Available online at www.sciencedirect.com
SCIENCE DIRECT®

Games and Economic Behavior 55 (2006) 112–130

GAMES and Economic Behavior

www.elsevier.com/locate/geb

Endogenous networks, social games, and evolution

Daniel A. Hojman [a,*], Adam Szeidl [b]

[a] *Department of Economics, Harvard University, Cambridge, MA 02138, USA*
[b] *Department of Economics, University of California, Berkeley, CA 94720, USA*

Received 27 December 2002

Available online 26 April 2005

Abstract

This paper studies a social game where agents choose their partners as well as their actions. Players interact with direct and indirect neighbors in the endogenous network. We show that the architecture of any nontrivial Nash equilibrium is minimally connected, and equilibrium actions approximate a symmetric equilibrium of the underlying game. We apply the model to analyze stochastic stability in 2×2 coordination games. We find that long-run equilibrium selection depends on a trade-off between efficiency and risk dominance due to the presence of scale effects arising from network externalities. Our results suggest a general pattern of equilibrium selection.

JEL classification: C72; C73

Keywords: Network formation; Stochastic stability; Scale effects; Coordination

1. Introduction

Individual decisions and strategic interaction are both embedded in a social network. The nature of outcomes that arise in economic, social and political environments may be seriously affected by the underlying interaction structure. This has been recognized in the literature on long-run equilibrium selection in coordination games. In fixed pop-

* Corresponding author at: Harvard University, Department of Economics, Littauer Center 200, Cambridge, MA 02138, USA.

E-mail address: hojman@fas.harvard.edu (D.A. Hojman).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2005.02.007

ulation structures, both global and local interaction favor the risk dominant equilibrium, though with very different convergence rates[1]; on the other hand, mobility favors payoff dominance.[2] This paper presents a model of social games in which agents simultaneously choose actions and partners in a network. This allows us to study evolutionary selection in an environment where the interaction structure is endogenous. Our main finding is that selection is determined by a trade-off between risk- and payoff-dominance, due to scale effects arising from network externalities.

In our model, agents choose whom to interact with, and maintain costly links connecting them to their opponents. Players are represented as nodes, partnership links as edges of the social network. A connection, direct or indirect, provides access to pairwise interaction, with payoffs derived from an underlying game Γ. We formalize the microstructure of interactions as a one-way flow model: links are directed, and agents receive payoffs only from interactions with those neighbors to whom there is a directed path in the network. Endogenous choice of neighborhoods gives rise to scale effects because of the non-rivalry in link formation.

To understand the dynamics of the social game, we first characterize its static behavior. We show that a (nontrivial) Nash equilibrium has to be minimally connected. In addition, any strict Nash equilibrium has the architecture of a wheel. Hence, our model extends the results of Bala and Goyal (2000) for a setting in which agents choose not only their neighbors, but also their behavior.[3] These results allow us to disentangle the architecture from the actions that obtain in a Nash equilibrium. Since the equilibrium is connected, actions correspond to some Nash equilibrium of an N-player game in which each player plays with everybody else. From this we conclude that for N large the equilibrium distribution of play in the social game resembles arbitrarily closely a symmetric equilibrium of the underlying game.

To study selection, we first characterize the limit sets of a myopic (unperturbed) best-response dynamics. We show that under weak conditions any nonempty limit set has the architecture of a wheel; moreover, if Γ is a coordination game, then actions also converge. In the case of 2×2 coordination games, this implies that the only nonempty limit sets are wheels where each player is coordinating on one of the two possible equilibria. Evolution will choose between these two equilibrium candidates, the A-wheel and the B-wheel.

The literature on stochastic stability in coordination games has identified two forces influencing long-run equilibrium selection. The first, *risk dominance* favors the less risky strategy for the intuitive reason that it yields a higher payoff against an average population. The second force, *mobility*, favors the Pareto optimal strategy. The reason is that when players are allowed to move, they will migrate to locations where the efficient action is played and away from clusters favoring the inefficient action. Hence, the interaction structure determines which equilibrium is selected. In our model, en-

[1] See Kandori et al. (1993) (hereafter KMR), Young (1993) and Ellison (1993).

[2] As in Ely (2002) and Oechssler (1999).

[3] In these equilibria, players typically have access to the benefits of pairwise interaction mostly through intermediaries, in line with the role played by ‘weak’ linkages in social networks, as emphasized by Granovetter (1973).

dogenous link formation provides a choice of opponents much like mobility does; and given a set of opponents, risk dominance matters in the choice of strategy. However, we also have a third effect influencing selection, namely *scale*. In the presence of scale effects, it may be optimal to stick to a large cluster of agents coordinating on the inefficient equilibrium instead of moving away to small (if efficient) clusters. Scale effects bring about the trade-off between risk dominance and efficiency.

Our results verify this intuition. First, in a game where lack of coordination is costly, payoff considerations will have a serious influence on link decisions. We prove that an equilibrium will be selected if and only if it is sufficiently risk dominant relative to its payoff disadvantage vis-à-vis the other equilibrium. An equilibrium that is not payoff dominant may be uniquely selected if it is not too risky; likewise, an equilibrium that is not risk dominant may be uniquely selected if its payoff is high enough. We have the exact same prediction if the costs of link formation are so high that agents cannot maintain more than one link. Intuitively, scale effects matter more when link formation is costly, because then who to link with is a crucial decision. Finally, if the costs of link formation are low and payoffs are always positive then scale effects do not matter: everybody can afford to link to everybody else. This case is similar to the classic KMR framework; each player optimally links to all opponents and the risk dominant equilibrium is uniquely selected.

As discussed in Section 5, the intuition for our results also leads us to identify a simple pattern driving selection in many of the models in the literature, namely that flexibility favors payoff dominance.

Beyond the literature on stochastic stability in coordination games mentioned above, this paper is also related to recent work on endogenous network formation. Bala and Goyal (2000) show in a general model that strict Nash networks are minimally connected, and are the long-run outcomes of a best-response dynamics under some conditions. Galeotti et al. (2004) extend the Bala–Goyal model to allow for heterogeneous populations. In a new paper, Hojman and Szeidl (2004) show that all equilibria have the star architecture in the presence of decreasing returns and arbitrary discounting. However, none of these papers study a social game and they do not analyze the role of endogenous network formation for equilibrium selection in coordination games.

