

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

This Draft: October 1, 2007

Abstract

We examine how groups of agents form trading networks in the presence of idiosyncratic risk and the possibility of contagion. Specifically, four agents play a two-stage finite repeated game. In the first stage, the network structure is endogenously determined through a noncooperative proposal game. In the second stage, agents play multiple rounds of a coordination game against all of their chosen ‘neighbors’ after the realization of a payoff relevant shock. While parsimonious, our four agent environment is rich enough to capture all of the important interaction structures in the networks literature: bilateral (marriage), local interaction, star, and uniform matching. Consistent with our theory, marriage networks are the most frequent and stable network structures in our experiments. We find that payoff efficiency is around 90 percent of the ex-ante, payoff dominant strategies and the distribution of network structures is significantly different from that which would result from random play.

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

JEL Codes: C72, C92, D85.

*We wish to thank Andreas Blume, Gary Charness, Ken Hendricks, David Levine, Stephen Morris, Jack Ochs, Thomas Palfrey, Dale Stahl, Max Stinchcombe, Anne van den Nouweland 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, ITAM, IUPUI, New York University, Oregon, Texas A&M, UC Irvine 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

Much of economic activity occurs not through centralized market mechanisms but rather via networks of agents. In this paper, we study how networks are formed in the face of idiosyncratic risk and the possibility of contagion. We design an experiment to test the predictions of our model and we find support for our model’s main predictions.

Our model consists of a group of four agents who repeat a two-stage game a finite number of times. All of our theoretical results can be extended directly to larger groups of agents, but four is the minimal number of agents required to admit the study of all of the networks that have appeared in the economic networks literature: bilateral, local interaction, star, and complete or uniform matching networks. In the first stage, the four players choose network links by playing Myerson’s (1991) simultaneous, noncooperative link formation game. The resulting network specifies the players with whom each player will interact – each player’s “neighbors”. In the second stage, each agent plays multiple rounds of a non-cooperative stag-hunt coordination game against the n neighbors in his network, earning an average payoff in each round that depends on his choice and the choices of his neighbors. N-player Stag hunt games, which have been studied by Carlsson and van Damme [6], characterize a wide variety of economic and social situations (including team production, public goods problems, and Keynesian coordination failures). One example is that of a household either consuming its own production (hunting safe-but-low-payoff hare) or going to a market place to trade its own production for more desired products (hunting riskier-but-higher-payoff stag). The attractiveness of the latter strategy depends on the set of other traders at the market and any fundamental uncertainty that affects market trade. Our paper endogenizes (in the first stage) the set of traders with whom the household interacts in the market place (in the second stage). Since we are interested in how agents form trading networks in the presence of idiosyncratic risk and potential crises (i.e. a contagious spread of the “bad” equilibrium), the stag hunt structure is a natural framework in which to work. We introduce idiosyncratic risk and the possibility of contagion into the stag-hunt structure by assuming that the action set of one randomly drawn agent in the economy is constrained to include only the payoff dominated action. The question then is how to form a network which insures against bad outcomes.

An example of a network model of risk sharing with idiosyncratic shocks and the possibility of multiple equilibria is Allen and Gale’s [1] model of financial contagion. They construct a network version of Diamond and Dybvig’s [11] banking model, which itself can be interpreted as an N-player Stag Hunt game: players simultaneously decide whether or not to run on a bank. In Allen and Gale’s paper there are four “regions” composed of ex-ante identical agents who receive unobservable preference shocks. While there is no aggregate uncertainty, the fractions of patient and impatient agents vary across the four regions so there is the potential for risk sharing between regions experiencing high versus low demand for liquidity. The transfers supporting an insurance arrangement depend on the network structure that links the four regions. The authors show that for incomplete (or what we call “local interaction”) networks a collapse in one region spreads to the other regions (i.e. a financial contagion), but this outcome does not occur in complete networks (or what we call “uniform matching”). The authors also show that a disconnected incomplete market structure (or what we call “marriage”) can implement a first best solution. While there are many differences between our framework and that in Allen and Gale, the primary one is that we endogenize the network structure while they simply take the network structure as given. Obviously the types of risk sharing networks that banks engage upon are not exogenous. For instance, if incomplete networks are more susceptible to contagions, why would they form in the first place? In this paper we seek to address this type of question by examining endogenous network formation, when contagions are possible. For the environment and parameterization we choose in this paper, we show that an efficient arrangement corresponds to Allen and Gale’s disconnected incomplete market structure or “marriage network” as it serves to minimize

the spread of contagion.

For the environment we consider, we show in a technical appendix [9] that the only strict, symmetric, ex-ante efficient, perfect Bayesian equilibrium (PBE) network is a bilateral or “marriage” network (where the four players form two pairs, 1 link each). Other symmetric and most other asymmetric networks that are possible in our environment are not *strict* PBE. We choose to focus on ex-ante, payoff dominant, perfect Bayesian strategies as an efficient benchmark that a planner subject to the same information restrictions would choose to implement. Since there is the possibility of multiple equilibria, we analyze this stability prediction in a number of experimental sessions, where the main treatment variable consists of the symmetric network structure that is initially exogenously imposed on subjects in the first two-stage game; in subsequent two-stage games, players are free to submit link proposals. Specifically, we ask whether subjects who start out playing the stag-hunt game in a certain network, say “uniform matching” (where each player is linked to every other player) decide to submit proposals in the subsequent game so as to re-implement that same network structure. This stability prediction is predicated on the assumption that players play the second stage, stag-hunt game in accordance with perfect Bayesian equilibrium predictions.

Our findings suggest that players frequently do play according to the exante payoff dominant perfect Bayesian equilibrium in the second-stage game. The frequencies with which action choices accord with these predictions exceed 75 percent, and subjects earn around 90 percent of the payoffs they could achieve if they played according to the exante payoff dominant perfect Bayesian equilibrium, that is, payoff efficiency is high. Regarding network formation in the first-stage proposal game, we find that the distribution of endogenously determined network types is significantly different from that which would be implied by random proposal choices. When subjects start out in marriage networks and are free to form links, they choose to implement marriage networks 77 percent of the time. However, when subjects start out in local interaction or uniform matching networks and are free to form links they choose to re-implement those network structures less than 3 percent of the time and they succeed in forming marriage networks 25-30 percent of the time.¹ Once a marriage network was endogenously formed, it was sustained for the duration of an experimental session. We regard the latter findings as strong support for our stability prediction, namely that marriage networks are the only stable symmetric networks in our environment.

The rest of the paper is organized as follows. After reviewing the literature in Section 2, we describe an economic environment (matching and productive technologies, as well as preferences and information structure) in Section 3. Section 4 describes the ex-ante, payoff dominant perfect Bayesian equilibrium for all possible network structures. Under the assumption that subjects hold beliefs that play in the second stage conforms to the above strategies, we state in Proposition 1 that the only network which is strictly immune to unilateral deviations is one where players form bilateral links, which we refer to as a “marriage” network. Proposition 1 provides us with a stark, testable hypothesis for the experiments which we take up in Section 5. Since there are multiple equilibria in both the proposal and stag-hunt stage games, we examine the data to determine whether subjects are playing according to the strategies considered in Proposition 1. Our experimental findings are fairly consistent with the predictions of the theory.

¹There was obvious experimentation with different network structures along the path and if we had let the subjects play longer perhaps the number of subsequent marriage networks would have grown even more. Indeed, we observe that the frequency of marriage networks generally increases as we increase the number of games played in a session from 5 to 9.

2 Literature

The literature on network economies is voluminous and we do not attempt to summarize it here. 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 [12], Kandori, Mailath, and Rob [24], Morris [28], and Young [33]. Papers that follow the second approach are Bala and Goyal [2], Jackson and Watts [21], Jackson and Wolinsky [23]. Jackson [20] surveys this literature.

There is now a small and growing experimental literature examining the impact of exogenous network configurations on behavior in games and another literature that considers endogenous network formation; Kosfeld [26] surveys this literature. Most closely related to this study are several experimental studies examining *endogenous* partner selection or network formation. For instance, Hauk and Nagel [18] 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. Several authors have experimentally examined network formation with the aim of testing the predictions of versions of Bala and Goyal’s [2] model of network-formation with unilateral link formation and one- or two-way information flow. (See, e.g., Callander and Plott [5], Falk and Kosfeld [13], Berninghaus et al. [4] and Goeree et al. [16]).

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’s network game, 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.² Unlike Jackson and Wolinsky, however, we implement two-sided link formation in a *non-cooperative* game. Third, we are not simply interested in the question of which networks emerge when agents are free to propose network links; we further examine how agents play a coordination game with their network neighbors, similar to the games studied by Keser et al. [25] and Berninghaus et al. [3] given the network structure they have implemented. Indeed, our study is among the first to unify the two different experimental literatures on network games.³ Finally, in our environment, 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 carefully explore the issue of the *contagious spread* of actions as a function of network structure.⁴ Such contagious behavior may be an important consideration in the design of financial market networks, as well as in other applications. Thus our paper adds to an exciting new literature that seeks to understand financial crises such as bank runs (Schotter and Yorulmazer [31], Garratt and Keister [15]), or speculative attacks, (Heinneman et al. [19]), using laboratory experiments.

