Learning, Network Formation and Coordination

Sanjeev Goyal* Fernando Vega-Redondo†

Econometric Institute Report 9954/A.
November 1999.This version: December 10th, 1999.

Abstract

In many economic and social contexts, individuals can undertake a transaction only if they are 'linked' or related to each other. We take the view that these links are costly, in the sense that it takes effort and resources to create and maintain them. The link formation decisions of the players define a network of social interaction. We study the incentives of individuals to form links and the effects of this link formation on the nature of social coordination.

Our analysis shows that equilibrium networks have simple architectures; they are either complete networks or stars. Moreover, the process of network formation has powerful effects on social coordination. For low costs of forming links all individuals coordinate on the the risk-dominant action, while for high costs of forming links individuals coordinate on the efficient action.

---

*Econometric Institute, Erasmus University, 3000 DR, Rotterdam, The Netherlands. E-mail address: goyal@few.eur.nl

†Facultad de Económicas and Instituto Valenciano de Investigaciones Económicas, Universidad de Alicante. E-mail address: vega@merlin.fae.ua.es.

1 Introduction

In recent years, several authors have examined the role of interaction structure – different terms like network structure, neighborhood influences, and peer group pressures, have been used – in explaining a wide range of economic phenomena. This includes work on social learning and adoption of new technologies, evolution of conventions, collective action, labor markets, and financial fragility.[1] The broad message of this literature is that interaction structure matters in important ways. This leads us to examine the reasonableness/robustness of different structures and is the primary motivation for developing a model in which the evolution of the interaction structure is itself an object of study.

We propose a general approach to study this question. We suppose that individual entities can undertake a transaction only if they are 'linked'. This link may refer to a social or a business relationship, or it may refer simply to awareness of the others. We take the view that links are costly, in the sense that it takes effort and resources to create and maintain them. This leads us to study the incentives of individuals to form links and the implications of this link formation for aggregate outcomes.

In the present paper, we apply this general approach to the study of social coordination. There is a group of players, who have the opportunity to play a 2×2 coordination game with each other. We start with the case where two players can only play with one another if they have a *direct pair-wise link*. These links can be made on individual initiative but are costly to form. So each player prefers that others incur the cost and form links with him. For simplicity, the game is assumed to yield only positive payoffs in every bilateral interaction. Individuals care about aggregate payoffs and, therefore, they always accept any link supported (i.e. paid) by some other player. The link decisions of different players define a network of social interaction. In addition to the choice of links, each player also has to choose an action that

[1] See e.g., Allen and Gale (1998), Anderlini and Ianni (1997), Bala and Goyal (1998), Chwe (1996), Coleman (1966), Ellison and Fudenberg (1993), Ellison (1993), Ely (1996), Goyal and Janssen (1997), Granovetter (1974), Haag and Lagunoff (1999), and Morris (1997), among others.

she will use in all the games that she will engage in. We are interested in the nature of networks that emerge and the effects of link formation on social coordination.

In our setting, links as well as actions in the coordination game are chosen by individuals on an independent basis. This allows us to study the social process as a non-cooperative game. We start by examining the Nash equilibrium of the static game. We find that a variety of networks – including the complete network, the empty network and partially connected networks – can be supported in equilibrium. Moreover, the society can coordinate on different actions and conformism as well as diversity with regard to actions of individuals is possible in Nash equilibrium. This multiplicity motivates an examination of the dynamic stability of different outcomes.

We develop a dynamic model in which, at regular intervals, individuals choose links and actions to maximize their payoffs. Occasionally they make errors or experiment. Our interest is in the nature of long run outcomes, when the probability of these errors is small. We find that the dynamics generate clear-cut predictions both concerning the architecture of networks as well as regarding the nature of social coordination.

In particular, *except for the case where costs of link formation are very high, every pair of players is linked and the complete network is the unique stable network.* Figure 1a gives an example of a complete network in a society with 4 players. This result also shows that partially connected networks are not stable. We also find that, if players are at all connected, they always coordinate in the long run on the same action, i.e. social conformism obtains. However, the nature of coordination depends on the costs of link formation. *For low costs of link formation, players coordinate on the risk-dominant action, while for high costs of link formation they coordinate on the efficient action.* Thus our analysis reveals that, even though the eventual network is the same in all cases of interest, the process of network formation itself has serious implications for the nature of social coordination.

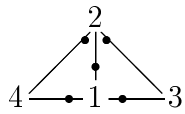

Figure 1a
Complete Network

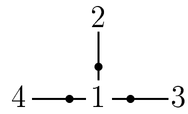

Figure 1b
Center-sponsored Star

These results are obtained in a model where two players can only play a game if they have a direct link between them. We also briefly examine a variation of this model, *in which two players can play a game if they are directly or indirectly linked with each other.*[2] In this setting, we find that *the center-sponsored star is the unique strict equilibrium architecture.* This is a network in which one player forms a link with every other player and pays for all the links. Figure 1b provides an example of this architecture for a society with 4 players. We also find that the costs of link formation matter for the nature of social coordination in a pattern which is similar to the basic model.

To summarize: we develop a framework to study the co-evolution of the network of interaction and the actions in a coordination game. Our analysis shows that this framework is tractable. In particular, we find that equilibrium networks possess simple architectures – they are complete networks or, in a variation of the basic model, stars. Moreover, the process of network formation has powerful effects on the nature of social coordination. For low costs of link formation everyone chooses the risk-dominant action, while for high costs everyone chooses the efficient action.

We now place the paper in context. Traditionally, sociologists have held the view that individual actions, and in turn aggregate outcomes, are in large part determined by interaction structure. By contrast, economists have tended to focus on markets, where social ties and the specific features of the interaction structures are typically not important. In recent years, economists have examined in greater detail the role of interaction structure and found that it plays an important role in shaping important economic phenomena (see the references given above, and also Granovetter, 1985). This

[2] More precisely, two players can play a game with each other if there is a path between them in the social network.

has led to a study of the processes through which the structure emerges. The present paper is part of this general research program.

We relate the paper to the work in economics next. We suppose that an individual players can form pair-wise links by incurring some costs, at their own initiative, i.e., link formation is one-sided. This allows us to model the network formation process as a non-cooperative game. This element of our model is similar to the work of Goyal (1993) and Bala and Goyal (1999). Related work on network formation includes Jackson and Wolinsky (1996) and van den Nouweland (1994). The primary contribution of the present paper is the development of a common framework within which the emergence of interaction networks and the selection of equilibrium can be suitably studied.[3]

There has been considerable work on equilibrium selection. From an evolutionary viewpoint, this work includes Blume (1993), Kandori, Mailath and Rob (1993), and Young (1993), among others.[4] The general insight of this work is that when players are uncertain about the strategies of their opponents, risk-dominance considerations tend to prevail over those of efficiency, and an inefficient but risk-dominant equilibrium becomes stable in the long run. However, this original finding has been re-examined by several authors, and the result has been shown to be sensitive to different assumptions, such as the nature of the strategy revision rule (Robson and Vega-Redondo,1996), the costs of flexibility (Galesloot and Goyal, 1997), precise modelling of mutation (Bergin and Lipman, 1996) and the role of mobility of players across locations (Bhaskar and Vega-Redondo, 1998; Ely, 1996; Mailath, Samuelson and Shaked,1994; Oechssler, 1997). Our paper is related to this last strand of work.

The basic insight to be gained from the evolutionary literature on player mobility can be summarized as follows. If individuals can separate/insulate themselves easily from those who are playing an inefficient action (e.g., the

[3]In independent work, Jackson and Watts (1999) have developed a related model which addresses similar concerns. We discuss their work in the concluding section.

[4]For a consideration of this same equilibrium selection problem from a different ("eductive") perspective, the reader may refer to the classical book of Harsanyi and Selten (1988) or the more recent paper by Carlson and van Damme (1993).

risk-dominant action), then efficient "enclaves" will be readily formed and eventually attract the "migration" of others (who will therefore turn to playing efficiently). Heuristically, one may be inclined to establish a certain parallelism between easy mobility and low costs of forming links. However, the considerations involved in each case turn out to be very different, as is evident from the sharp contrast between our conclusions (recall the above summary) and those of the mobility literature.