More related are Goyal and Vega-Redondo (2004) and Jackson and Watts (2002), both of which study the interaction between link formation and action choice in a social game with two-way (non-directed) links. Goyal and Vega-Redondo (2004) examine selection in 2×2 coordination games. Due to the two-way flow structure they do not find the trade-off that we do, but their results fit well into the general pattern of selection identified by this paper. In Jackson and Watts (2002), links are created when there is mutual agreement between the parties, which makes it difficult to compare their results to ours. Finally, Hojman (2003), in a model with endogenous location choice, also finds that selection depends on the interaction between scale effects, efficiency and risk, confirming our broader prediction.

2. A model of social games

2.1. Setup

The model has two main building blocks, the interaction structure and the underlying game. Formally, there is a set of N players, each of whom is represented as a node of a graph G. Directed links between agents correspond to the edges of this graph.

We say that player i is *connected* to player j and write $i \mapsto j$ if there is a directed path of links from i to j in the graph G. If the shortest path from i to j involves only one edge, we say that i is directly linked to j, or that the two players are adjacent. If the shortest path from i to j involves more than one edge, we say that j is an indirect neighbor of i. The set of (direct and indirect) neighbors of player i is denoted by $n(i, G)$. We define $\bar{n}(i, G) = n(i, G) \cup \{i\}$. The term *neighbor* will refer to both direct and indirect neighbors. A *component* of the social network G is a maximal subgraph of players connected to each other, ignoring the directionality of links. The social network G is *connected* if every player i has every other player j as her (direct or indirect) neighbor.[4] We say that the network is *minimally connected* if G is connected and, given any of its edges e, the network $G \setminus \{e\}$ is not connected. The network is called the *empty network* if there are no links.

The payoffs in the social game derive from two sources: the interactions between neighbors and the costs of maintaining links. The payoff from pairwise interactions comes from playing an underlying finite two-player symmetric game Γ with each neighbor. Specifically, each player i chooses a pure action of Γ, and earns the sum of (pairwise) payoffs from playing her action choice against all of her neighbors. Direct links are costly. The cost of maintaining links is assumed to depend only on the number of links a player maintains, but not on who the player is or who she links to.

Definition 1. A *social game* is characterized by a triplet $\Gamma^S = (\Gamma, \phi(\cdot), N)$, where

(i) the underlying game Γ is a symmetric finite two-player game with strategy space X and payoff function $u : X \times X \to \Re$,
(ii) the cost function is $\phi : Z_0^+ \to \Re \cup \{\infty\}$, defined on the set of non-negative integers, strictly increasing, weakly convex, satisfying $\phi(0) = 0$,
(iii) the set of players is N; with a slight abuse of notation we will also denote the number of players by N.

We assume that X has at least two elements.

Definition 2. A *pure strategy* of player i in the social game Γ^S is a pair $s_i = (x_i, N_i)$ where $x_i \in X$ is a pure action of Γ and $N_i \subset N \backslash \{i\}$ is the subset of players to whom i maintains a direct link. Given a strategy profile $s = (s_1, \ldots, s_N)$, the *payoff* of player i is

$$\Pi_i(s) = \sum_{j \in n(i,G)} u(x_i, x_j) - \phi\left(\sum_{j \in N_i} 1\right), \tag{1}$$

[4] This concept is considerably stronger than requiring connectedness of the undirected graph.

where $n(i, G)$, the set of (direct and indirect) neighbors of i, derives from the graph of links maintained by all players, G.

Each link is maintained by one particular player, and this player pays the full cost associated with her link. Also, the environment is completely *frictionless* in the sense that the distance between two indirect neighbors in terms of the number of intermediaries does not influence the payoff derived from playing the underlying game.

Throughout the paper we assume that the social game is generic,[5] and that Γ satisfies one of the following two conditions.

Assumption (A1) $u(x, x') > 0 \quad \forall x, x' \in X.$
Assumption (A1′) $u(x, x) + u(x', x') - u(x, x') > 0 \quad \forall x, x' \in X.$

Observe that (A1′) requires all diagonal payoffs to be positive (substitute $x = x'$), though it allows for negative off-diagonal payoffs. Also note that for a coordination game (A1′) is equivalent to the positivity of diagonal payoffs. Under (A1) pairwise interaction is always rewarding. Under (A1′) interaction is certainly rewarding when players coordinate their actions, but may be rewarding in other cases too.

A pure strategy Nash equilibrium profile s of the social game Γ^S is called a Nash network. If the equilibrium is strict we say that s is a strict Nash network.

2.2. Static equilibria

In order to understand the dynamic behavior of the social game, we need to study static equilibria first. We focus on both strict and non-strict Nash equilibria. The reason for this is that strict equilibria turn out to play an important role in the dynamics.

A directed graph G with N vertices is said to be a *wheel* if it is an N-cycle, that is, a fully connected directed graph with N edges. For any $x \in X$ and action profile $(x_1, \dots, x_N)$ let

$$\sigma_x(x_1, \dots, x_N) = \frac{1}{N} \sum_{i,\ x_i = x} 1$$

be the fraction of players playing x, and let $\sigma(x_1, \dots, x_N) \in \Delta(X)$ be the mixed strategy defined by these probabilities.

The following theorem is the main result in this section.

Theorem 1. *Any nonempty Nash equilibrium of the social game is minimally connected and any strict Nash equilibrium has the architecture of a wheel. Moreover, in the limit as N goes to infinity, the distribution of actions $\sigma(x_1, \dots, x_N)$ approximates a symmetric equilibrium of the underlying game.*

The proof is in Appendix A. Assumptions (A1) or (A1′) allow us to first characterize the equilibrium network architecture. The argument begins by showing that in equilibrium

[5] In the sense that the set of non-generic games is closed and has zero Lebesgue measure in the space of all social games.

each component is minimally connected; then we prove that either the network is empty, or there is a single component. With the additional assumption of strict equilibrium we get that the architecture is a wheel. The result about actions then follows from the fact that in a large, connected network, each agent effectively plays against the population distribution of actions. Thus each strategy played is an approximate best response against the population distribution $\sigma(x_1, \ldots, x_N)$.