²While mutual consent strikes us as a natural rule for link formation in economic and social networks, it may not be innocuous with regard to the equilibrium subjects choose in the second stage coordination game as shown in an experimental study by Charness and Jackson [7].

³See also Jackson and Watts [22].

⁴In the absence of a payoff shock, we found that subjects nearly always choose to coordinate on the payoff dominant equilibrium of the game, regardless of network structure -see Corbae and Duffy [9]. Garratt and Keister [15] report the same finding in their bank run experiments.

3 The Environment

The basic model is of a finite sequence of two-stage games. 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 take an action with payoffs similar to an n -person version of a stag hunt coordination game.

Specifically, there is a finite set of 4 players. There are κ repetitions of two-stage play. We call the first stage the *network proposal stage* and the second stage of τ rounds the *stag-hunt stage*. Thus, there are in total $\kappa(1 + \tau)$ discrete periods of play. Let t_p denote the times at which network proposal are made and t_a denote the times at which action choices are made.

3.1 Matching technology

One can think of economic interactions as being determined by a matching technology that assigns a weight to the link between any two agents i and j in a network. While we define networks in much the same way as Jackson and Wolinsky [23], we implement the network using a simultaneous, noncooperative game á la Myerson (1991, p. 448) as opposed to using cooperative, coalitional solution concepts such as pairwise or strong stability.⁵ We adopt a *non-cooperative* network formation game for several reasons. First, as the second-stage of our game involves the play of the non-cooperative Stag hunt game, it would seem inconsistent to mix cooperative and noncooperative solution concepts. Second, even if we did apply a cooperative solution concept to the first-stage network formation problem, implementation in the laboratory would be complex and problematic; for instance, we would have to allow pairs of players (pairwise stability) or larger coalitions (strong stability) to communicate with one another with regard to which, of the many network structures available in our environment, they were willing to collectively implement. Finally, while we recognize that Myerson’s non-cooperative network formation game can yield a multiplicity of equilibria, it *is* empirically interesting to ask how that multiplicity problem is overcome (in our theoretical analysis we propose ex-ante payoff efficiency as a refinement). It would seem that such coordination issues are a natural problem in group formation and we don’t want to gloss over this issue by appealing to coalitional stability considerations.

In our first stage, network proposal game, each agent i simultaneously chooses whether or not to link to each of the other agents in the economy. In particular, letting \mathcal{I} denote the set of four agents in the economy, agent i takes a *network proposal action* which is a 3-tuple $p_t^i = (p_{\{\mathcal{I} \setminus i, t\}}^i) \in P_t^i = \{0, 1\}^3$ 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 $t = t_p$. A *link* at time $t = t_p$ 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, $g_t = \{(i, j) \in \mathcal{I} : p_{j,t_p}^i p_{i,t_p}^j = 1\} \in \Gamma$, where Γ is the set of all possible networks.⁶ We assume that the network remains unchanged during the τ rounds of play in the stag-hunt stage until the next set of proposals are made. We define the *neighborhood* of agent i in network g_t to be the set of all agents to whom he/she is linked and denote it $N^i(g_t) = \{j : p_{j,t_p}^i p_{i,t_p}^j = 1, j \neq i\}$. If $N^i(g_t) = \emptyset$, then agent i is in autarky. The number of neighbors of agent i is simply the cardinality of $N^i(g_t)$ and is denoted $n^i(g_t)$.

⁵On this topic, see the discussion in Jackson [20]. Jackson and Watts[21] provide a dynamic version of network formation.

⁶In the earlier version of our paper [8] we included costly network formation.

The set of feasible graphs for our 4 agent economy is shown in Figure 1.⁷ The only complete graph is UM , a version of the *uniform matching* model of Kandori, Mailath, and Rob [24] or Young [33] where each agent has $n^i(g^{UM}) = 3$ direct links to all other agents in the economy. The graph LI is the 4-agent version of a *local interaction* model such as that studied by Ellison [12]. 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 just above, in the first row. The graph, $UM - LI$ has agents 1 and 3 in LI neighborhoods while agents 2 and 4 are in UM neighborhoods. The 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 (see, e.g. Jackson and Watts [21]) 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.




Figure 1: Illustration of all symmetric and asymmetric graph forms for 4-player groups

3.2 Payoffs

In each of the τ rounds of the second stage stag-hunt game, players may take one of two possible actions $\{X, Y\}$ that are payoff relevant. The action set for each agent in the stag-hunt game 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 $A^i(1) = \{Y\}$ and $A^i(0) = \{X, Y\}$ at $t = t_a$. The shocks, which arrive prior to the action being taken, provide the experimenter with some

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 .

control over the decisions in the stag-hunt stage. For instance, the shock can be used to randomly assign one agent in the economy a tremble in order to possibly start a contagion. While there are numerous stochastic processes that we can implement, we will focus on one particularly simple version.⁸ Specifically, in the first round of the second stage (i.e. $t = t_p + 1$), one agent (say i) receives $\omega_t^i = 1$ while all other agents receive $\omega_t^{-i} = 0$. Agents maintain their type in all τ rounds of the second stage (i.e. $t = t_a \in \{t_p + 1, \dots, t_p + \tau\}$). Simply put, before actions are taken in the first round of the stag-hunt stage, one out of the 4 agents will learn that he/she must take action Y for all τ rounds.

Taking as given the network g_t from the first stage, before agent i interacts with his neighbors and after he learns the state of his action set (i.e. ω_t^i), he takes action X or Y which is implemented in all of player i 's interactions with other players in his neighborhood $j \in N^i(g_t)$. The assumption, that actions cannot be made j contingent, is what makes network structure matter in our environment. Agent type contingencies remain possible in the first stage, network proposal game.

We assume agent i 's payoffs from his action choices, denoted $u^i(a_t^i, a_t^j)$, $j \in N^i(g_t)$, are given by:

$$\begin{aligned} u^i(X, X) &= a & u^i(X, Y) &= c \\ u^i(Y, X) &= b & u^i(Y, Y) &= b \\ u^i(a^i, \emptyset) &= d & \text{if } N^i(g_t) &= \emptyset \end{aligned} \tag{1}$$

where $a > b > c$. The last line of (1) simply says that if an agent is in autarky, he receives payoff $d \leq b$ independent of his actions. The assumption that $b \geq d$ will ensure that participation is weakly optimal. If there were only 2 players, the payoffs are simply those of a pure stag-hunt game with two pure strategy equilibria (all- X and all- Y) and a mixed strategy. Furthermore, while all- X is payoff dominant (for players receiving the favorable payoff shock), if $b > (a + c)/2$ then all- Y is the equilibrium with lowest *risk factor*.⁹ If $N^i(g_t) \neq \emptyset$, agent i 's payoff from playing action a^i in any round of the second stage game is given as the weighted sum of all the payoffs associated with actions taken by his neighbors,

$$\sum_{j \in N^i(g_t)} \frac{1}{n^i(g_t)} u^i(a^i, a^j). \tag{2}$$

and payoff d otherwise.

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 Y , local interaction would be preferred to marriage simply because it yields $2a$ rather than a . We also chose a simple average of payoffs, rather than some other weighting measure like a minimum function, since this corresponds more closely to the idea of maximizing expected profits in a regional banking problem as discussed in the introduction.¹⁰

3.3 Information and the Timing of Events

We assume that each agent knows g_t .¹¹ Since each agent interacts with every other agent in his neighborhood, we assume that in any round t_a of the second stage game since the latest network was

⁸In our earlier version [8] we consider other stochastic processes for ω_t^i . For instance, we show in a proposition that if $\omega_t^i = 1$ for one i and 0 for all others and this joint process is iid across time, then network structure does not matter.

⁹See Young [34] (p. 67) for a definition of risk factor. If the game were a perfectly symmetric 2×2 game, all- Y would be more familiar as the *risk dominant* equilibrium.

¹⁰A minimum weighting function under our parameterization would bias our results towards forming smaller networks.

¹¹An interesting extension of our work would take the decentralized nature of interactions literally and 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). While this is an interesting theoretical problem, we believe the resulting inference problem is virtually impossible for our human subjects to process and so leave this for future research.

formed in the proposal stage (i.e. for $t = t_a \in \{t_p + 1, \dots, t_p + \tau\}$), agent i knows the actions that have been played by their neighbors in all previous rounds (as well as his own). Consistent with the decentralized nature of agent interaction, we assume that players do not know the actions that have been played by agents outside of their neighborhood. That is, at t_a agent i knows $\{a_t^i, a_t^j\}_{j \in N^i(g_t), t < t_a}$.

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 $\omega_{t_a}^i$ he does not necessarily know the types of others $\omega_{t_a}^{-i}$ (of course, if $\omega_{t_a}^i = 1$, then i knows $\omega_{t_a}^{-i} = 0$). This assumption is consistent with the information partition associated with a given network structure.¹² This assumption is also consistent with work on incomplete information games by Morris [27]. Figure 2 summarizes the timing of events.




Figure 2: The timing of events in each session

4 Equilibrium Predictions

The section is intended to provide a theoretical basis for a set of predictions for our experimental study. The results are for a very simple game; there is one network proposal round ($\kappa = 1$) followed by τ rounds of the stag-hunt game.¹³ Unless stated explicitly, we make the following parametric assumption.