There are two main reasons for this contrast. First, in our case, players do not *indirectly* choose their pattern of interaction with others by moving across a *pre-specified* network of locations (as in the case of player mobility). Rather, they construct *directly* their interaction network (with no exogenous restrictions) by choosing those agents with whom they want to play the game. Second, the cost of link formation (which are payed per link formed) act as screening device that is truly effective only if it is high enough. In a heuristic sense, we may say that it is precisely the restricted "mobility" it induces in that case which helps insulate (and thus protect) the individuals who are choosing the efficient action. If the link-formation cost is too low, the extensive interaction it facilitates may have the unfortunate consequence of rendering risk-dominance considerations decisive.

The rest of this paper is organized as follows. Section 2 describes the basic model in which only directly linked players can play a game, while Section 3 presents the results. Section 4 studies the case where players can play a game if they are either directly or indirectly connected to each other. Section 5 concludes. For expositional simplicity, all proofs have been placed in the appendices.

2 Basic Model

2.1 Networks

Let $N = \{1, 2, \ldots, n\}$ be a set of players, where $n \geq 3$. We are interested in modelling a situation where each of these players can choose the subset of other players with whom to play a fixed bilateral game. Formally, let

$g_i = (g_{i,1}, \dots g_{i,i-1}, g_{i,i+1}, \dots g_{i,n})$ be the set of links formed by player i. We suppose that $g_{i,j} \in \{1, 0\}$, and say that player i forms a link with player j if $g_{i,j} = 1$. The set of link options is denoted by $\mathcal{G}_i$. Any player profile of link decisions $g = (g_1, g_2 \dots g_n)$ defines a directed graph, $\{N, \Gamma\}$, called a *network.* Abusing notation, the network will also be denoted by g.

Specifically, the network g has the set of players N as its set of *vertices* and its set of arrows, $\Gamma \subset N \times N$, is defined as follows:

$$\Gamma = \{(i, j) \in N \times N : g_{ij} = 1\}.$$

Graphically, the link (i, j) may be represented as an edge between i and j, a filled circle lying on the edge near agent i indicating that this agent has formed (or supports) that link. Every link profile $g \in \mathcal{G}$ has a unique representation in this manner. Figure 1 below depicts an example. In it, player 1 has formed links with players 2 and 3, player 3 has formed a link with player 1, while player 2 has formed no link.[5]

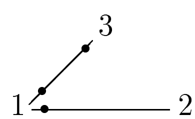

Figure 1

Given a network g, we say that a pair of players i and j are directly linked if at least one of them has established a linked with the other one, i.e. if $\max\{g_{i,j}, g_{j,i}\} = 1$. To describe the pattern of players' links, it is useful to define a modified version of g, denoted by $\bar{g}$, that is defined as follows: $\bar{g}_{i,j} = \max\{g_{i,j}, g_{j,i}\}$ for each i and j in N. Note that $\bar{g}_{i,j} = \bar{g}_{j,i}$ so that the index order is irrelevant. We say there is a *path* in g between i and j if either $\bar{g}_{i,j} = 1$ or there exist agents $j_1, \dots, j_m$ distinct from each other and i and j such that $\bar{g}_{i,j_1} = \dots = \bar{g}_{j_k,j_{k+1}} = \dots \bar{g}_{j_m,j} = 1$.

A subgraph $g' \subset g$ is called a *component* of g if for all $i, j \in g'$, $i \neq j$, there exists a path in g' connecting i and j, and for all $i \in g'$ and $j \in g$,

[5] Since agents choose strategies independently of each other, two agents may simultaneously initiate a two-way link, as seen in the figure.

$g_{i,j} = 1$ implies $g'_{ij} = 1$. A network with only one component is called connected. Given any g, the notation $g+ij$ will denote the network obtained by replacing $g_{i,j}$ in network g by 1. Similarly, $g - g_{i,j}$ will refer to the network obtained by replacing $g_{i,j}$ in network g by 0. A connected network g is said to be minimally connected if the network obtained by deleting any single link, $g - g_{i,j}$, is not connected. A special example of minimally connected network is the center-sponsored star: a network g is called a *center-sponsored star* if there exists some $i \in N$ such that, for all $j, k \in N \backslash \{i\}$, $j \neq k$, $g_{ij} = 1$ and $g_{jk} = 0$.

2.2 Social Game

Every pair of directly linked individuals plays a given 2×2 symmetric game in strategic form with common action (or strategy) set given by $A = \{\alpha, \beta\}$. For each pair of actions $a, a' \in A$, the payoff $\pi(a, a')$ earned by a player choosing a when the partner plays a' is given by the following table:

1 \ 2	α	β
α	d	e
β	f	b

Table I

The game is taken to be one of coordination, which amounts to postulating that

$$d > f, \quad b > e. \tag{1}$$

Without loss of generality, it is assumed that α is the efficient action, i.e.

$$d > b. \tag{2}$$

Furthermore, in order to focus on the most interesting scenario, it is supposed that efficiency and risk-dominance are conflicting criteria. That is:

$$d + e < b + f \tag{3}$$

Let $N^d(i; g) \equiv \{j \in N : g_{i,j} = 1\}$ be the set of players in network g with whom player i has established links, while $\nu^d(i; g) \equiv |N^d(i; g)|$ is its cardinality. Similarly, let $N^d(i; \bar{g}) \equiv \{j \in N : \bar{g}_{i,j} = 1\}$ be the set of players in network g with whom player i is connected (i.e. can play the coordination game), while $\nu^d(i; \bar{g}) \equiv |N^d(i; \bar{g})|$ is the cardinality of this set.[6]

An important feature of our approach is that links are assumed *costly*. Precisely, every agent who establishes a link with some other player is taken to incur a cost $c > 0$. Thus, we suppose that the cost of forming a link is independent of the number of links being established and is the same across all players. Another important feature of our model is that links are *one-sided*. This aspect of the model allows us to use standard solution concepts from non-cooperative game theory in addressing the issue of link formation. To simplify matters in this respect, we shall assume that the payoff of the bilateral game are all positive and, therefore, no player has any incentive to refuse links initiated by other players.

We are now in a position to define the payoff function for the complete game. Under the assumption that every player i is obliged to choose the same action in the (possibly) several bilateral games that she is engaged in, her strategy space can be identified with $S_i = \mathcal{G}_i \times A$, where recall that $\mathcal{G}_i$ is the set of possible link decisions by i and A is the common action space of the underlying bilateral game. Then, given the strategies of other players, $s_{-i} = (s_1, \ldots s_{i-1}, s_{i+1}, \ldots s_n)$, the payoff to a player i from playing some strategy $s_i = (g_i, a_i)$ is given by:

$$\Pi_i(s_i, s_{-i}) = \sum_{j \in N^d(i;\bar{g})} \pi(a_i, a_j) - \nu^d(i; g) \cdot c \tag{4}$$

[6] In section 4 below, we consider a model where two players can play a game if they are either directly or indirectly connected with each other.

This allows us to particularize the standard notion of Nash Equilibrium as follows. A strategy profile $s^* = (s_1^*, \ldots s_n^*)$ is said to be a *Nash equilibrium* for the game if, for all $i \in N$,

$$\Pi_i(s_i^*, s_{-i}^*) \geq \Pi_i(s_i, s_{-i}^*), \forall s_i \in S_i. \tag{5}$$

The set of Nash equilibria will be denoted by S^*. A Nash equilibrium is said to be strict if every player gets a strictly higher payoff with his current strategy than he would with any other strategy.