The fact that Nash architectures are connected is a consequence of network externalities: fixing actions, benefits are non-rival. However, non-strict Nash equilibria need not be wheels: One example is two equal sized wheels intersecting at a single agent who sponsors two links (with an appropriate underlying game). Strict Nash architectures are cost efficient, whereas Nash architectures, as this example shows, need not be. It is revealing to examine how the set of Nash equilibria changes as the underlying link costs change. If connection costs are low, the resulting network architecture is minimally connected. As link formation costs increase, the set of possible outcomes tends to polarize: we either have full connectivity or total fragmentation.

The following result will be used when we characterize evolutionary selection.

Corollary 1. *Let Γ be a coordination game. Then in any nonempty strict Nash equilibrium of the social game, all players coordinate on the same action x, where (x, x) is a pure strict Nash equilibrium of Γ.*[6]

3. Dynamics

We now turn to investigate the dynamic behavior of the social game. We focus on a myopic best response rule in the spirit of the unperturbed dynamics of KMR and Young (1993), and many other papers that followed. When a player comes to update, she chooses a static best response to the current play of her opponents, disregarding future adjustments. Thus agents are not fully rational.[7]

Definition 3. The best-response dynamics $Z(t)$ of the social game is defined as follows.

- At time t, the state of the system $Z(t)$ is the vector of the strategies played by each player:

  $$Z(t) = \big(s_1(t), \ldots, s_N(t)\big)$$

  where $s_i(t) = (x_i(t), N_i(t))$ is the strategy played by player i at time t.
- At each period one player is randomly selected from the population to adjust her behavior. Selection probabilities are strictly positive for each player, and selection is independent over time.

[6] Since the result is intuitive and the proof is tangential to the main focus of the paper, we omitted it. See Hojman and Szeidl (2003).

[7] A justification of the myopic best-response dynamics can be found in Fudenberg and Levine (1998, Chapter 4).

- When player i comes to revise, she chooses

$$s_i(t+1) \in \arg\max_{s_i'} \Pi_i\bigl(s_i', s_{-i}(t)\bigr)$$

while $s_j(t+1) = s_j(t)$ for $j \neq i$. When the best response of the updating player is not unique, we assume that she randomizes on the whole arg max with probabilities bounded away from zero.

This best-response dynamics defines a Markov process. We refer to the recurrent classes of this Markov process as *limit sets*. The *basin of attraction* of a limit set Ω, $D(\Omega)$, is the set of states from which the unperturbed dynamics converges to Ω with probability one.

Let us denote the set of all states where no links are maintained by E. We call this set of states the *empty network states*. In all of these states all players earn a payoff of zero regardless of the action played. Under certain conditions (large enough costs of maintaining links) the set of empty network states E will be a limit set of the dynamics. We denote by W the set of states such that the architecture of the social network is a wheel.

Theorem 2. *For N large enough, we have*:

(i) *Under Assumption* (A1)*, any limit set of the myopic best-response dynamics is either E, or a subset of W. If in addition Γ is a coordination game, then any nonempty limit set is a strict Nash equilibrium of the social game.*
(ii) *Under* (A1$'$)*, if Γ is a pure coordination game, then any nonempty limit set is a strict Nash equilibrium of the social game.*[8]

The proof of (i) is in Appendix A.[9]

In order to explain the main argument, we need to introduce some concepts. At a given state $Z(t)$, a player i is said to be a *maximal player* if her set of neighbors (including herself) $\bar{n}(i, G)$ is not strictly contained in the set of neighbors of any other player j. The *geodesic distance* between player i and j is the number intermediaries between i and j for the shortest directed path from i to j if such a path exists (infinity otherwise). Call a player j *minimal* if she is located at the maximal possible finite distance from some maximal player i. It is easy to see that the links maintained by a minimal player j are *irrelevant* for the corresponding maximal player i, in the sense that the set of neighbors of player i would not change if player j were to delete all her current links.

The idea of the proof in both (i) and (ii) is to construct a positive probability path that leads to the desired outcome. The argument relies on Lemma 3, also in Appendix A. The intuitive content of the lemma is that if a player is best responding to the current state of the population, under some circumstances it is optimal for some other agent to link to that player (directly or indirectly).

---

[8] Γ is a *pure* coordination game if it is a coordination game and $\forall x, y \in X$ we have that $u(x,x) \geqslant u(x,y)$.

[9] Since the argument for (ii) is similar, we omitted it from the paper. The details can be found in Hojman and Szeidl (2003).

The proof then proceeds as follows. Pick a player i who achieves the maximal possible payoff in a given limit set (that does not intersect with E). Clearly, i is best responding to the rest of the population. If she is not linked to the entire population, then by the lemma we can find another player who updates to link to i. Under both (i) and (ii), the latter player will achieve a strictly larger payoff than i, a contradiction. This shows that our selected player is linked and best responding to the entire population.

Next fix this player i, and select a corresponding minimal player j. The links of j are irrelevant to i who is a maximal player, hence her update does not influence who i is connected to. Because of this, under (A1), our minimal player j will find it optimal to maintain a single link to i. This makes j become a new maximal player. The same argument can now be repeated, to get that another minimal player will link our new maximal player. This chain of inductive steps is easily seen to lead to a wheel architecture. Finally, it is not difficult to show that if achieved, the wheel architecture is maintained indefinitely: once in a wheel, players adjust actions only. In the case of a coordination game, actions can be shown to converge to homogeneous play.

It is easy to check that if $\max_{x,x' \in X} u(x,x') \geqslant \phi(1)$ holds, then the empty network is never a limit set.

4. Stochastic stability in 2 × 2 coordination games

Throughout the section we focus on the case when Γ is a 2×2 coordination game. In the environments we analyze, the myopic best-response dynamics may converge to two possible limit sets corresponding to the two pure strategy equilibria of Γ. Equilibrium selection with persistent randomness along the lines of KMR and Young (1993) will help us determine which of the two social equilibria is more likely to prevail.

The action set of Γ is just $X = \{A, B\}$. The payoffs of this game are shown in Fig. 1.

We have $a > c$ and $b > d$. In line with both (A1) and (A1′) we assume that $a, b > 0$. Let α denote the mixing probability of action A for the mixed equilibrium of Γ, that is,

$$\alpha = \frac{b-d}{a-c+b-d}. \tag{2}$$

The higher α, the riskier playing A is. The profile (A, A) is the risk dominant equilibrium if and only if $\alpha < 1/2$, or equivalently $b + c < a + d$. Otherwise the profile (B, B) is the risk dominant equilibrium.[10]

	A	B
A	a, a	d, c
B	c, d	b, b

Fig. 1. Coordination game.