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

The assumption that $a > b > c$ is standard in coordination games where coordinated risky (X) play yields a higher payoff than safe (Y) play. The assumption that $\frac{2a+c}{3} > b$ ensures that in neighborhoods of three players, coordinated X play yields a higher payoff than Y 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, X play is suboptimal. This assumption is necessary for contagion to get started. It is also consistent with all- Y play being risk dominant in a two player game.¹⁴

The logic of our analysis is to first characterize continuation equilibria of the stag-hunt 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 stag-hunt game

¹²Of 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.

¹³Our results on existence of equilibrium can be extended to multiple rounds of proposals ($\kappa > 1$), simply by constructing an equilibrium where agents disregard the history of the prior proposal and second stage games. Furthermore, since the shock process is independent across proposal times, such equilibria can be ex-ante payoff dominant in the full game.

¹⁴A pair of strategies is risk dominant (Harsanyi and Selten [17]) 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 X yields payoff $\frac{1}{2}a + \frac{1}{2}c$ while choice of Y yields b , then we have Y being the risk dominant action.

(e.g. all- Y), we characterize the ex-ante payoff dominant, pure strategy, symmetric Perfect Bayesian equilibrium (PBE) of the stag-hunt game, taking as given the network structure.¹⁵¹⁶ In virtually all cases, the equilibrium strategies are such that players who do not receive the shock play X 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 stag-hunt 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.¹⁷ We then use these results to construct LI and M equilibria under the assumption that agents coordinate upon payoff dominant equilibria in each stag-hunt 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, since we characterize stability for the entire set of possible networks, we have shown that M networks are the only networks that are strictly immune to unilateral deviations.

To gain intuition for why M equilibria are likely to arise, consider a social planner who weights agents equally and wants to maximize economywide utility but is subject to the same information frictions as agents. Since the planner can avoid coordination problems that can arise in both the proposal stage and stag hunt game, this amounts to choosing a network that maximizes ex-ante payoffs and is individually rational. A reader might be tempted to think the planner would implement a UM network and direct any unshocked agent to play the high level X action. Under the above parametric assumption, conditional on being in a UM network, this structure results in the highest frequency of individually rational, high level X play across the τ rounds of the stag-hunt game. On the other hand, the UM network implies that every neighborhood contains the shocked agent playing the safe, low level Y strategy and while it is individually rational for an agent to participate in such an arrangement, the ex-ante payoff over the τ rounds of the stag hunt game is $\frac{1}{4}(b\tau) + \frac{3}{4}\left(\frac{2a+c}{3}\right)\tau$. On the other hand, if the planner implements two marriages and directs unshocked agents to play the high level X action unless their partner plays Y , such an arrangement implies that any spread of the low level Y strategy is contained within a subset of the population. The ex-ante payoff of this arrangement over the τ rounds of the stag hunt game is $\frac{1}{4}(b\tau) + \frac{1}{4}(c + b(\tau - 1)) + \frac{1}{2}(a\tau)$, which exceeds

¹⁵In the extended version of the paper [8], we provide an explicit definition of the equilibrium concept we use. In particular, we study *perfect Bayesian equilibria*, which are behavior strategy–belief pairs such that (i), given beliefs, the behavior strategies are a best response to all others strategies after any possible history and (ii), wherever possible, posteriors satisfy Bayes’ rule.

¹⁶By 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 $p_{i+2,0}^i = 0$, mod 4.

¹⁷The 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, in a Technical Appendix to this paper (Corbae and Duffy [9]), there are only two cases of 11 we analyze (lemmas 10 and 11) where this is the case.

the expected payoff from the UM arrangement by $\frac{(b-c)(\tau-1)}{4}$ associated with the containment of the contagion to a single “bad marriage” which involves the shocked player and his individually rational partner playing the low level Y strategy in rounds 2 through τ of the stag-hunt game.

In a technical appendix to this paper (see Corbae and Duffy [9]), we analyze this problem in two steps. The first step establishes certain results of the stag-hunt game under a given network structure.¹⁸ The only real issue for players in the stag-hunt game is to infer who has the shock and best respond accordingly. The results are used 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. The second step is to use the preceding results to establish predictions for play in the proposal game that determines equilibrium networks.¹⁹ These two steps are summarized in this paper by Proposition 1. 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.²⁰ 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.²¹ 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 X in the first round, players who have played X in each previous round continue to do so if their partner has played X in each previous round, and any player who has herself played Y or whose partner has played Y in some round plays Y in each subsequent round.²² Since agents are ex-ante more likely to be in an unshocked marriage and Y play invokes a Y response according to the subgame perfect strategy, in the first round it is optimal to play X until one knows whether one’s partner is the shocked player, in which case it is optimal to play Y 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.²³ 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 Y play in the subsequent rounds as dictated by the ex-ante payoff dominant PBE strategy in an LI network.²⁴ This strategy (where unshocked agents

¹⁸That is, lemmas 1 to 11 of the Technical Appendix.

¹⁹See, in particular, lemmas 12 to 20 of the Technical Appendix.

²⁰There 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 stag-hunt game under a given network structure before sending their proposals in a subsequent stage.

²¹This result is established in lemma 12 in the Technical Appendix.

²²This result is established in lemma 6 in the Technical Appendix.

²³This result is established in lemma 13 in the Technical Appendix.

²⁴These results are established in lemmas 3 and 8 in the Technical Appendix.

in an LI network play X 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 Y and hence plays Y . Notice that this result would be very different if we were using a solution concept like naive best response. Thus, the “contagious” Y 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.²⁵ This result is despite the fact that the ex-ante payoff dominant, pure strategy PBE in UM results in each unshocked agent playing X 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 X) 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 $(2a + c)/3 < a$.²⁶ The payoff gain in this event is offset by the event when either agents 2 or 4 receive the shock, in which case the ex-ante payoff dominant, pure strategy PBE in UM-LI calls for play of action Y after the first round, insulating all agents from receiving a fraction of the payoff c .²⁷ In ex-ante terms, the gains more than offset the cost so the deviation is profitable. 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.

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 $\frac{1}{4}b\tau + \frac{1}{2}a\tau + \frac{1}{4}(c + b(\tau - 1)) = \frac{2a+b+c}{4} + \frac{a+b}{2}(\tau - 1)$ from being in a marriage are strictly higher than the expected payoffs by being the middleman $\frac{1}{4}b\tau + \frac{3}{4}(\frac{2a+c}{3})\tau = \frac{2a+b+c}{4}\tau$ since in the event that he is unshocked, he provides insurance against the shock with probability one each period.²⁸ 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.²⁹ We summarize the results for all possible network configurations in the following proposition.

Proposition 1 *When τ is sufficiently large, and we restrict play in the second stage 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. The ex-ante efficient, strict PBE network is M.*

²⁵ See lemma 15 of the Technical Appendix.

²⁶ This payoff is consistent with the ex-ante payoff dominant, perfect Bayesian equilibrium strategy in UM where each unshocked agent plays X in the first round and thereafter plays X in each round in which at least three agents play X in the previous round, and plays Y 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 in the Technical Appendix.

²⁷ See lemma 4 of the Technical Appendix.

²⁸ See lemma 16 in the Technical Appendix for this result.

²⁹ See lemma 9 in the Technical Appendix.

Our restriction to an economy comprised of $I = 4$ agents makes feasible a complete characterization of all possible symmetric and asymmetric network structures. Here we discuss the sensitivity of our results to raising the number of agents while maintaining the assumption that only one out of I agents receive the shock which restricts their action set to playing Y . It can be shown that under Assumption 1, the strategies that result in the ex-ante payoff dominant 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 payoff V_M to a given agent of being in an M network is

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

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

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

where the first term is the payoff if the agent is shocked, and the second term is the payoff if he is not shocked. 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 $I = 4$ occurs in the LI network. In this case, the strategy that resulted in the ex-ante payoff dominant equilibrium generalizes as follows: each unshocked agent plays X in the first round, and then plays X either until one of his neighbors has played Y , or until he can infer that one of his neighbors will play Y in the current round, and he plays Y thereafter.³⁰ For I even with $I/2 > \tau$, the ex-ante payoff V_{LI} to a given agent of being in an LI network following this strategy is

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

where the first term is the payoff if the agent is shocked, the second term is the payoff if the shocked agent is in his neighborhood, the third term is the payoff if the shocked agent is in his neighbor's neighborhood, etc.³¹ It is straightforward to show that $V_{UM}(I) > V_{LI}(I)$ for all finite $I > 2\tau$.³²

It should 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.

³⁰This is a generalization of the strategy used in lemma 3 of the Technical Appendix.

³¹When $\tau \geq I/2$, the expected payoff from the strategy described is

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

³²Subtracting V_{LI} from V_{UM} gives $\frac{1}{I} (a - b) \sum_{i=1}^{\tau} 2(i - 1) > 0$.

5 Experimental Design and Findings

Our experimental design focuses on the stability of network structures when players are free to choose links.