2.3 Dynamics

Time is modelled discretely, $t = 1, 2, 3, \ldots$. At each t, the state of the system is given by the strategy profile $s(t) \equiv [(g_i(t), a_i(t))]_{i=1}^n$ specifying the action played, and links established, by each player $i \in N$. At every period t, there is a positive independent probability $p \in (0, 1)$ that any given individual gets a chance to revise her strategy. If she receives this opportunity, we assume that she selects a new strategy

$$s_i(t) \in \arg\max_{s_i \in S_i} \Pi_i(s_i, s_{-i}(t-1)). \tag{6}$$

That is, she selects a best response to what other players chose in the preceding period. If there are several strategies that fulfill (6), then any one of them is taken to be selected with, say, equal probability. This strategy revision process defines a simple Markov chain on $S \equiv S_1 \times ... \times S_n$. In our setting, which will be seen to display multiple strict equilibria, there are several absorbing states of the Markov chain. This motivates the examination of the relative robustness of each of them.[7]

To do so, we employ the, by now, standard techniques used by Kandori, Mailath and Rob (1993), and Young (1993). We suppose that, occasionally, players make mistakes, experiment, or simply disregard payoff considerations in choosing their strategies. Specifically, it is postulated that, conditional on

[7]We note that the set of absorbing states of the Markov chain coincides with the set of strict Nash equilibria of the one-shot game.

receiving a revision opportunity, a player chooses her strategy at random with some small "mutation" probability $\epsilon > 0$. For any $\epsilon > 0$, the process defines a Markov chain that is aperiodic and irreducible and, therefore, has a unique invariant probability distribution. Let us denote this distribution by μ_ϵ. We analyze the structure of μ_ϵ as the probability of mistakes becomes very small, i.e. formally, as ϵ converges to zero.

Define $\lim_{\epsilon\to 0} \mu_\epsilon = \hat{\mu}$. We shall say that a state $s = (s_1, s_2, \ldots, s_n)$ is *stochastically stable* if $\hat{\mu}(s) > 0$. Heuristically, we may view a stochastically stable state as one which is played a significant fraction of the time in the long run, even when the perturbation probability ε becomes arbitrarily low.

3 Analysis

This section is divided into three parts. The first one characterizes the networks as well as action profiles that define a Nash equilibrium of the induced social game. The second part studies the unperturbed learning dynamics and establishes results concerning convergence of the process. In the third part, we examine the stochastic stability of different equilibrium configurations.

3.1 Equilibrium outcomes

Let $s^* = [(g_i^*, a_i^*)]_{i=1}^n$ be a Nash equilibrium of the population game described and denote by $g^* \equiv (g_i^*)_{i=1}^n$ the corresponding equilibrium network.[8]

Our first result concerns the nature of networks that arise in equilibria. When costs of link formation are low ($c < e$), then a player has an incentive to link up with other players irrespective of the actions the other players are choosing. On the other hand, when costs are quite high (specifically, $b < c < d$) then everyone who is linked must be choosing the efficient action. This, however, implies that it is attractive to form a link with every other player and we get the complete network again. Thus, for relatively low and high costs, we should expect to see the complete network.

[8]The fact that links are costly immediately implies the absence of superfluous links, i.e. if $g_{i,j}^* = 1$ then $g_{j,i}^* = 0$.

In contrast, when costs are at an intermediate level ($f < c < b$) a richer set of configurations is possible. On the one hand, since $c > f(> e)$, the link formation is only worthwhile if other players are choosing the same action. On the other hand, since $c < b(< d)$, coordinating at either of the two equilibria (in the underlying coordination game) is better than not playing the game at all. This allows for networks with two disconnected components in equilibria.

The above considerations are summarized and completed in the following result.

Proposition 3.1 *Suppose (1)-(3) and (4) hold. (a) If $c < \min\{f, b\}$, then an equilibrium network is complete. (b) If $f < c < b$, then an equilibrium network is either complete or can be partitioned into two complete components.*[9] *(c) If $b < c < d$, then an equilibrium network is either empty or complete. (d) If $c > d$, then the unique equilibrium network is empty.*

We now characterize the Nash equilibria of the static game. First, we introduce some convenient notation. On the one hand, recall that g^e denotes the empty network characterized by $g^e_{ij} = 0$ for all $i, j \in N$ ($i \neq j$). On the other hand, let

$$G^c \equiv \{g : \forall i, j \in N,\ \bar{g}_{ij} = 1,\ g_{ij} g_{ji} = 0\}$$

stand for the set of complete and essential networks on the full player population N. Analogously, for any given subset $M \subset N$, denote by $G^c(M)$ the set of complete and essential subgraphs on M. Given any state $s \in S$, we shall say that $s = (g, a) \in S^h$ for some $h \in \{\alpha, \beta\}$ if $g \in G^c$ and $a_i = h$ for all $i \in N$. More generally, we shall write $s = (g, a) \in S^{\alpha\beta}$ if there exists a partition of the population into two subgroups, N^α and N^β (one of them possibly empty), and corresponding components of g, g^a and g^β, such that:

(i) $g^a \in G^c(N^\alpha)$, $g^\beta \in G^c(N^\beta)$;

(ii) $\forall i \in N^\alpha,\ a_i = \alpha$; $\forall i \in N^\beta,\ a_i = \beta$.

With this notation in hand, we may state the following result.

[9] Our parameter conditions allow both $f < b$ and $b < f$. Of course, if the latter inequality holds, Part (b) of Proposition 3.1 (and also that of Proposition 3.2 below) applies trivially.

Proposition 3.2 *Suppose (1)-(3) and (4) hold. (a) If $c < \min\{f, b\}$, then the set of equilibrium states $S^* = S^\alpha \cup S^\beta$. (b) If $f < c < b$, then $S^\alpha \cup S^\beta \subset S^* \subset S^{\alpha\beta}$, the first inequality being strict for large enough n. (c) If $b < c < d$, then $S^* = S^\alpha \cup \{(g^e, (\beta, \beta, ..., \beta))\}$. (d) If $c > d$, then $S^* = \{g^e\} \times A^n$.*

The above result indicates that, whenever the cost of links is not excessively high (i.e. not above the maximum payoff attainable in the game) Nash equilibrium conditions allow for a genuine outcome multiplicity. For example, under the parameter configurations contemplated in Parts (a) and (c), such an equilibrium multiplicity permits alternative states where either of the two actions is homogeneously chosen by the whole population. Under the conditions specified in Part (b), the multiplicity concerns a wide range of possible states where neither action homogeneity nor full connectedness necessarily prevails. Therefore, the model raises a fundamental issue of equilibrium selection.[10] To address this issue, we approach matters dynamically in the subsequent subsections.

3.2 Long-run learning dynamics

All (strict) Nash equilibria are rest points in terms of the unperturbed dynamics proposed in section 2.3. Here we examine if the *set* of such equilibria are a dynamically sound prediction in the sense that, given any initial condition, the learning dynamics is bound to lead (a.s.) to *some* state in S^*. This issue is addressed by the following result.

Proposition 3.3 *Suppose (1)-(3) and (4) hold. Then the best response dynamics converges to one of the rest points indicated in Proposition 3.2.*

Since the (unperturbed) process of social learning converges to the set S^*, long-run states of the *perturbed* dynamics will also be in this set. (Heuristically, if the perturbation probability is small, "most" of the time the process will be in the set S^*.) This observation, however, does not by itself settle the

[10] We note that the equilibria identified in parts (a)-(b) and those in S^α are also strict equilibria.

issue of how the perturbed process addresses the equilibrium selection problem, i.e. the question of what specific states in S^* will prevail (again, most of the time) independently of initial conditions. This is the issue addressed in the next subsection.

3.3 Stochastic Stability

The states which are selected in the long run by the perturbed learning dynamics are those in the support of $\hat{\mu}$, the limit invariant distribution. Formally, the set of such *stochastically stable states* is defined as follows:

$$\hat{S} \equiv \{s \in S : \hat{\mu}(s) > 0\}.$$

The main result of this paper characterizes the set of stochastically stable states.

Theorem 3.1 *Suppose (1)-(3) and (4) hold. There exists some $\tilde{n}$ such that, if $n \geq \tilde{n}$, the set of stochastically stable sets is as follows. (a) If $c < \min\{f, b\}$, then $\hat{S} = S^{\beta}$. (b) If $f < c < b$, then $\hat{S} = S^{\alpha}$, provided $d - b \geq b - f$. (c) If $b < c < d$, then $\hat{S} = S^{\alpha}$. (d) If $c > d$ then $\hat{S} = \{g^e\} \times A^n$.*

This result establishes an essentially unique pattern of play for (the interior) of each of the four regions resulting from a comparison of the cost of link formation and the payoffs of the underlying game.[11] In particular, unless the costs are excessively high, we should expect to see the complete network. However, within this range of permissible cost levels, the precise level of the costs matters crucially. When they are low, everyone chooses the risk-dominant action, while if they are high everyone chooses the efficient action. Thus even though the eventual network is the same, the process of

[11] The maintained assumptions on the game payoffs are just as before, i.e. the underlying game is a coordination game where the efficient and risk-dominant actions differ. However, for Part (b) – which is void if $f > b$ – an additional condition on payoffs is required, which may be interpreted as imposing some lower bound on the "efficiency premium" enjoyed by action α at equilibrium. We conjecture that this condition may well be unnecessary, but have been incapable of dispensing with it in the proof of the result. In any case, note that it is fully compatible with (3), i.e. the assumption that efficiency and risk dominance conflict.

link formation has a powerful influence on the quality of social coordination obtained.