[10] We abstract away from the knife-edge case of equality.

4.1. Evolutionary dynamics

As usual, we choose the unperturbed dynamics to be our best-response dynamics. The evolutionary dynamics obtains by adding persistent randomness to the system. We distinguish between link mutations and action mutations, and allow the frequency of these two types of mutations to differ.

Definition 4. The perturbed or evolutionary dynamics $Z_\varepsilon^\lambda(t)$ for $\varepsilon, \lambda \geqslant 0$ is defined as follows. Typically the system evolves according to the unperturbed dynamics. At each date t, however, a mutation may take place:

- With probability ε an action mutation occurs.
- With probability ε^λ a link mutation occurs.

When a mutation occurs, one of the players is randomly selected to mutate. In the case of a link mutation she chooses randomly the set of players to link to. Her choice is such that a set of direct neighbors with infinite cost will never be chosen. In the case of an action mutation, she chooses randomly the action to play. Except for the restriction on infinite costs, all possibilities occur with positive probability in both types of mutations. Also, all players are chosen to mutate with probability bounded away from zero.

The parameter λ measures the frequency of an action mutation relative to that of a link mutation. The higher λ, the less frequent link mutations are. The smaller ε, the less frequent both kinds of mutations are.

Definition 4 implies that the dynamics defines an irreducible Markov process on the subset of states where no player maintains infinitely costly links.[11] We are interested in the limit of the invariant distribution of the Markov process as the noise level ε goes to zero. As is well known, the limit distribution will be concentrated on the limit sets of the unperturbed dynamics. The set of states that have positive probability under the limit distribution are called *stochastically stable*. In our case, then, the limit distribution will be concentrated on the A-wheel and the B-wheel. We say that a limit set is uniquely selected if the limit distribution is concentrated (with probability one) on that limit set only. Intuitively this means that for ε small the perturbed dynamics spends most of the time in that equilibrium.

Our analysis makes use of the *radius/coradius* approach introduced by Ellison (2000). Some preliminary definitions are necessary. The cost $c(Z, Z')$ of a mutation from state Z to state Z' is a measure of the unlikeliness of that mutation occurring.[12] The cost of a path $Z_1, Z_2, \ldots, Z_k$ is the sum of the costs of individual transitions. The radius of a limit set Ω, denoted by $R(\Omega)$, is the minimum cost along any path that escapes $D(\Omega)$. The coradius of Ω, denoted by $CR(\Omega)$, is the maximum over all other states of the minimum cost of mutations necessary to reach $D(\Omega)$. The radius measures how difficult it is to leave a limit

[11] States where a player maintains infinitely costly links never occur; they are avoided by mutation as well as by updating. This is in line with the idea that players do not make infinitely costly mistakes.

[12] Formally, it equals $\lim_{\varepsilon\to 0}[\log P(Z' \mid Z)/\log\varepsilon]$ where $P(Z' \mid Z)$ is the probability of a transition from Z to Z' in one period.

set, the coradius provides a measure of its attractiveness. Ellison's theorem establishes that a limit set Ω is stochastically stable if $R(\Omega) > CR(\Omega)$; if there are only two limit sets this condition is also sufficient. Intuitively, if it is less costly to enter the basin of Ω than to leave it, Ω is stochastically stable. The theorem also provides a measure of the speed of convergence of the perturbed process in terms of the coradius.

4.2. Equilibrium selection

We focus on three environments that illustrate how selection is influenced by both payoffs and the cost structure of links.[13] The first environment, which we call *costly mis-coordination with cheap links*, assumes that off-diagonal payoffs in the underlying game Γ are negative, that is, $c, d < 0$. Lack of coordination in a game like this is heavily punished, and no player would maintain a link to another one for the sole reason of their pairwise interaction if they are not coordinating. We also assume that link costs are very small, so that for all practical purposes they can be taken as equal to zero (but all else equal, less links are preferred to more).[14]

Our second environment, *positive payoffs with cheap links*, assumes that the payoffs in the underlying game are all positive. In this specification, it is beneficial for whoever comes to update to link to all opponents (directly or indirectly) because links are costless and pairwise interaction is unquestionably good.

In the third environment, *expensive links*, we assume that although maintaining one link is very cheap, maintaining two or more links is prohibitively costly. Regarding the underlying game we assume all payoffs are positive in order to shut down the effect of costly miscoordination. In this environment, we also assume that the unperturbed dynamics is conservative. That is, along the dynamics the number of different maximal sets does not increase (where a maximal set is just the set of neighbors of a maximal player).[15]

For N large, Corollary 1 implies that for each of our three environments, the social game has two strict Nash equilibria, each having all players coordinate on one of the two actions: the 'A-wheel' and the 'B-wheel.' Moreover, Theorem 2 implies that these are the only limit sets of the best-response dynamics: in the first environment the underlying game is a pure coordination game satisfying (A1$'$); in the second and third environments (A1) holds.

The main result of the paper is the following theorem.

Theorem 3. *For N large enough, in both the costly-miscoordination/cheap-links and the expensive links environment, the A-wheel is the unique long-run equilibrium if and only if*

$$\frac{a}{b} > \frac{\alpha}{1-\alpha}. \tag{3}$$

[13] A general characterization seems very difficult to obtain.

[14] The effects of this assumption are easily reproduced by choosing a strictly increasing cost function that is very small relative to the absolute magnitude of all payoffs in the underlying game.

[15] We believe that all results continue to hold if we relax the conservative assumption. However, we have not been able to prove this.

In the positive-payoffs/cheap-links environment for N large enough, the risk dominant equilibrium is uniquely selected.

The proof of the costly-miscoordination/cheap-links case is in Appendix A. We save the reader from the details for the expensive links case, as it relies on the same ideas (see Hojman and Szeidl, 2003). The third case turns out to be quite straightforward.