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 1: in the presence of permanent shocks, the only strict pure-strategy perfect Bayesian equilibrium networks satisfying the ex-ante payoff dominance criterion in the second stage continuation game are M, M-A, and LI-A networks. Our use of the ex-ante payoff dominance criterion is an obvious benchmark equilibrium prediction. The evidence on whether subjects playing coordination games are more likely to coordinate on payoff-dominant as opposed to other (e.g., risk dominant) equilibria is mixed (see Devetag and Ortmann [10] for a survey). However, some of the more careful experimental work (Rankin, Van Huyck and Battalio [30]) suggests that payoff dominance is indeed the most relevant selection criterion.³³

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 subsequent two-stage games. Our main finding is that, consistent with the prediction of Proposition 1, only M networks appear to be stable.

5.1 Representation of Payoffs in the Stag-Hunt Game

We work with a specific parameterization for payoffs in the second stage stag-hunt 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)	X	Y
X	60	0
Y	35	35

Payoffs are shown only for the row player, i . In the case of a symmetric M network, the other player's payoffs can be inferred from such a representation. We note that for this benchmark, symmetric, 2×2 case, if players' action sets are unrestricted, there are two pure strategy Nash equilibria: all-X and all-Y. It is easily verified that all-X is the payoff dominant equilibrium, while all-Y is the risk dominant equilibrium.³⁴

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).

³³Regarding group size, there is not much evidence on $2 \times N$ coordination games where $N=4$. The closest parallel to our experimental environment is found in Berninghaus, Erhart and Keser [3]. They experimentally examine 3-player Stag Hunt games under an average payoff rule (as in our design) and report that only around 10 percent of the 3-player groups play the risk-dominant (safe action), with the rest coordinating on the payoff dominant action. Thus we believe that our group size of 4 subjects is not so large as to inhibit the play of payoff dominant strategies. See also footnote 4.

³⁴There is also a mixed strategy equilibrium to the symmetric, 2-player game where each player plays action X with probability $\frac{7}{12}$ and earns an expected payoff of 35. We focus here on pure strategy equilibria.

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)		2X	1X1Y	2Y
X		60	30	0
Y		35	35	35

where 2X means that 2 of the $j = 2$ neighbors chose X , 1X1Y means that 1 of the two neighbors chose X and the other chose Y , and 2Y means that both of the 2 neighbors chose Y . The payoffs for these outcomes are consistent with the calculation in (2).³⁵

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)		3X	2X1Y	1X2Y	3Y
X		60	40	20	0
Y		35	35	35	35

where again, the different payoff amounts reflect the weighting scheme in (2).³⁶ 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 Y . 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 two-stage games played, and in addition were awarded a fixed, \$5 participation payment. Average total earnings over all sessions involving 5 two-stage games was \$14.43 per subject (including the \$5 participation payment); these sessions lasted approximately 75 minutes. The comparable average total earnings over all sessions involving 9 two-stage games was \$22.10 per subject; these sessions lasted approximately 100 minutes.

5.2 Experimental procedures

The experiments were implemented using networked personal computers in the University of Pittsburgh Experimental Economics Laboratory. The subject pool consisted of inexperienced undergraduates, recruited from the population of undergraduates at the University of Pittsburgh.

³⁵For example, if a player with two neighbors chose X , and his two neighbors' choices were X and Y (i.e. 1X1Y), 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.

³⁶Note that, in the case where a player has 3 neighbors (and in this case *alone*), the outcome where the player and all three of his neighbors plays action X , each earning a payoff of 60, is not possible in our environment, as one player in every four player group is shocked (restricted to playing action Y) in every round. All other payoff outcomes are possible. This fact was carefully pointed out to subjects in the experimental instructions.

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.³⁷ 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 5.1. 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 Y in all rounds of the subsequent second stage 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 stag-hunt 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.³⁸ 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 $\kappa = 5$ or $\kappa = 9$ 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 stag-hunt 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 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 re-emerge 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.

³⁷Copies of the instructions used in our experiment can be viewed or downloaded at <http://www.pitt.edu/~jduffy/networks/>

³⁸The 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.

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. 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 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 *Y* 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 stag-hunt 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* 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.³⁹ In round 3, player 1 chose *Y* 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.

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 three 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, and that link proposals were costless. We chose not to attach a payoff cost to link proposals

³⁹In theory, of course, player 1 should have chosen action *Y* in round 2; the screenshots are just an illustration of what could happen.

as we did not want to create any bias in favor of networks where players have low numbers of links. On the other hand, to the extent that there is some mental/physical cost to checking boxes (making proposals), subjects would want to check the minimal number of boxes necessary to implement a desired network structure. If subjects in fact behaved in this manner, it would serve to validate our focus on equilibria with minimal proposal strategies, and we check (below) for whether the minimal proposal strategy in fact obtains in the experimental data.

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 other players in their four-player group were facing. Of course, their own payoff table was prominently featured on their 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.

5.2.1 Hypotheses

When analyzing the data, we examine two main hypotheses that underlie the lemmas which make up Proposition 1.

Hypothesis 1 *In the continuation game following the implementation of a network, subjects play according to the ex-ante, payoff dominant PBE strategies.*

Hypothesis 2 *In the proposal game, subjects implement strict-PBE networks.*

5.2.2 Experimental Findings

[Insert Figures 7, 8, and 9 here.]

Figures 7, 8 and 9 provide an illustration of the raw data collected from three of the four sessions where subjects played a total of 5 two-stage games. As noted above, in each of these sessions, one 4-player group started out in M, one in LI and one in UM, but here we have rearranged the data, so

that the results for three groups (different sessions) starting out in M are presented in Figure 7 and, analogously, the results for three groups starting out in LI or UM are presented in Figures 8 and 9.⁴⁰

In Figures 7–9, 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 proposal 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 (the one forced to play action Y 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 stag-hunt game P2 as played against that player’s network neighbors (if s/he had any). Thus, for example, $XYYYY$ means that the player chose action X in the first round and action Y 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 *actually observed* given his knowledge of the network structure and assuming he was playing according to the PBE strategies (as described in lemmas 2 to 11 of Corbae and Duffy[9].) We then divided this count by the total number of choices made by all unshocked and linked players. For example, in Figure 7, Group 2, game 1, player 1 chose action Y in the first round counter to the PBE prescription, but then chose Y 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. (see Lemma 6 of Corbae and Duffy [9]). Hence, in game 1, subject 1 chose the right action (played a best response) 4 times. Player 3 also started out playing Y counter to the PBE prescription. Given that choice, player 3 should have expected that player 4 would resort to playing Y in the four remaining rounds, and so player 3 should have continued playing action Y in the remaining 4 rounds. Instead, player 3 played X in the next four rounds. Therefore, we conclude that in game 1, subject 3 played 0 best responses. Finally, player 4 started out playing X as prescribed by the PBE strategy. Once player 4 observed that player 3 played Y in round 1, player 4 should have resorted to playing Y in the remaining 4 rounds of the game. In fact, player 4 played Y in rounds 2 and 4 and X 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 7. Notice that in Game 2 and those following it, players begin playing exactly according to the ex-ante, payoff dominant PBE strategy for M networks (as described in lemma 6 of Corbae and Duffy [9]). The other best response frequencies are calculated in a similar fashion.⁴¹

Our discussion of our experimental results is divided up into two parts, corresponding to our two main hypotheses: (1) players’ behavior in the second stage stag-hunt game, and (2) link proposals and network configurations in the first proposal stage.

⁴⁰Space constraints prevent us from presenting the raw data from the fourth session of the 5-game treatment, or from any of the 9-game treatments, however all of this data is considered in the various aggregate statistics reported on below in the text. Readers interested in the complete, raw dataset from all sessions may want to examine the Technical and Data Appendix to this paper, Corbae and Duffy [9], which is available at: <http://www.pitt.edu/~jduffy/networks/>

⁴¹Notice 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 address in further detail below.

Treatment	Average Frequency of X in First Round	Average Game Payoff Per Subject (Std. Dev.)	Ratio of Avg. Game Payoff to PBE Payoff (Std. Dev.)
M, 5-Games	0.917	\$2.166 (.387)	0.925 (.090)
M, 9-Games	0.909	\$2.196 (.235)	0.965 (.056)
M-all	0.913	\$2.181 (.297)	0.945 (.073)
LI, 5-Games	0.766	\$1.498 (.020)	0.825 (.053)
LI, 9-Games	0.682	\$1.852 (.323)	0.875 (.056)
LI-all	0.724	\$1.675 (.288)	0.850 (.057)
UM, 5-Games	0.867	\$1.935 (.339)	0.902 (.062)
UM, 9-Games	0.465	\$1.735 (.225)	0.843 (.089)
UM-all	0.666	\$1.835 (.287)	0.872 (.077)

Table 1: Actions and Payoffs of Unshocked Players With at Least One Link

5.2.3 Behavior in the second-stage stag-hunt game

Given a network configuration, our theory prescribes how play should evolve in every subgame of the stag-hunt game (see lemmas 2-11 in the Technical Appendix). These theoretical predictions serve as the basis for Hypothesis 1.

In examining that hypothesis, we note first that, regardless of the network structure implemented, *all* of our equilibria have unshocked players with one or more links choosing action X in the first round of every second-stage stag-hunt game. It is therefore of interest to consider what actions players *actually* choose in the first round of these games.