In our framework, the dynamic process involves link formation decisions as well as action choice in the underlying bilateral coordination game. Since the analysis of such co-evolution of links and actions requires some new arguments, it may be worthwhile at this point to elaborate on the intuition underlying the result.

First, we derive the minimal size of a group of α and β action choosers which can sustain itself, i.e., no agent has an incentive to switch actions. We denote this size by K_α and K_β, respectively. Since $d > b$ it follows that $K_\alpha < K_\beta$.

Second, we study the transition process from a state in which everyone chooses action β and the network is complete to a state in which there is a sustainable group of α choosers. Roughly speaking, we show that this transition only requires a single event of simultaneous mutations involving K_α agents.

Third, the argument derives a lower bound on the number of mutations needed to transit out of a state in which everyone chooses action α and the network is complete. Denoting this lower bound by H_β, it is shown that $H_\beta < K_\beta$.

Fourth, we study the transition from an intermediate stage in which there are two components conforming to different actions to a complete network in which everyone chooses the same action. We show that this transition can be accomplished via a sequence of single mutations interspersed with best-response play.

Finally, we note that the first, second, and fourth steps describe a transition tree, from a state $s \in S^\beta$ to a state $s \in S^\alpha$. Roughly speaking, this requires K_α simultaneous mutations. On the other hand, the third and fourth steps show that the transition from a state $s \in S^\alpha$ to a state $s \in S^\beta$ requires at least H_β simultaneous mutations. Since, under the maintained assumptions on the parameters, we find that $H_\beta > K_\alpha$, the arguments developed by Kandori, Mailath and Rob (1993) and Young (1993) can then be used to complete the proof of the theorem.

4 Extension: Indirect Links

In the basic model we assumed that two players can undertake a transaction (play the game) if and only if they have a direct link. In many settings, individual players can also transact if they are indirectly linked, for example through common acquaintances. There are a variety of ways in which the benefits of indirect links can me modeled. To get a first idea of the impact of such indirect links, we consider a model in which indirect links are as good as direct links and focus on its Nash equilibria. Our analysis shows that those networks which yield a *strict* Nash equilibrium have a simple architecture, while the nature of social coordination is consistent with the pattern observed in our basic model.

We say that two players are linked with each other if there is a path between them in the network. Thus, given a network of interaction g, two players can play a game if either $\overline{g}_{i,j} = 1$ or there exists a sequence of distinct players $i_1, i_2, i_3, \ldots, i_m$, such that all $\overline{g}_{i,i_1} = \overline{g}_{i_1,i_2} = \cdots ... = \overline{g}_{i_m,j} = 1$. We write $i \overset{\bar{g}}{\longleftrightarrow} j$ to indicate a path between players i and j in network $\bar{g}$. Let $\hat{N}(i;g) \equiv \{j \in N : i \overset{\bar{g}}{\longleftrightarrow} j\}$ be the players with whom player i is directly or indirectly linked in a network g. Also recall that $N^d(i;g) \equiv \{j \in N : g_{i,j} = 1\}$ is the set of players in network g with whom player i has formed links, while $\nu^d(i;g) \equiv |N^d(i;g)|$ is the cardinality of this set.

Given the strategies of other players, $s_{-i} = (s_1, \ldots s_{i-1}, s_{i+1}, \ldots s_n)$, the payoff to a player i from playing some strategy $s_i = (g_i, a_i)$ is now given by:

$$\hat{\Pi}_i(s_i, s_{-i}) = \sum_{j \in \hat{N}(i;\bar{g})} \pi(a_i, a_j) - \nu^d(i;g) \cdot c \tag{7}$$

Given these payoffs, the definition of Nash equilibrium is as before. Our first result derives some basic properties of equilibrium networks and actions.

Proposition 4.1 *Suppose (1)-(3) and (7) hold. There exists some $\tilde{n}$ such that, if $n \geq \tilde{n}$, an equilibrium network is either minimally connected or empty. If the equilibrium network is connected then everyone chooses the same action and social conformism obtains.*

The proofs of results in this section are provided in Appendix B. We note that an implication of allowing for indirectly linked players to play with each other: any non-empty equilibrium network is connected and in such a network everyone chooses the same action. Recall that in the basic model, where only directly connected players are allowed to play with each other, unconnected networks can be supported in a strict Nash equilibrium and social diversity can also be sustained.

Minimal connectedness, however, is a permissive requirement and allows for a a wide range of network architectures. Our main interest is in the robustness of the results of the basic model. With this aim in view, we now approach matters more simply than before and adopt a purely static perspective. Specifically, we refine the notion of equilibrium and examine the nature of strict Nash networks. It turns out that, in our setting, the requirement of strict best response is quite effective. The following proposition provides a characterization of strict Nash architectures. Recall that a center-sponsored star is a network in which a single agent forms a link with every other player and therefore pays for every link formed.

Proposition 4.2 *Suppose (1)-(3) and (7) hold. There exists some $\tilde{n}$ such that, if $n \geq \tilde{n}$, then a strict Nash network is either empty or a center-sponsored star.*

We now provide the intuition behind this result. Consider a minimally connected network as identified in Proposition 4.1. Suppose that player i forms a link with a player j in this network. Then, the requirement of strict Nash has the implication that player j cannot be directly linked to any other player. For, if player j were linked to some other player k, then player i could interchange his link with player j for a link with player k and get the same payoffs. This would mean that the link of player i with player j is not a strict best response. Since the network in question must be connected this argument also implies that player i is directly linked to every other player and that he forms all the links. In other words, the only candidate for strict Nash networks is the center-sponsored star.[12]

[12] This argument is similar to those used in Bala and Goyal (1999, Proposition 4.2).

The above result on strict Nash networks helps us in providing a characterization of strict Nash equilibria for different values of cost of forming links. This is contained in the following result, which parallels, for the present context, the insights obtained in Proposition 3.2. The set of all center-sponsored stars will be denoted by G^{cs}. Also let S^{**} be the set of strict Nash equilibrium.

Proposition 4.3 *Suppose (1)-(3) and (7) hold. There exists some $\tilde{n}$ such that, if $n \geq \tilde{n}$, the following is true: (a) If $0 < c < b$ then $S^{**} = [G^{cs} \times \{(\alpha, .., \alpha)\}] \cup [G^{cs} \times \{(\beta, .., \beta)\}]$. (b) If $b < c < d$ then $S^{**} = G^{cs} \times \{(\alpha, .., \alpha)\}$. (c) If $c > d$ then there is no strict Nash equilibrium.*

5 Concluding Remarks

In many economic and social contexts, individuals can undertake a transaction only if they are 'linked' or related to each other. It is therefore natural to postulate that individual players invest effort and resources in forming links with others, their link decisions then defining the network of social interaction. In this paper, we have studied the nature of networks that form and the effects of link formation on social coordination.

We start with a basic model in which two players can transact or play a game only if they have a direct link between them. We then consider a variation of it where two players can play a game if they are directly or indirectly linked. Our analysis reveals that, in both settings, equilibrium networks have simple architectures. In the former model the unique equilibrium architecture is the complete network, while in the latter model, the unique equilibrium architecture is the star. Moreover, in both settings the cost of forming links has a crucial effect on the nature of social coordination. For low costs of forming links, individuals coordinate on the risk-dominant action, while for high costs of forming links individuals coordinate on the efficient action. These results suggest that the process of network formation may have serious implications for the nature of social coordination.