Inequality (3) is a parsimonious expression of a trade-off between risk dominance and payoff dominance. The left-hand side, a/b, is a measure of the payoff advantage of strategy A relative to strategy B. The right-hand side, $\alpha/(1-\alpha)$, is a measure of the riskiness of strategy A. The inequality says that the A-wheel is uniquely selected if and only if action A is sufficiently advantageous in terms of payoffs to compensate for its riskiness. That is, the inequality is an arithmetic expression of the trade-off between payoff dominance and risk dominance. If the inequality fails to hold, the reverse inequality is equivalent to the corresponding condition for the B-wheel. Thus abstracting away from the knife-edge case of equality, we always have unique equilibrium selection for N large.

The trade-off between risk dominance and payoff dominance means that if there is a strategy that is both risk dominant and payoff dominant, it will be uniquely selected. However, a strategy that is not risk dominant may be uniquely selected if it delivers sufficiently high payoffs; and conversely, a payoff dominated strategy may also be uniquely selected if it is not too risky.

In the costly-miscoordination/cheap-links case, the intuition for the trade-off can best be understood by means of the following example. Imagine a player facing a large cluster of opponents, mostly coordinating on the inefficient action, but some of them coordinating on the efficient one. Moreover, the geometry of the network is such that not only is our player able to connect to all other players, she is also able to connect to the efficient players only, avoiding interaction with the inefficient ones.

Connecting to the whole set of players would provide the benefits of scale stemming from the large number of inefficient players but at a cost of suffering from miscoordination with the efficient opponents. On the other hand, avoiding the inefficient players would provide a benefit of coordinating with a smaller number of players on the efficient equilibrium. The payoff of the former depends on risk-dominance considerations; that of the latter is related to payoff dominance. Thus the player is effectively trading off risk-dominance against payoff dominance in her choice, as captured by (3).

In the expensive links environment, the basis for (3) is the decision problem of a player who is facing two clusters of players, a large one coordinating mostly on the inefficient equilibrium, and a smaller one coordinating on the efficient one. She can only link one of those, thus trading off the scale effects stemming from the large but inefficient cluster against the benefits of the efficient cluster. As the large cluster is not completely coordinated, risk-dominance considerations also influence the choice.

In the positive-payoffs/cheap-links case, whenever a player updates, she optimally links the whole population. It follows that in terms of actions chosen the environment is the same as that studied by KMR: each updating player is best responding to the distribution of play in the whole population. Thus the unique equilibrium selected will correspond to the strategy that performs best against an average population, that is, the risk dominant one.

5. Conclusion

This paper built a model of social games with endogenous network formation. In the static case, Nash equilibria are minimally connected, with play approximating a symmetric equilibrium. A strict Nash network must be a wheel, which is also the unique cost efficient architecture. We also find that a myopic best-response dynamics eventually selects a wheel architecture, though convergence in actions is not always guaranteed.

In our analysis of stochastic stability we identified a trade-off between risk dominance and payoff dominance due to the presence of scale effects. Our findings differ from those obtained by Goyal and Vega-Redondo (2004) for a two-way flow model. In their model, for low costs of link formation the risk dominant equilibrium is selected; for medium costs the payoff dominant equilibrium is selected; and for high costs only the empty network is stochastically stable. They do not get the trade-off identified by the this study. Nevertheless, the selection pattern in their paper and in ours is similar. For low costs of links and positive payoffs we both get the risk dominant equilibrium. Additionally, as link costs increase, payoff dominance starts to matter more for selection in both models.

Standard selection results in the literature, as well as our results, seem to fit into the following general pattern. As the models restrict the selection of opponents less, so that choosing them becomes a critical decision, risk-dominance considerations matter less for equilibrium selection; payoff considerations, on the other hand, matter more. With fixed interactions, or low costs and positive payoffs, the choice of opponents is completely determined and risk dominance provides equilibrium selection. Indeed, with positive-payoffs/cheap-links, and likewise in the low cost case of Goyal and Vega-Redondo (2004), everybody will link to everybody else anyway and the risk dominant equilibrium is selected. When players' link decisions become important, either because links are more expensive or because miscoordination is costly, payoff considerations start to matter. Our first and third environments fit into this category. Finally, when players can essentially choose the equilibrium they play, as in Ely (2002), payoff considerations will dominate. In Ely (2002) there are no scale effects, so if one isolated player is playing the efficient equilibrium, everybody will connect and coordinate with her. In short, these models show a tendency that flexibility favors payoff dominance. In this sense our study provides an interesting qualification to the standard results of both KMR and Ely (2002) and finds a unifying pattern that incorporates and extends those results.

Acknowledgments

We are grateful to Drew Fudenberg for helpful discussions, Markus Mobius, Miklos Koren and participants of the Harvard Theory Lunch and the 2002 Latin American Meeting of the Econometric Society for comments.

Appendix A

Proof of Theorem 1

Lemma 1. *If Γ satisfies either* (A1) *or* (A1′)*, then in any Nash equilibrium the network architecture G is either the empty network, or minimally connected. Moreover, any strict Nash equilibrium is a wheel.*

Proof. We start by introducing some terminology. A subset of players $\mathcal{V}$ is said to be *recurrent* if: (i) for any $i, j \in \mathcal{V}$, $i \mapsto j$; and (ii) for any $i \in \mathcal{V}$ and $k \notin \mathcal{V}$, $i \nrightarrow k$. In words, every player in the set is connected to any other player in the set and there are no connections exiting the set. There are two differences between a recurrent set and a component. First, a recurrent set has to be connected in the directed sense, whereas a component only needs to be connected in the non-directed sense. Second, there may be connections entering a recurrent set, but no links enter or exit a component. Note that $\mathcal{V} = \{i\} \cup n(i, G)$ for any $i \in \mathcal{V}$.

Step 1: *Any component of a Nash network is a recurrent set.*

Let l designate a player in the component who has the smallest number of neighbors, i.e., $|n(l, G)|$ is minimal within the component to which l belongs. Denote by $x_l \in X$ player l's equilibrium action. Since $|n(l, G)|$ is a component minimum it follows that the set of players $H(l) = \{l\} \cup n(l, G)$ is recurrent.[16] We need to show that this subset of players includes the whole component. Assume not; then we consider two subcases depending on whether l maintains any link or not.