Finding 1 *In treatments where players start out in M or LI networks, most (more than 50%) linked and unshocked players choose action X in the first round of each second-stage game. Further, the frequency of first round X choices is increasing over time. These findings do not hold in the treatment where players start out in UM networks.*

Support for Finding 1 is found in Table 1 and Figure 10. The first column of Table 1 shows the average frequency of play of the risky action X by linked and unshocked players in the first round of all second-stage stag-hunt games played in all sessions of a given treatment. Figure 10 disaggregates these first-round frequencies of choosing X by game number to give some sense of how these frequencies change with experience. In the treatment where players started out in M networks, the frequency of choosing X in the first round is greater than 90% on average 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 frequency of choosing X in the first round is greater than 70% on average and increases slightly with experience; the mean frequency of play of X 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 X is lowest, averaging 67%, but with decreasing considerably with experience. Indeed the average frequency of first round play in the 9-game treatments is less than 50%. Figure 10 reveals that this low average is due to a large drop-off in the frequency of initial play of X in the last four games played in the 9-game sessions. Indeed, the frequency of first round X play falls from an average of 68% over the first five games to an average of 47% over the last four games.

[Insert Figure 10 here.]

The high frequency of play of action X 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, and all ex-ante PBE strategies prescribe the play of action X by unshocked players in the first period.⁴² Our assumption of ex ante payoff maximization – all unshocked players playing X in round 1– comes closest to being realized in sessions where players start out in M networks. If we disaggregate the results presented in Table 1 and Figure 10 by group (session) we find that for 6 of the 8 groups starting out in M networks, the frequency of play of action X in the first round exceeds 90% over all games, and for all 8 groups that started out in M networks, this frequency always exceeds 50%. A binomial test confirms that we may reject the null hypothesis that players were equally likely to choose action X or Y in the first round of all games in favor of the alternative that they were more likely to choose action X ($p=.004$). 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 X in the first round exceeds 50%. Again, using a binomial test of the null hypothesis that players were equally likely to choose action X or action Y , we can reject this hypothesis in favor of the alternative that they were more likely to choose action X ($p=.035$). 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 X 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 X or action Y in the first rounds of this treatment ($p>.10$). As noted above, for the UM treatment there is an observed decrease over time in the mean frequency of first round X choices, as illustrated in Figure 10. One possible explanation for this finding 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 of the initial imposition of a UM network configuration. Indeed, as we show below, the frequency of players with 2 or 3 links is highest (and the frequency of players with just 1 link is lowest) in groups that started out in exogenously imposed UM networks. As a consequence, we surmise that players in such groups may have become wary of playing action X in any round, including the first, given the knowledge that they were frequently indirectly linked to the shocked player.

We next consider the payoffs that players earned in all rounds of the second stage game relative to perfect Bayesian equilibrium payoffs. Specifically, for each game, we consider the network that was actually implemented by the subjects at the completion of the proposal stage. We then calculated two statistics: 1) the actual total payoff earned by all unshocked and linked players in that network over the 5-repetitions of the second-stage game and 2) the total payoff that each of these same unshocked and linked subjects could have earned had they played according to the ex-ante, payoff dominant PBE strategy for the network implemented (as described in lemmas 2 to 11 of Corbae and Duffy [9]). For example, suppose an M network was implemented. We first calculate the total payoff actually earned by the three unshocked and linked subjects in the 5 rounds played in that M network. We then calculate what the total payoff to these three unshocked players over the 5 repetitions of the stag-hunt game would have been had they played according to the PBE strategy. In an M network, this potential total payout equals \$7.40 (or an average payoff of \$2.47 for each of the three unshocked players).⁴³ We then calculated the ratio of actual total payoffs earned by the three unshocked subjects

⁴²The frequency of play of action X in rounds two through five is a far less informative statistic, as the ex-ante PBE strategies prescribe the play of action X in those rounds only in certain networks and not in others (e.g. UM vs. LI) and only conditional on a certain history of play by unshocked players. We will characterize the extent to which play in rounds 1-5 of the second stage game accords with the ex-ante PBE strategies below in Findings 2-3.

⁴³Following the PBE strategy, two players would earn \$.60 in each of the five repetitions of the game (from both playing X), and the other player matched to the shocked player would earn 0 in the first round (since he would have started out playing X against the shocked player's Y in that first round) and \$.35 in the remaining 4 rounds— $.60 \times 10 + .35 \times 4 = \7.40 .

to the total potential PBE payout as a measure of payoff efficiency for each game. In Table 1 we report the averaged values (and standard deviations) of these efficiency ratios across all games and sessions of the various treatments. In addition we report average amounts actually earned by each linked and unshocked subject in all 5-round games across games and sessions of various treatments.

Finding 2 *Subjects’ achieve payoffs that are, on average, 82% to 95% of the payoffs they could have achieved by playing according to the ex-ante payoff dominant perfect Bayesian equilibrium strategy in the second-stage game, given the network they implemented in the first stage. Further, there are no significant differences in these efficiency ratios across treatments.*

Support for finding 2 is found in the last column of Table 1. While it appears that the efficiency ratio is slightly higher in sessions where subjects started out in M networks, nonparametric rank order tests reveal these differences to be statistically insignificant ($p > .10$) in all pairwise comparisons between treatments using session-level data. This finding suggests that the PBE strategies we use to characterize play in the second-stage game may have some explanatory power in line with Hypothesis 1.

In addition to considering payoff efficiency, we can also examine the extent to which subjects’ *action choices* are in accordance with PBE predictions. In particular, consider the frequency of best response play by the unshocked members of each group over all games of a given treatment. The bar charts at the bottom of Figures 7, 8, and 9 show mean best response frequencies over all rounds of each game for the unshocked members of each 4-player group (who have at least one link). While these Figures illustrate less than half of the groups for which we have data, a perusal of these charts might again lead one to the conclusion that players who started out in M networks played more frequently in accordance with the PBE predictions than players who started out in LI or UM networks. However, the evidence from the entire dataset does not warrant such a strong conclusion. Indeed, with a single exception, no group in any treatment had a best response frequency of 100% over all games played.⁴⁴

Consider Figure 11, which consists of two charts showing mean best response frequencies by unshocked players over all games for each group in the 5-game (top chart) and 9-game (bottom chart) sessions. Each bar in these charts represents a single group starting out in M, LI or UM, and sessions are collected together for comparison purposes. The mean best response frequencies are calculated over all games played in the session.⁴⁵ Consider first the top chart showing mean best response frequencies over all games in the 5-game sessions. There we observe that the mean best response frequencies are always slightly higher for the four 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, and as a consequence, the rank order test does not allow rejection of the null hypothesis of no difference in mean best response frequencies between these two treatments. 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 4, has a higher best response frequency than its 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 11 here.]

⁴⁴The exception was a group in the 9-round treatment that started out in a M network.

⁴⁵Thus, for example, group 1 in session 1 of the 5-game treatment, which started out in a M network, achieved a mean best response frequency of 98.67% over all 5 games, which is just the mean of the five best response frequencies in games 1-5 as illustrated in the bar chart in the bottom left corner of Figure 7.

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 11). A one-sided binomial 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, in support of Hypothesis 1. We summarize our observation concerning best response frequencies as follows.

Finding 3 *In the second-stage stag-hunt game, unshocked subjects play best responses at least 75% of the time across all treatments. These frequencies of best response play are not significantly different across treatments.*

5.2.4 Behavior in the first-stage proposal game

We now turn our attention to behavior in the first-stage proposal game. In our theoretical analysis we considered whether any single agent had an incentive to deviate proposing a link necessary to implement a given network proposal profile. Since there is a coordination issue involved in implementing networks, whether or not agents are actually able to implement a given network is an open question. Recall also that link proposals of the form $\{0\}$ (i.e. refuse all links) were always a valid proposal available to subjects. Here we analyze both data on link proposals as well as the data on links actually implemented (mutually agreed upon) conditional on the different treatments associated with the initial network in which agents started.

Our first finding is that subjects do propose to form at least one link.

Finding 4 *Independent of the initial imposed network structure, subjects nearly always proposed to link to at least one other player.*

When players were given the opportunity to choose links, 95% of all players – 91 out of 96 players – chose to submit at least one link proposal in every game played.⁴⁶ 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, $d = b = 35$ did not cause players to avoid proposing links, i.e., submitting $\{0\}$ proposals. Indeed, it also serves to highlight the difficult coordination problem that players faced in the link formation stage. There were several instances where all 4 players in a group submitted link proposals but there were no *mutually agreed upon* links, resulting in an A network.

We next examine the *number* of link proposals that players made over time and across the three treatments. In Table 2 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 5 *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.*

⁴⁶The five exceptions to this rule are found in the 9-game M, LI, or UM treatments.

Support for this finding is found in Table 2. We see that the mean frequency of a single link proposal increases steadily as players gain experience (play more games) in treatments where players started out in M networks. By contrast, in treatments where subjects started out in LI networks, experienced 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 in the UM, 9-game treatment; each member of this group proposed a single link in Game 6 and the resulting network was the autarchic one (A).