An important aspect of our model is that link formation is *one-sided.* This formulation has the advantage that it allows us to study the social process of link formation and coordination as a non-cooperative game. In some settings, it is perhaps more natural to think of link formation as a *two-sided* process: a link is formed when the two players involved both incur some costs. This implies that both players must acquiesce in the formation of the link. In independent work, Jackson and Watts (1999) study such a model.

In their work, Jackson and Watts assume that two players can play a game only if they are directly linked. This assumption corresponds to our basic model. In this setting they find, like us, that the equilibrium network is complete. However their results on social coordination are quite different. For instance, they find that, if the costs of link formation are high, then both of the states where players choose a common action are stochastically stable. In contrast, we find that, when the costs are high (but below the maximum achievable payoff), the only stable states involve players choosing the efficient action. This difference arises out of differences in the way we model the link formation as well as the accompanying assumptions on the timing of moves. They postulate that individuals choose links and actions separately, i.e., players choose links taking actions as given while they choose actions taking the links as a given. In our view, this feature of their model is intimately linked to the two-sided link formation formulation. Individuals contemplating the formation of a link have to form expectations regarding the actions that will follow. Jackson and Wolinsky assume that past actions should be taken as fixed, for the time being, when decisions on links are being considered. By contrast, in our setting, any individual can independently change *both* her action and her supported links.

These observations suggest that it is important to study the effects of different levels of flexibility in the two choice dimensions, links and actions– for example, it might be natural to postulate that link revision is more rigid than action change.

Finally, we note that our analysis has examined the two polar cases with regard to the value of indirect links: either only directly linked players can interact (and indirect links are irrelevant), or indirect links are as good as direct links, irrespective of the length of the path between the players. We would like to examine a more general formulation which allows for 'distant' links to be less valuable as compared to 'near by' links.

6 Appendix A

Proof of Proposition 3.1: We take up part (a) first. Consider an incomplete network g. There are then two players, say i and j, such that $\overline{g}_{ij} = 0$. If $a_i = a_j$ then the minimum payoff to player i from a link with player j is b while the cost is c. Since $c < b$, player i has an incentive to form a link with player j, which implies that g cannot be an equilibrium network. Next, suppose that $a_i \neq a_j$ and let player i (without loss of generality) be the player who chooses action β. If she forms a link with player j then she will get a payoff of f while the cost is c. Since $c < f$, this is clearly profitable. Hence an incomplete network cannot be sustained in equilibrium.

We now prove part (b). Consider an equilibrium network g and suppose that g' is a component of g. Suppose i and j belong to g' and $\overline{g}_{ij} = 1$. Then, it follows that $a_i = a_j$ because $c > f > e$ and, if i and j are choosing different actions, then the player who forms the link can do strictly better by deleting her link. A similar argument can be used to show that if i and j belong to g' but are indirectly linked then they must be choosing the same action in the underlying coordination game. These arguments establish that every player in the same component of an equilibrium network must choose the same action.

Consider next an equilibrium network with two components g' and g''. We claim that $a_i \neq a_j$ if $i \in g'$ and $j \in g''$. The reason for this is that if $a_i = a_j$ then the minimum payoff to i from playing the coordination game with j is b. Since $c < b$ this means that it is profitable for player i to form the link $g_{ij} = 1$, which implies that g is not an equilibrium network. The final step is to note that since there are only two actions in the coordination game, there can be at most two distinct components.

We finally show that each of the components in an equilibrium network must be complete. From the above considerations, we know that every component must involve homogeneous actions. This immediately implies that the payoff to any given player from playing the coordination game with another player (in the same component) is at last b. Since $c < b$, it is clearly profitable for every player to form a link with any other player in the same component. Thus a component in an equilibrium network must be complete.

We now prove part (c). There are two subcases to consider: $c > \max\{b, f\}$ or $f > c > b$. (Note, of course, that the former subcase is the only one possible if $b > f$.) Suppose first that $c > \max\{b, f\}$, and let g be an equilibrium network which is non-empty but also incomplete. Note that if $\bar{g}_{i,j} = 1$, then $a_i = a_j = \alpha$. Moreover, if $a_j = \beta$, then player j can have no links in the network. (These observations follow directly from the hypothesis that $c > \max\{b, f\}$.) However, since g is assumed incomplete, there must exist a pair of agents i and j such that $\overline{g}_{ij} = 0$. There are three cases possible with respect to the actions of players i and j. First, suppose that $a_i = a_j = \alpha$. Then, since $c < d$, it is clearly profitable for either of the two players to deviate and form a link with the other player. Suppose next that $a_i = a_j = \beta$. Then, players i and j can have no links and, furthermore, since g is non-empty, there must be at least two other players $k, l \in N$ such that $a_k = a_l = \alpha$. The payoff to players i and j in the network g is simply zero. In contrast, if, say, player i deviates towards choosing action α and forming a link with k (or l), she will increase her payoff strictly. Finally, consider the case where $a_i \neq a_j$ and let player i choose β. Then, if this player deviates to action α and forms a link with player j she increases her payoff strictly. We have thus shown that $\overline{g}_{ij} = 0$ cannot be part of an equilibrium network. This proves that a non-empty but incomplete network cannot be an equilibrium network in the first subcase considered.

Consider now the second case with $b < c < f$ and suppose, for the sake of contradiction, that g is an equilibrium network that is non-empty but incomplete. Then there is a pair of players i and j such that $\bar{g}_{i,j} = 0$. If $a_i = a_j = \alpha$, then it is immediate that each of the players has an incentive to form a link with the other. Thus g cannot be part of a Nash equilibrium. The other possibility is that one of these two players, say i, chooses β. We note that since $b < c < f$, the non-empty network g is not consistent with everyone choosing action β. Thus, player j must be choosing action α. Moreover, since $c < f$, it follows that player i must find it optimal to connect with every other player choosing α. Finally, note that no player $k \neq i$ will form a link with player i since $e < b < c < f$. These observations immediately imply that player i must, in equilibrium, support a link with another player if and only if this other player chooses action α. It follows, therefore (since $f < d$), that this player can do strictly better by deviating to action α and supporting a link with every player choosing α (in particular, player j), which leads to the desired contradiction for the second subcase of part (c).

Part (d) is immediate from the hypothesis that $c > d$.

□

Proof of Proposition 3.2: We start proving Part (a). In view of Part (a) of Proposition 3.1 and the fact that the underlying game is of a coordination type, the inclusion $S^{\alpha} \cup S^{\beta} \subset S^*$ is obvious. To show the converse inclusion, take any state s such that the sets $A(s) \equiv \{i \in N : s_i = \alpha\}$ and $B(s) \equiv \{i \in N : s_i = \beta\}$ are both non-empty. We claim that it cannot be an equilibrium state.

Two subcases need to be considered: $c < e$ and $e < c < \min\{f, b\}$. We start by proving the first subcase. Assume, for the sake of contradiction, that such a state s is a Nash equilibrium of the game and denote $m \equiv |A(s)|$, $0 < m < n$. Then, the following optimality conditions must apply. On the one hand, for players $i \in A(s)$, we must have:

$$(m-1)d + (n-m)e - \nu^d(i;g) \cdot c \geq (m-1)f + (n-m)b - \nu^d(i;g) \cdot c \quad (8)$$

and for players $j \in B(s)$:

$$(n-m-1)b + mf - \nu^d(j;g) \cdot c \geq (n-m-1)e + md - \nu^d(j;g) \cdot c. \quad (9)$$

Note that the above two conditions rely on the fact that, since $c < e$, every deviating player would want to keep supporting her previous links. It is immediate to see that these conditions are jointly incompatible.

Assume now that $c > e$. Then, the counterpart of the previous optimality condition for all players $i \in A(s)$ is as follows: $\forall r \leq \nu^d(i;g)$,

$$\begin{aligned} &(m-1)d + (n-m)e - \nu^d(i;g) \cdot c \\ &\geq (m-1-r)f + (n-m)b - (\nu^d(i;g) - r) \cdot c \end{aligned} \tag{10}$$

where we implicitly use the fact that, at equilibrium, all players in $A(s)$ will only support links to other players in $A(s)$ and, furthermore, that all players in $B(s)$ will support a complete set of links to players in $A(s)$. Clearly, the maximum for the RHS of (10) is attained for $r = 0$ (since $c < f$), which implies that the present optimality condition is identical to (8).