If l maintains at least one link, then because players in $H(l)$ have the minimum number of neighbors within the component, it must be the case that for some $j \notin H(l)$ in the component and some $w \in H(l)$, j is adjacent to w. Without loss of generality, suppose l has a direct link to w. We claim that l is better off breaking his link with w and establishing a direct link with j. Under (A1), it is immediate to check that player l does strictly better by doing so and playing action x_l. If (A1) does not hold but (A1′) does, the same conclusion applies replacing x_l by x_j, the latter being the action played by j. To see this, let Π_j and Π_l respectively denote the payoffs of player j and player l in the current configuration. Note that in a Nash equilibrium, we must have that

$$\Pi_j \geqslant \Pi_l + u(x_l, x_l),$$

where the right-hand side is exactly what j would receive if she deviated by removing all her links, connecting l and playing x_l. Similarly, if player l removed her link to w and connected to j instead, her payoff, Π_l would satisfy

$$\Pi_l \geqslant \Pi_j + u(x_j, x_j) - u(x_j, x_l),$$

where the right-hand side is what she would obtain by choosing x_j. Using these inequalities, a sufficient condition for the above deviation to be profitable for player l is (A1′).

[16] Observe that given $l' \in n(l, G)$, $l \to l''$ for any $l'' \in n(l', G)$. Note the set may consist of only l.

In the case where l maintains no links, so that $H(l)$ is a singleton, assume there is another player j in the component who links to l. If j maintains no other links, then equilibrium requires that $\phi(1) < u(x_j, x_l)$ (equality can only occur in the non-generic case) which implies that l would be better off linking back j. If j maintains other links, then the net payoff from these other links is strictly positive; but then l could exactly mimic these other links of j and achieve the same or higher payoff by the convexity of $\phi(\cdot)$.

Step 2: *Any Nash equilibrium is either the empty network, or minimally connected.*

We have seen that any Nash equilibrium is a disjoint union of minimally connected subgraphs, since each component is a recurrent set, and nobody is paying for irrelevant links. If not all of these are empty networks, then none of them are, otherwise any player with no links would have a profitable deviation of linking somebody who already maintains a link. Next suppose that w and z are two players belonging to disjoint nonempty components. Let x_w and x_z be their equilibrium actions and Π_w and Π_z denote the players' equilibrium payoffs. If player w breaks all her links, connects to z and chooses the same action played by z, she gets a payoff of at least $\Pi_z + u(x_z, x_z)$. Nash equilibrium requires $\Pi_w \geqslant \Pi_z + u(x_z, x_z)$ or $\Pi_w - \Pi_z \geqslant u(x_z, x_z)$. Observe that (independently) (A1) and (A1$'$) imply $u(x_z, x_z) > 0$, from which $\Pi_w - \Pi_z > 0$. The analogue deviation for player z entails $\Pi_z \geqslant \Pi_w + u(x_w, x_w)$ or $\Pi_w - \Pi_z \leqslant -u(x_w, x_w)$. Hence, $\Pi_w - \Pi_z < 0$, a contradiction. Thus if a Nash equilibrium is not the empty network, it is minimally connected.

Step 3: *In a strict Nash network every player is connected to some other player.*

Assume not and consider a player, say i, who is not connected to any other player and therefore receives a payoff equal to 0 regardless of the action she chooses. Since i is indifferent between any action this cannot be a strict Nash equilibrium.

Step 4: *A strict Nash network is a wheel.*

From Step 3, we know that a strict Nash equilibrium is not the empty network. From Step 2, therefore, it has to be minimally connected. If it is not a wheel, there exists a player i who is paying for links to at least two other players, j and k. Since the set is connected, $j \to i$ and $k \to i$. Suppose that the paths connecting j to i and k to i are disjoint and let j_i denote the player adjacent to i in the path from j to i. It is clear that j_i is indifferent between paying for a connection with i and connecting k instead. Similarly, if the paths connecting j to i and k to i are not disjoint, minimality implies that when these paths first intersect at some node l they coincide thereafter. Choose j_l to be adjacent to l in the path from j to i. The previous indifference argument applies to (directed) links (j_l, l) and (j_l, i). □

Definition 5. The *reduced game* $\Gamma(N)$ has the same set of players N, with pure strategy spaces X each. The payoff function $U_i : \times_{i=1}^{N} X \to \Re$ of player i is defined to be

$$U_i(x_i, x_{-i}) = \sum_{j \neq i} u(x_i, x_j). \tag{A.1}$$

Clearly, for any nonempty Nash equilibrium of the social game, the action profile chosen has to be a Nash equilibrium of the reduced game; moreover, for any strict Nash equilibrium of the social game, the action profile has to be a strict equilibrium of the reduced game.[17]

For any action profile $(x_1, \dots, x_N)$ of the reduced game, the payoff of player i is given by

$$U_i(x_i, x_{-i}) = \sum_{j=1}^{N} u(x_i, x_j) - u(x_i, x_i) = Nu(x_i, \sigma) - u(x_i, x_i)$$

where $\sigma = \sigma(x_1, \dots, x_N)$.

Lemma 2. *The action profile $(x_1, \dots, x_N)$ is a Nash equilibrium of the reduced game $\Gamma(N)$ if and only if for all x_i and $x' \in X$ we have*

$$u(x_i, \sigma^*) \geqslant u(x', \sigma^*) + \frac{1}{N}\bigl[u(x_i, x_i) - u(x', x_i)\bigr]. \tag{A.2}$$

Here $\sigma^* = \sigma(x_1, \dots, x_N)$.

Proof. Player i gets $Nu(x_i, \sigma^*) - u(x_i, x_i)$ if she sticks to playing action x_i. If she deviates and plays $x' \in X$, she gets $Nu(x', \sigma^*) - u(x', x_i)$. Thus Nash equilibrium requires

$$Nu(x_i, \sigma^*) - u(x_i, x_i) \geqslant Nu(x', \sigma^*) - u(x', x_i).$$

Rearranging gives the result. □

Then it is easy to see that for any Nash equilibrium $(x_1, \dots, x_N)$ of the reduced game $\Gamma(N)$, the mixed strategy $\sigma(x_1, \dots, x_N)$ is a symmetric ε-equilibrium of the underlying game Γ, where ε is of order $1/N$. This is because the last term in (A.2) is of order $1/N$.

Moreover, if we have a selection of Nash equilibria $(x_1, \dots, x_{N_k})$ of the reduced game with N_k players where N_k goes to infinity, such that the corresponding mixed strategies $\sigma(x_1, \dots, x_{N_k})$ converge to some mixed strategy σ^*, then σ^* is a symmetric equilibrium of the underlying game Γ. Thus for a large population, we have that the typical Nash equilibrium of the reduced game has to correspond to symmetric equilibria of the underlying game Γ.