Finding 5 suggests that players were learning over time to move away from non strict-PBE network structures (LI, UM) and towards strict PBE networks (M, M-A, LI-A) in the sense that players who start out in LI or UM networks move away from proposing 2 or 3 links respectively, while those who start in M networks typically propose a single link. As the number of links actually formed cannot exceed the number of link proposals, it follows that similar findings should hold for the number of links actually implemented by subjects, by mutual consent. We now turn our attention to this issue of link formation.

Table 3 summarizes the link formation results 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 (5-game) and the 4 long (9-game) sessions, while for games 5–9, the results are from the 4 long sessions alone. Figure 12 illustrates the mean link frequencies reported in Table 3 over time, i.e., for each game of a treatment (players beginning in M, LI or UM).

[Insert Figure 12 here.]

As a severe stability test, we use the nonparametric, Kolmogorov-Smirnoff goodness-of-fit test to ask whether the sample cumulative distribution function, $F^S(Z)$, over links $Z = 0, 1, 2, 3$ differs from a theoretically predicted cumulative distribution function, $F^T(Z)$, where $F^i(Z)$ represents the proportion of observations that are less than or equal to Z for $i \in \{S, T\}$. In the case where players started out in an M network, our hypothesis is that each player should maintain their single link (as M networks belong to the set of strict-PBE networks) 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 $F^S(Z)$ and $F^T(Z)$. 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 3. 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 6 *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.*

As seen in Table 3, 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.

However link proposals are not conclusive evidence of the networks actually formed, as link formation in our environment requires mutual consent. We next turn to an analysis of the actual networks implemented and the main question of our paper, namely whether there is support for Hypothesis 2?

[Insert Table 4 here.]

Table 4 reports the frequency with which all possible network types (as listed in Figure 1) were actually implemented when subjects were free to choose their own links (i.e., we ignore the networks formed in the first games of each session, where networks were exogenously imposed on subjects). These frequencies are divided up according to whether subjects began the session starting in M, LI or UM networks and whether they played 5 or 9 games (5-G or 9-G). The first three rows show the frequencies with which the three strict-PBE network types M, M-A, and LI-A (as identified in Proposition 1) are observed across all sessions of the various treatments. The remaining rows report the observed frequencies of the other eight (non-strict PBE) network types. The overall frequencies of the various network types, (i.e., all sessions of all treatments) is shown in the penultimate column. The final column, “Simul. Random”, reports the frequency of network types obtained from a simulation

of random link proposals by four agents (frequencies are based on a simulation of 5 million, 4-player network proposal games).⁴⁷

There are several interesting findings in this table. First, over all treatments/sessions, the class of strict-PBE networks accounts for 49 percent of the endogenously determined networks. Further, M networks are the most frequently observed network-type at 28 percent overall. By contrast, in the simulation of random network formation, M networks are only the third most frequent network type at just 5 percent, and strict-PBE networks are observed only 43 percent of the time; in the random proposal simulations, there is a much greater frequency (36 percent) of M-A networks than in the human subject data. Second, the other two symmetric network types that were exogenously imposed on groups in the first games of sessions are almost never observed again when subjects are free to choose links; overall, just 1 percent of all endogenously formed networks were of the LI-type and the UM network is *never* observed when players have the opportunity to choose links. Third, among the set of non-strict PBE networks, the three most frequently observed – LI-M-A, LI-M and UM-LI-M in that order – are quite “close” to being strict-PBE networks, (M, M-A, LI-A), in the sense that removal of a *single* link in LI-M-A, LI-M or UM-LI-M could result in M-A, M or LI-A, respectively (see Figure 3); the other four non-autarkic, non-strict PBE networks (UM, UM-LI, LI and UM-M) would require more than a single deletion of a link to become one of the strict-PBE networks. Finally, the autarkic (A) network arises in just 5 percent of all human subject determined networks but arises in 18 percent of the random proposal simulated networks. Indeed, using a Chi square goodness-of-fit test, we can easily reject the null hypothesis that the frequencies of the 11 network types in the human subject data are the same as would be obtained by random proposal links ($p < .001$; $N = 144$).

We summarize these findings as follows:

Finding 7 *Over all treatments and sessions, the set of strict-PBE networks (M, M-A, LI-A) accounts for 49 percent of all endogenously determined networks. The most frequently observed network is the M-network. The UM and LI networks are never or rarely observed. Among non-strict PBE networks, those that are “close” to being strict-PBE networks (as defined above) are the most frequently observed. The observed frequency of the 11 network types is significantly different than would be predicted by random proposal links.*

If we condition the frequency of the various networks observed in Table 4 on the initial conditions, i.e., whether subjects started out in M, LI or UM networks, we have following additional finding:

Finding 8 *Strict-PBE networks (M, M-A, LI-A) are observed with a high frequency (75 percent or greater) in treatments where players initially began in M networks. Strict-PBE networks are less frequent (less than 50 percent) but still prominent in treatments where players initially began in LI or UM networks.*

There is a straightforward explanation for Finding 8: the problem of coordinating on a strict-PBE network is considerably less difficult for subjects who start out in M networks, as M networks are among the set of strict-PBE networks. Indeed, as Table 4 reveals, in treatments where subjects began (exogenously) in M networks, the frequency of endogenously formed networks that were of type M averaged 50 percent; adding M-A networks this frequency jumps to 77 percent (LI-A networks are never observed when players start out in M networks). By contrast, in treatments where subjects began in non-strict-PBE networks such as LI and UM, the process of coordinating on a strict-PBE network

⁴⁷For each simulated proposal game, each of four players randomly decides whether or not to propose a link to the other three players. Mutual consent rules determined the networks actually formed, and a tally was kept of which of the 11 different network types obtained in 5 million repetitions of this random proposal game.

appears to have been more difficult, though not impossible. In treatments where subjects began in LI (UM) networks, Table 4 reveals that 25 (30) percent of the time, the endogenously formed network was a marriage network (M or M-A); more generally, 27 (42) percent of the time the endogenously formed network was a strict-PBE network (M, M-A, or LI-A).

In addition to considering whether subjects were able to implement strict-PBE networks, we want to further explore whether subjects were able to sustain or re-implement such networks once they were achieved. After all, once a strict-PBE network has been implemented, there is no reason to deviate from that network structure and we should see it implemented again and again, i.e., it should remain *stable*. Indeed, one reason that we started subjects out in exogenous M, LI or UM networks was to test such a stability hypothesis.⁴⁸

In particular, suppose agents were able to determine the set of strict-PBE networks in advance of play (as we have done in Proposition 1). Then if subjects were exogenously placed in a strict-PBE network (e.g. the M network), they should want to reimplement that same type of network, and if they were not placed in a strict-PBE network (e.g. the LI or UM network), they should want to move away from that type of network. The Kolmogorov-Smirnoff tests reported above in the discussion of Table 3 suggested that the distribution of *links* among players who began in M networks was often statistically indistinguishable from a distribution where every player had exactly one link while the distribution of links among players who began in LI or UM networks was considerably more diffuse. Here, we focus not on the distribution of links over time, but rather on the stability of given network structures from one game to the next. In particular, we have:

Finding 9 *M networks are frequently stable (sustainable) while LI and UM networks are not.*

Consider first the case where players initially start out in exogenous M networks and played 5-games, e.g., as illustrated in Figure 7. After some experimentation, players in this treatment always settled on link choices that resulted in implementation of M networks by the last game of the session; in groups 2 and 3 this had happened by game 3 while in group 1 by game 4. In the remaining 5-game, M-treatment observation (not shown in Figure 7), an M network was only achieved by the last game. Similar results were obtained in the 9-game sessions, which are not illustrated here due to space constraints.⁴⁹ Indeed in all four of the 9-game sessions, we find that players always end up in an M network, and this M network is *sustained* for at least two sequential games at the end of each 9-game session (e.g. games 8-9) and sometimes for more than two sequential games. 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 in all M-observations – group 2 of the 5-game-treatment, shown in Figure 7 – managed to coordinate on a different network structure and this one was also a symmetric M network!

On the other hand, as Table 4 reveals and Figures 8 and 9 illustrate, when players started out in exogenously imposed LI or UM networks, they did not consistently choose links so as to continue to maintain a LI or UM network, even after some period of experimentation. Indeed, with a single exception, *no group that started out in LI or UM ever succeeded in re-implementing that same type of network following the first game, when network formation was endogenously determined.*⁵⁰ This finding stands in stark contrast to the findings for the groups starting out in M networks who *always* succeeded in implementing a symmetric M network by the last few games of the session, and suggests strong support for hypothesis 2

⁴⁸ Another reason was to consider the impact of initial conditions.

⁴⁹ But see Figures M9.a–M9.d in the Technical and Data appendix, Corbae and Duffy [9]

⁵⁰ LI was endogenously implemented by one group in one of the 9-game sessions of the LI treatment and by another group in one of the 9-game sessions of the UM treatment, but this LI network was never re-implemented (i.e. it was not sustained).

The result of the endogenous link choices by groups who started out in LI or UM is typically, though not always, some kind of *asymmetric network* – see the frequency of network types given in Table 4 – though none of these asymmetric networks is *sustained* for more than a game or two – see Figures 8–9, for example. When a group that started out in LI or UM succeeded in sustaining a network for more than three sequential games (which happened only in the 9-game sessions), the network was *always* a M network. Furthermore, once achieved, these strict-PBE networks were sustained for the duration of the session! Again, such behavior would appear to support our hypothesis 2 that M is stable.