Consider now the players choosing action β. For each $j \in B(s)$, the relevant optimality condition in the present case can be written as follows:

$$(n-m-1)b + mf - \nu^d(j;g)c \geq (n-m-1-(\nu^d(j;g)-m))e + md - mc \tag{11}$$

It may be checked that (10) implies:

$$\frac{m}{n} \geq \frac{b-e}{b+d-e-f} + \frac{1}{n}\frac{d-f}{b+d-e-f} \tag{12}$$

and (11) implies:

$$\frac{m}{n} \leq \frac{b-e}{b+d-e-f} - \frac{1}{n}\frac{b-e}{b+d-e-f}. \tag{13}$$

Again, (12) and (13) are incompatible, thus proving the desired conclusion.

Now, we turn to Part (b). As in Part (a), the inclusion $S^\alpha \cup S^\beta \subset S^*$ is trivial, in view of Part (a) of Proposition 3.1. To show that the inclusion $S^* \subset S^{\alpha\beta}$ holds strictly for large enough n, consider a state s where both $A(s)$ and $B(s)$, defined as above, are both non-empty and fully connected components. Specifically, focus attention on those configurations that are symmetric within each component, so that every player in $A(s)$ supports $\frac{m-1}{2}$ links and every player in $B(s)$ supports $\frac{m-n-1}{2}$ links. (As before, m stands for the cardinality of $A(s)$ and we implicitly assume, for simplicity, that m and $n-m$ are odd numbers.) For this configuration to be a Nash equilibrium, we must have that the players in $A(s)$ satisfy:

$$d(m-1) - \frac{m-1}{2}c \geq f\frac{m-1}{2} + b(n-m) - c(n-m) \tag{14}$$

where we use the fact that, in switching to action β, any player formerly in $A(s)$ will have to support herself all links to players in $B(s)$ and will no longer support any links to other players in $A(s)$ – of course, she still anticipate playing with those players from $A(s)$ who support links with her.

On the other hand, the counterpart condition for players in $B(s)$ is:

$$(n-m-1)b - \frac{n-m-1}{2}c \geq dm + e\frac{n-m-1}{2} - cm \tag{15}$$

where, in this case, we rely on considerations for players in $B(s)$ that are analogous to those explained before for players in $A(s)$. Straightforward algebraic manipulations show that (14) is equivalent to:

$$\frac{m}{n} \geq \frac{1}{n}\frac{2d-c-f}{2b+2d-3c-f} + \frac{2(b-c)}{2b+2d-3c-f} \tag{16}$$

and (15) is equivalent to:

$$\frac{m}{n} \leq \frac{1}{n}\frac{c+e-2b}{2b+2d-3c-e} + \frac{2b-c-e}{2b+2d-3c-e}. \tag{17}$$

We now check that, under the present parameter conditions:

$$\frac{2b-c-e}{2b+2d-3c-e} > \frac{2(b-c)}{2b+2d-3c-f}.$$

Denote $U \equiv 2b-c$, $V \equiv 2b+2d-3c$, and rewrite the above inequality as follows:

$$\frac{U-e}{U-c} > \frac{V-e}{V-f} \tag{18}$$

which is weaker than:

$$\frac{U-e}{U-f} > \frac{V-e}{V-f} \tag{19}$$

since $c > f$. The function $\zeta(x) \equiv \frac{x-e}{x-f}$ is uniformly decreasing in x since $f > e$. Therefore, since $U < V$, (19) obtains, which implies (18). Hence it follows that, if n is large enough, one can find suitable values of m such that (16) and (17) jointly apply. This completes the proof of Part (b).

We now present the proof for part (c). We know from Proposition 3.1 that the complete and the empty network are the only two possible equilibrium networks. Since $c > b > f > e$, it is immediate that, in the complete network, every player must choose α and this leads to a Nash equilibrium. Concerning the empty network, for it to be consistent with equilibrium, it must be the case that no player has an incentive to form a link with any other player. This implies that no player must choose α. Hence the only action profile consistent with equilibrium in the empty network is the one where everyone chooses the action β. Trivially, this outcome defines a Nash equilibrium.

The proof of part (d) follows directly from the hypothesis $c > \max\{d, b, f, e\}$.

$\square$

Proof of Proposition 3.3: It is enough to show that, from any given state ω_0, there is a finite chain of positive-probability events (bounded above zero, since the number of states is finite) that lead to a rest point of the best response dynamics.

Choose one of the two strategies, say β, and denote by $B(0)$ the set of individuals adopting action β at ω_0. Order these individuals in some pre-specified manner and starting with the first one suppose that they are given in turn the option to revise their choices (both concerning strategy and links). If at any given stage τ, the player i in question does not want to change strategies, we set $B(\tau + 1) = B(\tau)$ and proceed to the next player if some are still left. If none is left, the first phase of the procedure stops. On the other hand, if the player i considered at stage τ switches from β to α, then we make $B(\tau + 1) = B(\tau)\backslash\{i\}$ and, at stage $\tau + 1$, re-start the process with the first-ranked individual in $B(\tau + 1)$, i.e. *not* with the player following i. Clearly, this first phase of the procedure must eventually stop at some finite τ_1.

Then, consider the players choosing strategy α at τ_1 and denote this set by $A(\tau_1) \equiv N\backslash B(\tau_1)$. Proceed as above with a chain of unilateral revision opportunities given to players adopting α in some pre-specified sequence, restarting the process when anyone switches from α to β. Again, the second phase of the procedure ends at some finite τ_2.

By construction, in this second phase, all strategy changes involve an increase in the number of players adopting β, i.e. $B(\tau_2) \supseteq B(\tau_1)$. Thus, if the network links affecting players in $B(\tau_1)$ remain unchanged throughout, it is clear that no player in this set would like to switch to α if given the opportunity at $\tau_2 + 1$. However, in general, their network links will also evolve in this second phase, because individual players in $A(\tau_1)$ may form or delete links with players in $B(\tau_1)$. In principle, this could alter the situation of individual members of $B(\tau_1)$ and provide them with incentives to switch from β to α. It can be shown, however, that this is not the case. To show it formally, consider any given typical individual in $B(\tau_1)$ and denote by $\hat{r}^h$, $h = \alpha, \beta$, the number of links received (but not supported) by this player from players choosing action h. On the other hand, denote $\hat{m} \equiv |A(\tau_1)|$. Then, since the first phase of the procedure stops at τ_1, one must have:

$$\begin{aligned} &\max_{q^\alpha, q^\beta} b(q^\beta + \hat{r}^\beta) + f(q^\alpha + \hat{r}^\alpha) - c(q^\alpha + q^\beta) \\ \geq &\max_{q^\alpha, q^\beta} e(q^\beta + \hat{r}^\beta) + d(q^\alpha + \hat{r}^\alpha) - c(q^\alpha + q^\beta) \end{aligned} \tag{20}$$

for all q^α, q^β such that $0 \leq q^\alpha \leq \hat{m} - \hat{r}^\alpha, \quad 0 \leq q^\beta \leq n - \hat{m} - 1 - \hat{r}^\beta$. Now denote by $\tilde{r}_h$ and $\tilde{m}$ the counterpart of the previous magnitudes ($\hat{r}^h$ and $\hat{m}$) prevailing at τ_2. By construction, we have $\tilde{m} \leq \hat{m}$, $\tilde{r}^\alpha \leq \hat{r}^\alpha$, and $\tilde{r}^\beta \geq \hat{r}^\beta$. We note that $\tilde{m} \leq \hat{m}$ by construction of the process. Next note that if $\tilde{r}^\alpha > \hat{r}^\alpha$ then this implies that some player who chooses action α has formed an additional link with player i in the interval between τ_1 and τ_2. This is only possible if $c < e$. It also implies that player i did not have a link with this player at τ_1. This is only possible if $c > f$, a contradiction. Thus $\tilde{r}^\alpha \leq \hat{r}^\alpha$. Finally note that $\tilde{r}^\beta \geq \hat{r}^\beta$ follows from the fact that the all the players choosing β at τ_1 do not revise their decisions in the interval between τ_1 and τ_2.