The statement of the theorem follows from Lemmas 1 and 2.

Proof of Theorem 2

Lemma 3. *Under either* (A1) *or* (A1′), *if some player i, who is best responding to the rest of the population, maintains at least one link, and is not linked to player j, and now j comes to update, then player j's best response involves linking some player connected to player i. Furthermore, the payoff of j after updating will be strictly higher than the payoff of i before updating and the payoff of i does not change.*

[17] In Hojman and Szeidl (2003) we show that a strict Nash equilibrium $(x_1, \dots, x_N)$ of the reduced game will induce a strict Nash equilibrium of the social game if and only if $\min_i U_i(x_i, x_{-i}) > \phi(1)$.

Proof. Suppose i has updated at time $t = 0$. Towards a contradiction, assume j links some player k_j such that $k_j \to i$.[18] This means that j does not make use of player i's links. Note that this strategy was also available for player i when updating at $t = 0$. It follows that

$$\Pi_i\big(s_i(1), s_{-i}(0)\big) \geqslant \Pi_j\big(s_j(2), s_{-j}(1)\big). \tag{A.3}$$

On the other hand, by deleting her current links, initiating a link with i and using action $x_i(1)$, player j gets a payoff equal to $\Pi_i(s_i(1), s_{-i}(0)) + u(x_i(1), x_i(1))$. Thus, since we assumed j's best response to $s_{-j}(1)$ does not include this possibility, it must be the case that

$$\Pi_j\big(s_j(2), s_{-j}(1)\big) > \Pi_i\big(s_i(1), s_{-i}(0)\big) + u\big(x_i(1), x_i(1)\big). \tag{A.4}$$

Adding (A.3) and (A.4) it follows that $u(x_i(1), x_i(1)) < 0$ which violates both (A1) and (A1′). We conclude that if player j, as previously defined, comes to update at $t = 2$ (with positive probability) her revised strategy involves linking some player k_j such that $k_j \to i$ or $k_j = i$. □

Lemma 4. *Under Assumption* (A1)*, any limit set of the myopic best-response dynamics is either E, or a subset of W.*

Proof. For any state Z, denote the maximal payoff in the population by $M(Z)$. Let the set of states achievable along the dynamics from the current state be H. Denote the maximum of $M(Z)$ on H by $\overline{M}$. Because the underlying game is finite, there exist (potentially many) states in H where $M(Z)$ assumes $\overline{M}$. Let Z be one of them. By assumption, Z has the property that the maximal payoff in any state that can be achieved from Z is at most $\overline{M}$.

Let player 1 be a player who receives $\overline{M}$ in Z. It is clear that she is best responding. If player 1 is maintaining no links, then $\overline{M}$ equals zero, and with probability one the system will get to a state in E. This is because all players can achieve a zero payoff by dropping all links, and by assumption the maximal payoff any one of them can achieve is also zero. Once the system is in an empty network state, it will never leave it (in the generic case), because no player can get a positive payoff by maintaining a link (indifference would be non-generic). On the other hand, in E, any state can be reached from any other state. Thus, when player 1 maintains no links, the limit set the system converges to is E.

Suppose 1 maintains at least one link. If she is not linked to the whole population, there would be a player 2 she is not linked to; if that player updates next, by Lemma 3, her payoff will be strictly higher then that of 1, which cannot happen. Thus a player who receives $\overline{M}$ is always linked to the whole population and best-responding.

It follows that state Z has players being linked to the entire population who are best responding. Let 1 be such a player.

Let 3 be a player at maximum geodesic distance from 1 at state Z. Pick 3 to update now. Note that by construction, 3's current links are irrelevant for player 1. Also note that for N large enough, one best response of 3 is to link to 1 and play a best response to the

[18] Given player i chose to link somebody, it cannot be optimal for player j not to maintain any links: by linking i and playing x_i player j attains a payoff that exceeds i's payoff by $u(x_i, x_i)$, which is strictly positive.

rest of the population. Indeed, as all pairwise payoffs are positive, for N large maintaining a single link that connects to the whole population is the most profitable thing to do.

Thus 3 now links to 1 with positive probability. Now 3 is linked to the rest of the population and is best responding. Picking a player 4 who is at maximal geodesic distance from 3, and repeating the above argument, we find that with positive probability 4 will drop all her links and link to 3. And so on. Hence, under (A1), we can find a positive-probability path to a wheel by repeating this argument. It takes at most $N - 1$ more iterations to reach a state in W.

At $Z(t) \in W$, any player i is connected to the entire population by a single link. An argument similar to the one just presented ensures that a best response of i at $Z(t)$ must involve maintaining her current link: by doing so player i remains connected to the entire population paying for a single link. Hence, $Z(t + 1) \in W$. □

The best-reply graph of a game is *weakly acyclic* if starting from any strategy profile there exists a directed path in the best-reply graph to a strict Nash equilibrium. It is well known (Young, 1993) that if this property holds then any limit set of the best-response dynamics is a singleton.

Lemma 5. *Under* (A1)*, if the best-reply graph of the reduced game* $\Gamma(N)$ *is weakly acyclic, then any nonempty limit set of the myopic best-response dynamics in the social game is a strict Nash equilibrium.*

This is because under (A1), the previous lemma implies that a nonempty limit set has a wheel architecture. Since in a wheel all players interact with each other, to establish convergence of actions, we can focus on the best-response dynamics of the reduced game.

The second statement in Theorem 2(i) follows from the fact that if Γ is a coordination game, then the reduced game is weakly acyclic. For a proof, see Hojman and Szeidl (2003).

Proof of Theorem 3

Proof. We only deal with the costly-miscoordination /cheap-links environment. The strategy of the proof is as follows. First we give a lower bound on the radius of the A-wheel. Then we construct a path leaving the basin of attraction of the A-wheel that demonstrates that our lower bound is essentially exact. Finally, we make use of Ellison's radius-coradius theorem.

Suppose we start the system from the A-wheel, and under the perturbed dynamics, at time t, the total cost of mutations taken place so far is $m = \lambda l + q$, where l is the number of link mutations and q is the number of action mutations. Now, if from this state the unperturbed dynamics reverts back to the A-wheel with probability one, then m is clearly a lower bound for the radius.