The unraveling of the UM network may seem somewhat surprising in light of the results we obtained for the exogenous UM network (lemma 2 of Corbae and Duffy [9]) i.e., that all unshocked players play X thus avoiding any contagion) and the fact that in ex-ante terms players are better off than in LI. Still, we know from Proposition 1 that UM is not a PBE network (weak or strict), and the experimental findings are consistent with this result. The unraveling of the LI network may also seem surprising in light of our theoretical finding that LI networks are weakly stable (Lemma 14 of the Technical appendix). However, if we use the strict PBE refinement, as we do in Proposition 1 then we wouldn't expect to see LI networks as an equilibrium outcome.

6 Conclusions

Broadly consistent with the theory we have developed, 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 in the stag-hunt stage game where payoffs in a given network can be ranked $M > UM > LI$. Note that this expected payoff ranking is nonlinear in the number of one's neighbors. 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).

Our findings also shed some light on the difficulty of implementing a network in a noncooperative manner as in Myerson [29]. While we imposed a given network structure in the first proposal stage as a way of coordinating agents' beliefs over subsequent proposal profiles, the coordination problem when subjects are free to choose links in a noncooperative simultaneous move game is rather challenging. Nevertheless, once groups had coordinated on a marriage network (after some experimentation), that same network configuration was always sustained (i.e. re-implemented) in all subsequent two-stage games for the duration of the experimental session, that is, the coordination problem was effectively solved. Furthermore, we found that subject behavior was quite purposeful and appeared to be largely consistent with our assumptions of Bayesian updating and the use of ex-ante payoff dominant strategies; for instance in the second-stage game, payoff efficiency is around 90 percent of the PBE prediction (Finding 2) and in the first stage game, the distribution of network structures is significantly different from that which would result from random play (Finding 7).

An interesting extension would be to allow some form of communication or possibly allow sequential moves before the proposal game to ease the coordination problem. Alternatively, one could give players more than a single game of experience with a single, exogenously imposed network structure before setting them free to form networks endogenously. We leave these extensions to future research.

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] Berninghaus, S.K., M. Ott, and B. Vogt (2004): “On Networks and Stars – Recent Results in Network Experiments,” working paper no. 04-34 Sonderforschungsbererich 504, Universität Mannheim.
- [5] Callander, S., and C.R. Plott (2005): “Principals of Network Development and Evolution: An Experimental Study,” *Journal of Public Economics* 89, 1469-1495.
- [6] Carlsson, H. and E. van Damme (1993): “Equilibrium Selection in Stag Hunt Games,” in K. Binmore, A. Kirman, and P. Tani (Eds.), *Frontiers of Game Theory*, Cambridge, MA: MIT Press.
- [7] Charness, G., and M.O. Jackson (2007): “Group Play in Games and the Role of Consent in Network Formation,” *Journal of Economic Theory* 136, 417-445.
- [8] Corbae, D. and J. Duffy (2002): “Experiments with Network Economies,” working paper, July, available at <http://www.pitt.edu/~jduffy/networks/>
- [9] Corbae, D. and J. Duffy (2007): “Technical and Data Appendix to ‘Experiments with Network Formation,’” available at <http://www.pitt.edu/~jduffy/networks/>
- [10] Devetag, G. and A. Ortmann (2007): “When and Why? A Critical Survey on Coordination Failure in the Laboratory,” *Experimental Economics* 10, 331-344.
- [11] Diamond, D. and P. Dybvig (1983): “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy* 91, 401-19.
- [12] Ellison, G. (1993): Learning, Local Interaction, and Coordination,” *Econometrica* 61, 1047-1071.
- [13] Falk, A. and M. Kosfeld (2003): “It’s All About Connections: Evidence on Network Formation,” University of Zurich, IZA Discussion Paper No. 777.
- [14] Fudenberg, D. and J. Tirole (1991): *Game Theory*, Cambridge, MA: MIT Press.
- [15] Garratt, R. and T. Keister (2005): “Bank Runs: An Experimental Study,” working paper, UC Santa Barbara.
- [16] Goeree, J.K., A. Riedl and A. Ule (2005): “In Search of Stars: Network Formation among Heterogeneous Agents,” working paper, CalTech.
- [17] Harsanyi, J. and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, Cambridge, MA: MIT Press.
- [18] Hauk, E. and R. Nagel (2001): “Choice of Partners in Multiple Two–Person Prisoner’s Dilemma Games,” *Journal of Conflict Resolution* 45, 770–793.
- [19] Heinneman, F., R. Nagel and P. Ockenfels, (2004): “The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information,” *Econometrica* 72, 1583–1599.

- [20] Jackson, M.O. (2005): A Survey of Models of Network Formation: Stability and Efficiency,” in G. Demange and M. Wooders (Eds.), *Group Formation in Economics: Networks, Clubs and Coalitions*, Cambridge: Cambridge University Press, 11–57.
- [21] Jackson, M.O. and A. Watts (2002): “The Evolution of Social and Economic Networks,” *Journal of Economic Theory* 106, 265–295.
- [22] Jackson, M.O. and A. Watts (2006): “Social Games: Matching and the Play of Finitely Repeated Games,” working paper, Caltech.
- [23] Jackson, M.O. and A. Wolinsky (1996): “A Strategic Model of Social and Economic Networks,” *Journal of Economic Theory* 71, 44–74.
- [24] Kandori, M., G. Mailath, and R. Rob (1993): “Learning, Mutation, and Long-Run Equilibria in Games,” *Econometrica*, 61, 29–56.
- [25] Keser, C., K-M. Ehrhart, and S.K. Berninghaus (1998): “Coordination and Local Interaction: Experimental Evidence,” *Economics Letters* 58, 269–275.
- [26] Kosfeld, M. (2004): “Economic Networks in the Laboratory: A Survey,” *Review of Network Economics* 3, 20–41.
- [27] 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.
- [28] Morris, S. (2000): “Contagion,” *Review of Economic Studies*, 67,57–78.
- [29] Myerson, R. B. (1991): *Game Theory: Analysis of Conflict* Cambridge: Harvard University Press.
- [30] Rankin, F., J.B. Van Huyck and R.C. Battalio (2000): “Strategic Similarity and Emergence of Conventions: Evidence from Payoff Perturbed Stag Hunt Games,” *Games and Economic Behavior* 32, 315–337.
- [31] Schotter, A. and T. Yorulmazer (2004), “On the Severity of Bank Runs: An Experimental Study,” working paper, New York University.
- [32] 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.
- [33] Young, H.P. (1993): “The Evolution of Conventions,” *Econometrica*, 61, 57–84.
- [34] Young, H.P. (1998): *Individual Strategy and Social Structure*, Princeton: Princeton University Press.




Figure 3: Illustration of all unilateral deviations for 4-player groups

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

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



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 are free to choose X or Y.
1 other member(s) of economy A must choose Y.

X Y

Results

Game	Round	My Choice	Neighbor Choices	My Payoff
1	4	Y	2Y4Y	35.00
1	3	Y	2Y4Y	35.00
1	2	X	2Y4Y	0.00
1	1	X	2X4X	60.00

Figure 5: Illustration of Unshocked Player's Decision Screen

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

PAYOFF				
Your Neighbors				
	2X	1X1Y	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