Therefore, (20) implies:

$$\begin{aligned} &\max_{q^\alpha, q^\beta} b(q^\beta + \tilde{r}^\beta) + f(q^\alpha + \tilde{r}^\alpha) - c(q^\alpha + q^\beta) \\ \geq &\max_{q^\alpha, q^\beta} e(q^\beta + \tilde{r}^\beta) + d(q^\alpha + \tilde{r}^\alpha) - c(q^\alpha + q^\beta) \end{aligned}$$

for all q^α, q^β such that $0 \leq q^\alpha \leq \tilde{m} - \tilde{r}^\alpha, \quad 0 \leq q^\beta \leq n - \tilde{m} - 1 - \tilde{r}^\beta$. This allows us to conclude that the concatenation of the two phases will lead the process to a rest point of the best response dynamics, as desired.

□

Proof of Theorem 3.1: In view of Propositions 3.2 and 3.3, the proof of Parts (a), (c) and (d) are straightforward applications of the standard arguments used in received evolutionary theory. Therefore, we shall simply sketch very briefly the main ideas. For Part (a), since $c < f$, the only rest points of the best-response dynamics involve states with a single complete component and a given action being played uniformly. Thus, from any given such state in S^h $(h \in \{\alpha, \beta\})$ the possibility of implementing a transition to some state in $S^{h'}$ $(h' \in \{\alpha, \beta\},\ h' \neq h)$ depends on enough simultaneous mutations taking place that shifts the system to the basin of attraction of the set $S^{h'}$. In the end, whether S^α or S^β is the set that includes the stochastically stable states depends on which of the two can be accessed from the other one through the minimum number of simultaneous mutations. This is determined by risk-dominance considerations, which are postulated to favor action β by (3). Therefore, the set of stochastically stable states must lie in the set S^β. In fact, since all states in this set may be connected by steps involving a single mutation and the ensuing operation of the best-response dynamics, the set S^β defines what Samuelson (1994) has labelled a mutation-connected component. A direct implication of this fact is that *all* states in S^β must be stochastically stable.

Concerning Part (c), refer again to Propositions 3.2 and 3.3, which indicate that the only candidates to being stochastically stable are the states in S^α and the state $(g^e, (\beta, \beta, ..., \beta))$. But, clearly, the latter state is extremely fragile to mutation since only one individual mutating to action α may trigger a chain of best responses by others which leads the system into some state in S^α. Since, reciprocally, the transition from any state in S^a to the state $(g^e, (\beta, \beta, ..., \beta))$. requires many simultaneous mutations, only the states in S^α (in fact, all of them since they define a mutation-connected component) are stochastically stable. And next, for Part (d), the fact that all absorbing states of the unperturbed best-response dynamics define the mutation-connected component $\{g^e\} \times A^n$ yields the stated conclusion.

Now, we present the detailed argument for Part (b). The proof for this case builds upon the following four lemmas, all of them stated under the general assumptions of the theorem and the specific parameter configurations contemplated in its Part (b).

Lemma 1 *Let $s \in S^*$ be an equilibrium state that displays two (complete) components, g^α and g^β, that partition the population into two non-empty subsets, $A(s)$ and $B(s)$, with $a_i = \alpha\ \forall i \in A(s)$ and $a_i = \beta\ \forall i \in B(s)$. Then, if their respective cardinalities satisfy $0 <\mid A(s) \mid, \mid B(s) \mid< n$ one must have:*

$$\frac{\mid A(s) \mid -1}{\mid B(s) \mid} \quad \geq \quad K_\alpha \equiv \frac{2b-2c}{2d-e-c} \tag{21}$$

$$\frac{\mid B(s) \mid -1}{\mid A(s) \mid} \quad \geq \quad K_\beta \equiv \frac{2d-2c}{2b-f-c}, \tag{22}$$

provided that $\mid A(s) \mid$ and $\mid B(s) \mid$ are odd. If either of them is even, the respective conditions are as follows:

$$\frac{\mid A(s) \mid -1-(c-e)}{\mid B(s) \mid} \quad \geq \quad K_\alpha \tag{23}$$

$$\frac{\mid B(s) \mid -1-(c-f)}{\mid A(s) \mid} \quad \geq \quad K_\beta. \tag{24}$$

One of these lower bounds is exactly attained at some state $s \in S^*$.

Proof: Let g^α and g^β be two graph components, as described in the statement of the Lemma. Focus on g^α and denote by $m \equiv\mid A(s) \mid$ the cardinality of g^α. For concreteness, assume that m and $n-m$ are odd, so that (21) applies – if one of them were even, the argument is adapted in a straightforward manner to obtain (23). Denote by q_i denote the number of links supported by player $i \in A(s)$ and let $\hat{q}^\alpha \equiv \max_{i \in A(s)} q_i$ stand for the maximum number of links supported by some individual in $A(s)$. Since g^α is complete, note for future reference that the minimum possible $\hat{q}$ is attained when every agent in $A(s)$ is connected to exactly the same number $\frac{m-1}{2}$ of other agents.

If the state under consideration is to be a rest point of the best-response dynamics, it must be the case that every individual in $A(s)$ displays no incentives to switch actions and link herself to the other β-component. This requires that, for every $i \in A(s)$, we have:

$$d(m - q_i - 1) + (d - c)q_i \geq e(m - q_i - 1) + (b - c)(n - m). \qquad (25)$$

By rearranging terms, we can write this expression as follows:

$$d(m - q_i - 1) + (d - c)q_i - e(m - q_i - 1) - (b - c)(n - m) \geq 0. \qquad (26)$$

Since $c > f > e$, this expression is decreasing in q_i. Thus the most favorable circumstances under which this expression may be satisfied for all $i \in A(s)$ is when each q_i attains the minimum value for $\hat{q}^\alpha$, namely $\frac{m-1}{2}$. Introducing this value in (25), we obtain:

$$(d - e)\frac{m-1}{2} + (d - c)\frac{m-1}{2} \geq (b - c)(n - m)$$

or

$$\frac{m-1}{n-m} \geq \frac{2b - 2c}{2d - e - c}.$$

Proceeding analogously for the β-component, we obtain the similar condition:

$$\frac{m' - 1}{n - m'} \geq \frac{2d - 2c}{2b - f - c}$$

where m' now stands for the cardinality of the set $B(s)$. This completes the proof of the Lemma. $\|$

Lemma 2 *Consider any equilibrium state $s \in S^*$ such as that considered in Lemma 1. The following statements hold.*

(a) State s can be reached from some state $s' \in S^\alpha (\subset S^*)$ by a suitable chain of single mutations together with one simultaneous mutation of $\mid B(s) \mid$ agents, the best-response dynamics operating in between each of these mutation events.

(b) *State s can be reached from some state $s'' \in S^{\beta}(\subset S^*)$ by a suitable chain of single mutations together with one simultaneous mutation of $\mid A(s) \mid$ agents, the best-response dynamics operating in between each of these mutation events.*

Proof: Consider Part (a) and suppose, for simplicity, that the cardinality $\mid B(s) \mid$ of component g^{β} in s is odd. Choose some state s' as described where the links among the agents in $N \backslash B(s)$ are exactly as those that these agents display in s. First, through a sequence of single mutations and ensuing best-response adjustment, it is clear that the process may make a transition from any s' as described to an (equilibrium) state $s_1 \in S^{\alpha}$ where every $i \in B(s)$ supports *exactly* $\frac{|B(s)|-1}{2}$ links to other agents in $B(s)$. For example, to achieve this final outcome, it is enough that, in sequence, each of the agents i which has more than $\frac{|B(s)|-1}{2}$ links at that point experiences a unilateral mutation who removes one of her links to some agent j with a lower number of supported links at that stage. If every such mutation is followed by a revision opportunity received by agent j, the latter will support a link to i. Eventually, the system reaches a state with the property desired.

Once state s_1 has been reached, a similar concatenation of single mutations may lead the process to a state where all players still play action α but with no agent in $B(s)$ supporting a link to agents in $N \backslash B(s)$. Call this state $s_2 \in S^{\alpha}$.

