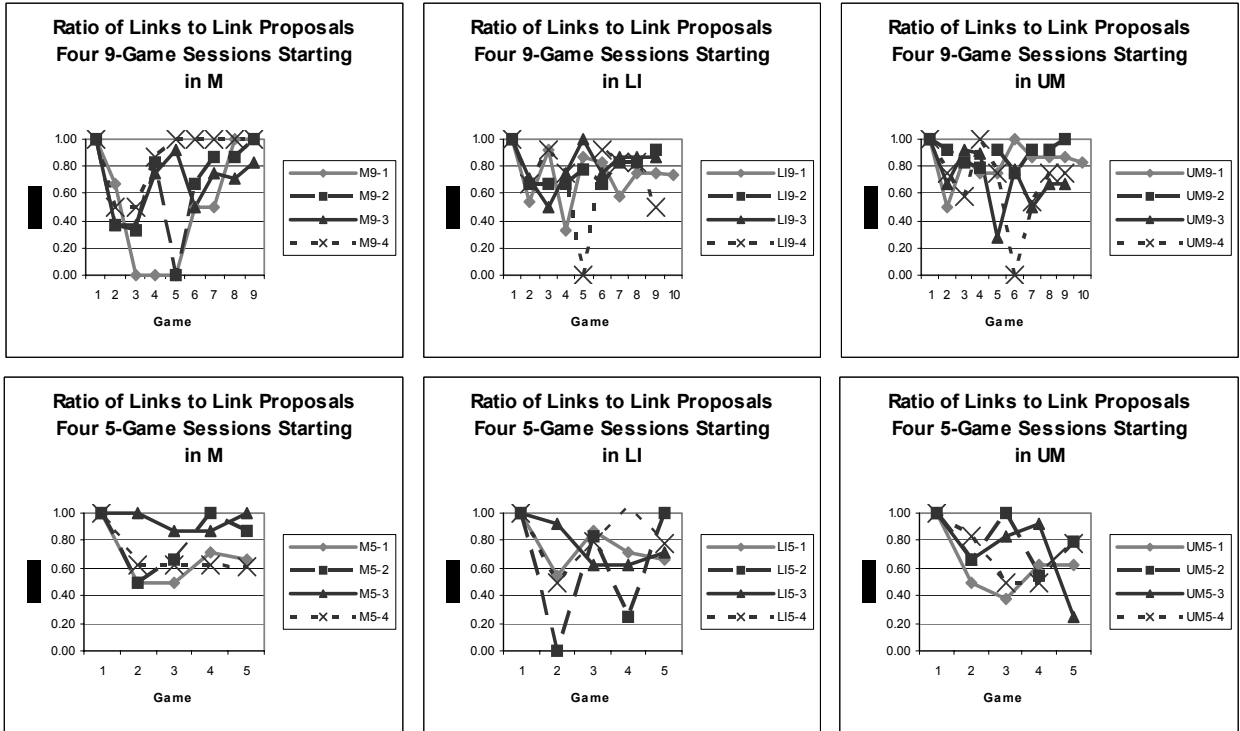


Figure 8: Ratio of Links to Link Proposals: All treatments




Figure 9: Frequency of Play of Action R in the First Round of the Second Stage Investment Game





Figure 10: Frequency of Best Response Play: All Treatments