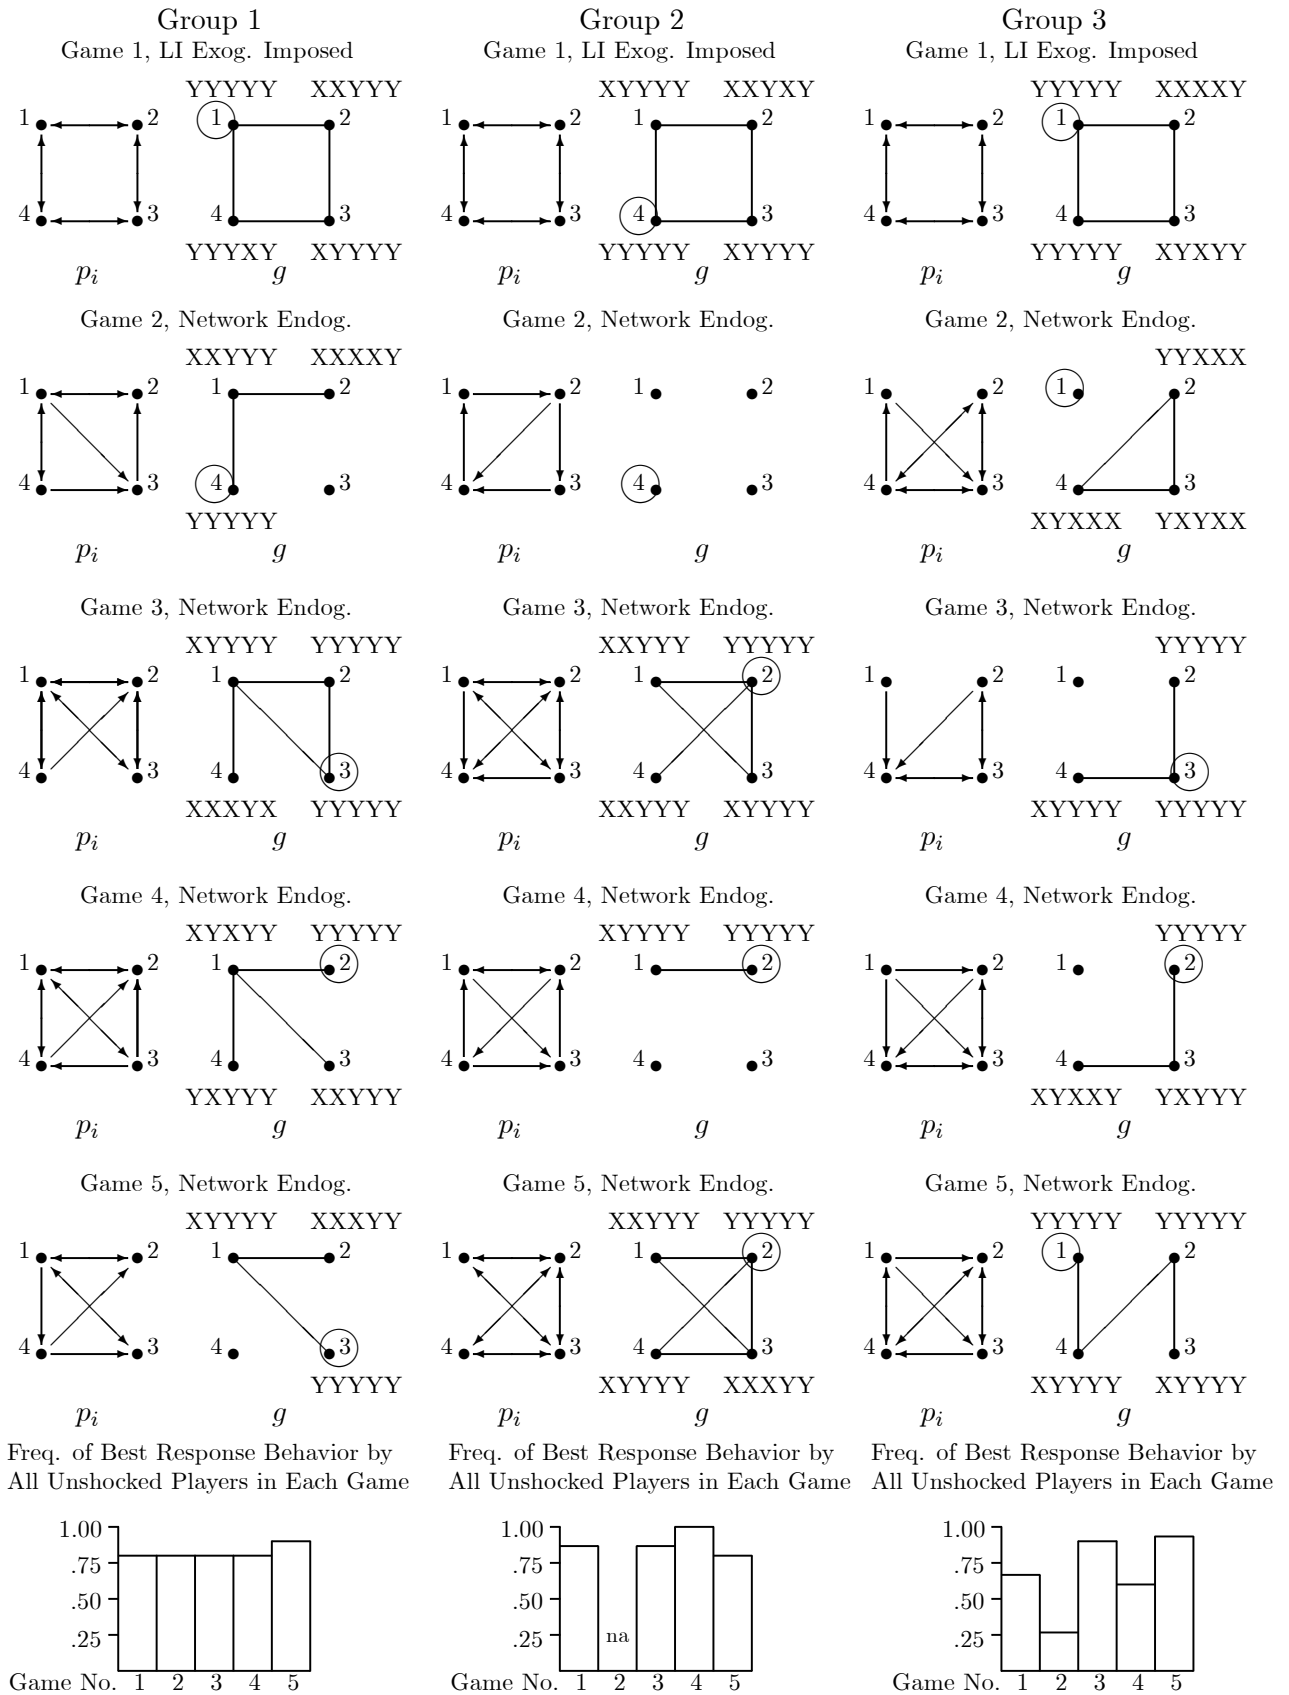


Figure 7: Three Groups Initially in M Play 5 Games






Figure 9: Three Groups Initially in UM Play 5 Games




Figure 10: Frequency of Play of Action X by Unshocked Players in the First Round of the Second Stage Stag-Hunt Game (Using Pooled Data From All Sessions of a Given Treatment)

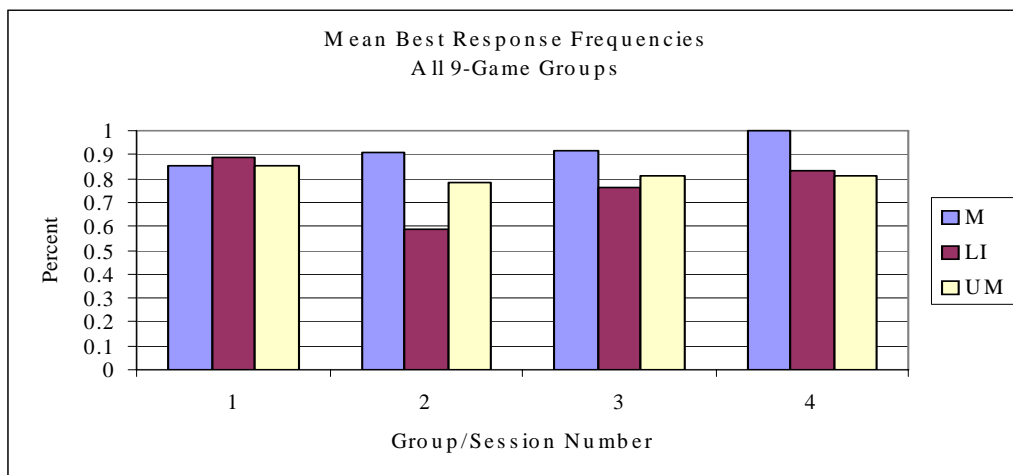



Figure 11: Mean Best Response Frequencies, All Treatments

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

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 3: Link Frequencies by Game and Across Treatments






Figure 12: Frequencies of Players Having 0, 1, 2 or 3 Links in Each Game: All Groups Starting out in M, LI or UM

Network Type	Starting in M			Starting in LI			Starting in UM			Over All Sess.	Simul. Random
	5-G	9-G	All	5-G	9-G	All	5-G	9-G	All		
M	0.56	0.47	0.50	0.06	0.22	0.17	0.06	0.22	0.17	0.28	0.05
M-A	0.19	0.31	0.27	0.06	0.09	0.08	0.31	0.03	0.13	0.16	0.36
LI-A	0.00	0.00	0.00	0.06	0.00	0.02	0.06	0.16	0.13	0.05	0.02
All PBE	0.75	0.78	0.77	0.19	0.31	0.27	0.44	0.41	0.42	0.49	0.43
LI	0.00	0.00	0.00	0.00	0.03	0.02	0.00	0.03	0.02	0.01	0.01
LI-M-A	0.19	0.03	0.08	0.25	0.31	0.29	0.31	0.16	0.21	0.19	0.24
LI-M	0.06	0.06	0.06	0.25	0.06	0.13	0.13	0.09	0.10	0.10	0.08
UM	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
UM-LI-M	0.00	0.00	0.00	0.13	0.13	0.13	0.13	0.13	0.13	0.08	0.03
UM-M	0.00	0.00	0.00	0.06	0.06	0.06	0.00	0.00	0.00	0.02	0.03
UM-LI	0.00	0.00	0.00	0.06	0.06	0.06	0.00	0.16	0.10	0.06	0.00
A	0.00	0.13	0.08	0.06	0.03	0.04	0.00	0.03	0.02	0.05	0.18
All Other	0.25	0.22	0.23	0.81	0.69	0.73	0.56	0.59	0.58	0.51	0.57

Table 4: Frequency With Which Network Types are Endogenously Implemented Across Treatments