Additionally, if after a total mutation cost of m no updating under the unperturbed dynamics makes the number of agents playing B increase, then with probability one the system will revert to the A-wheel. This is because there are only two absorbing states, the A-wheel and the B-wheel, and the unperturbed dynamics settles the system in one of them in the long run. So if the number of Bs does not increase (and $m < N$) then the absorbing state selected has to be the A-wheel. Indeed, there cannot be more than m players playing B because by our assumption no mutation takes place and nobody wants to switch to B under the unperturbed dynamics.

So if we can find an m such that after a total cost of mutations less than or equal to m no player wants to switch to B, we have found a lower bound for the radius. Suppose we have had l links and q action mutations, with total cost at most m. Then the best payoff a player who is to update can get from switching to B is at most qb. Indeed, if she is able to connect to only the B players, she gets this amount; but if she has to connect (possibly indirectly) to some A players too, she will get less (since by assumption $c, d < 0$).

Now, the best she can get if she keeps on playing A is at least $(N - q - 1)a + qd$. She gets this payoff if she connects to all other players and continues to play A.

She will certainly play A if the latter of these two numbers is larger. We are looking for the maximal q where this inequality is not yet violated, that is, the $\bar{q}$ that makes the two numbers equal: $\bar{q}b = (N - \bar{q} - 1)a + \bar{q}d$ or

$$\bar{q} = (N - 1)\frac{a}{a + b - d}.$$

It follows that $\bar{q}$ serves as a lower bound for the radius.

To get an upper bound, consider the following path. The system is initially in the A-wheel, then at date $t = 0$ simultaneously $[\bar{q}] + 1$ (the brackets stand for integer value) consecutive members of the wheel mutate to B, and the first of them (who is seen by all others) drops her link to the next player. More formally, players $1, 2, \dots, [\bar{q}] + 1$ mutate to B, and player 1 drops her link to player N. This mutation has a total cost of $[\bar{q}] + 1 + \lambda$.

We claim that the resulting configuration is not in the basin of attraction of the A-wheel anymore. To see why, consider the following pattern of updating. First player N comes to update. Her two best options are either to connect to the B-sequence of the configuration and switch to B, which gives a benefit of $([\bar{q}] + 1)b$, or to continue with A, but then her best pick is to stay where she is, i.e., indirectly being connected to all other players. This gives a payoff of $(N - [\bar{q}] - 2)a + ([\bar{q}] + 1)d$. Since $[\bar{q}] + 1 > \bar{q}$, it follows that she prefers the first choice, and hence switches to B and connects directly to the end of the B-sequence.

Suppose that player $N - 1$ updates next. She has a similar choice, only that the length of the B-sequence has increased by one, and the length of the A-sequence has decreased by one. Thus she is going to switch to B and connect to the B-sequence too. If then player $N - 2$, then $N - 3$, and so on, are to update, in the end we will have everybody playing B. Then they will form a wheel by the usual properties of the unperturbed dynamics. Finally, this pattern of updating has positive probability under the unperturbed dynamics. Thus $[\bar{q}] + 1 + \lambda$ is an upper bound for the radius.

So we have proved that the radius of the A-wheel, $R(A)$, satisfies

$$(N - 1)\frac{a}{a + b - d} \leqslant R(A) \leqslant (N - 1)\frac{a}{a + b - d} + 1 + \lambda.$$

Next note that since there are only two absorbing states, we also have that the radius of the B-wheel is equal to the coradius of the A-wheel. By symmetry, our previous arguments can be repeated for the basin of attraction of the B-wheel to give bounds for $R(B)$ or equivalently $CR(A)$. These bounds turn out to be

$$(N-1)\frac{b}{a+b-c} \leqslant CR(A) \leqslant (N-1)\frac{b}{a+b-c} + 1 + \lambda.$$

By Ellison's radius–coradius theorem, we have that if $R(A) > CR(A)$ then the A-wheel is the unique long-run equilibrium. Using our bounds, we see that as long as

$$(N-1)\frac{a}{a+b-d} > (N-1)\frac{b}{a+b-c} + 1 + \lambda$$

we will have the A-wheel selected. Now, for any fixed game and λ this last formula is certainly implied by $a/(a+b-d) > b/(a+b-c)$ for N sufficiently large.

By a completely symmetric argument one can prove that if this inequality is reversed, then for N large enough the B-wheel is uniquely selected in the long run. Ignoring the knife-edge case of equality, reorganizing these inequalities gives that as long as (3) holds, for N large enough the A-wheel is uniquely selected in the long run. The proof is complete. □

References

Bala, V., Goyal, S., 2000. Self-organization in communication networks. Econometrica 68, 1181–1230.

Ellison, G., 1993. Learning, local interaction and coordination. Econometrica 61, 1047–1071.

Ellison, G., 2000. Basins of attraction, long-run equilibria, and the speed of step-by-step evolution. Rev. Econ. Stud. 67 (1), 17–45.

Ely, J., 2002. Local conventions. Adv. Theor. Econ. 2 (1).

Fudenberg, D., Levine, D., 1998. The Theory of Learning in Games. MIT Press, Cambridge, MA.

Galeotti, A., Goyal, S., Kamphorst J., 2004. Network formation with heterogeneous players. Games Econ. Behav. In press.

Goyal, S., Vega-Redondo, F., Network formation and social coordination. Games Econ. Behav. In press.

Granovetter, M., 1973. The strength of weak ties. Amer. J. Sociology 78 (6), 1360–1380.

Hojman, D.A., Interaction structure and the evolution of conventions. Working paper. Harvard University.

Hojman, D.A., Szeidl, A., 2003. Endogenous networks, social games and evolution. Working paper. Harvard University and UC at Berkeley.

Hojman, D.A., Szeidl, A., 2004. Core and periphery in endogenous networks. Working paper. Harvard University and UC at Berkeley.

Jackson, M.O., Watts, A., 2002. On the formation of interaction networks in social coordination games. Games Econ. Behav. 41, 265–291.

Kandori, M., Mailath, G., Rob, R., 1993. Learning, mutation and long-run equilibria in games. Econometrica 61, 29–56.

Oechssler, J., 1999. Competition among conventions. Math. Computat. Organ. Theory 5, 31–44.

Young, H.P., 1993. The evolution of conventions. Econometrica 61, 57–84.