Finally, assume that a *simultaneous* mutation of *all* agents in $B(s)$ makes them switch from action α to action β, no other features of the situation being affected (in particular every pre-existing link remaining in place). If now the agents in $N \backslash B(s)$, and only them, receive a revision opportunity, they will quit supporting their links to agent in $B(s)$. This follows from the assumption that $c > f$. This leads to state s, as desired, thus proving Part (a). Part (b) is proven analogously. $\|$

Lemma 3 *Let* $s = (g, a)$ *be any arbitrary state (not necessarily an equilibrium) such that the subgraph defined on the set* $A(s) \equiv \{i \in N : a_i = \alpha\}$ *(not necessarily a proper component of* g*) is complete. Then, if*

$$\frac{n - |A(s)|}{|A(s)| - 1} < H_\beta \equiv \frac{d - c}{b - e} \tag{27}$$

every state $s' = (g', a')$ *which is reachable from* s *with positive probability through the exclusive operation of the best-response dynamics has* $|A(s')| = \{i \in N : a'_i = \alpha\} \geq |A(s)|$.

Proof: It is enough to show that given any state $s = (g, a)$ which induces a complete subgraph on $A(s)$ and satisfies (27), no player $i \in A(s)$ would like to switch actions, independently of what is the pattern of connections among the players in $A(s)$ and those in $N \backslash A(s)$. Take any player $i \in A(s)$ and denote:

- q_i^h, as the set of links supported by i to other individuals in the set $\{j \in N : a_j = h\}$, for each $h = \alpha, \beta$;
- r_i^h, as the set of links to i supported by individuals in the set $\{j \in N : a_j = h\}$, for each $h = \alpha, \beta$.

Of course, note that by the hypothesis we make on s, $q_i^\alpha + r_i^\alpha = |A(s)| \equiv m$. On the other hand, q_i^β and r_i^β are a priori unrelated, except that $q_i^\beta + r_i^\beta \leq |N \backslash A(s)| = n - m$. Suppose now that i receives a revision opportunity. With the former notation in hand, the (optimal) expected payoff of i to staying with α is given by

$$\pi_\alpha(q_i^\alpha, r_i^\beta) \equiv d(m-1) - cq_i^\alpha + er_i^\beta, \tag{28}$$

where we are implicitly relying on the fact that, if player i stays with strategy α, in the new state $\tilde{s}$ induced by i's adjustment, she will support no links to players in $N \backslash A(s)$ That is, adapting previous notation in the obvious way, $\tilde{q}_i^\beta = 0$. On the other hand, if i were to decide switching to strategy β, her expected payoffs are bounded above by the following expression:

$$\pi_\beta(q_i^\alpha, r_i^\beta) \equiv e(m - 1 - q_i^\alpha) + b(n-m) - c(n - m - r_i^\beta). \tag{29}$$

where, analogously as before, we implicitly use the fact that if player i switches to β, the new state $\tilde{s}$ will display $\tilde{q}_i^{\alpha} = 0$. Denote $\psi(q_i^{\alpha}, r_i^{\beta}) \equiv \pi_{\alpha}(q_i^{\alpha}, r_i^{\beta}) - \pi_{\beta}(q_i^{\alpha}, r_i^{\beta})$. Mere inspection of (28) and (29) shows that

$$\min_{q_i^{\alpha}, r_i^{\beta}} \psi(q_i^{\alpha}, r_i^{\beta}) = \psi(m-1, n-m) = (d-c)(m-1) - (b-e)(n-m).$$

By construction, if $\psi(m-1, n-m) > 0$, no agent $i \in A(s)$ would want to switch actions at s, no matter what its specific details might be (in particular, the pattern of links). Thus, if

$$\frac{n - |A(s)|}{|A(s)| - 1} < H_{\beta} \equiv \frac{d-c}{b-e},$$

the best-response dynamics is sure not to decrease the number of players choosing α. This completes the proof of the Lemma. $\|$

Lemma 4 *Consider any equilibrium state s involving two non-degenerate (α and β) components with respective cardinalities $\mid A(s) \mid > 0$ and $\mid B(s) \mid > 0$. Then, there is another equilibrium state s' with cardinality for the resulting α component $\mid A(s') \mid > \mid A(s) \mid$ that can be reached from s by a suitable chain of single mutations, the best-response dynamics operating in between each of them. An identical conclusion applies to some equilibrium state s'' with $\mid A(s'') \mid < \mid A(s) \mid$.*

Proof: Let s be any state with non-degenerate α- and β-components, whose respective player sets $A(s)$ and $B(s)$ have cardinalities $m > 0$ and $n - m > 0$. To address the first part of the Lemma, consider any given agent $i \in B(s)$ and let q_i be the number of links supported by her. We need to contemplate different possibilities:

1. Suppose $q_i < n - m - 1$. Then, there must be some other agent $j \in B(s)$ with $g_{ji} = 1$ (i.e. supporting a link to i). Take any such j and suppose that this individual experiences a mutation, whose only effect is to remove her link to i. Assume then that only player i receives a revision opportunity. Two subcases must be considered:

(a) She might decide to keep playing β and establish a link to j, in which case her number of supported links becomes $q_i + 1$ and a new rest point of the best response dynamics is attained.

(b) She might decide to switch to α, in which case she will also establish direct links to all agents in $A(s)$, thus eventually inducing a.s. a new rest point of the best response dynamics with the desired properties (i.e. a larger number of players in the α-component) if the players in $B(s)\backslash\{i\}$ alone then receive a (say, simultaneous) revision opportunity.

2. Now suppose that $q_i = n - m - 1$, which implies that agent i supports links to all other agents in $B(s)$. Assume that she experiences a mutation, whose effect is to remove all her supported links and switch to action α. Furthermore, suppose that all agents in $A(s)$ (and only them) are given a revision opportunity, thus making all of them establish links to agent i. Again, the final outcome for the ensuing operation of the best response dynamics is a.s. a new rest point of the best response dynamics with the desired properties.

Combining the above considerations, it follows that, after a finite chain of events involving only single mutations and the operation of the best-response dynamics in between mutations, the process may reach a state with a larger size of the α-component. This proves the first conclusion stated in the Lemma. A symmetric argument on the β-component shows that a similar chain of events may lead the process to increasing the size of this component, thus establishing the second stated conclusion as well. ||

We can now complete the proof of part (b) of Theorem 1. Let

$$m_\alpha \equiv \frac{K_\alpha}{K_\alpha + 1} n + \frac{1}{K_\alpha + 1} \tag{30}$$

stand for the minimum size of a stable α-component (for concreteness, with an odd number of players), where K_α is defined in Lemma 1. On the other hand, note that, from Lemma 3, the number of simultaneous mutations ℓ_β

that must occur at a state in S^α before a stable β component could be possibly formed through the operation of the best-response dynamics alone is bounded below as follows:

$$\ell_\beta \geq \frac{H_\beta}{H_\beta + 1} n - \frac{1}{H_\beta + 1}. \tag{31}$$

We now argue that:

$$H_\beta \equiv \frac{d-c}{b-e} > K_a \equiv \frac{2b-2c}{2d-e-c}. \tag{32}$$

To see this, simply note that our assumption that $d - b > b - f$ implies

$$d - c > 2b - 2c$$

and that, on the other hand, since $b < d$,

$$b - e < 2d - e - c.$$

Combining (30), (31), and (32), we conclude that, for large n,

$$\ell_\beta > m_\alpha.$$

This fact and the former Lemmata allow us to prove the desired conclusion, by relying on the by-now standard techniques of Freidlin and Wentzel (1984). In essence, these techniques require identifying, for each $s \in S^*$, the s-tree h that displays the minimum resistance cost $c(h)$, where these costs may be computed as the minimum number of mutations required for *all* the transitions contemplated in h. Once these minimum-cost trees are found, the stochastically stable states are simply those which have such minimum-cost tress with a cost that is itself minimum across all states in S^*. (See Kandori, Mailath and Rob (1993), Young (1993), or Kandori and Rob (1995) for a detailed description of the procedure.)

To apply this approach to our case, denote by $u \equiv |S^*|$ the cardinality of the set of equilibria. In view of Lemmas 1, 2, and 4, for any state in $s \in S^\alpha$, there is an s-tree h whose cost is

$$c(h) = \left[\frac{K_\alpha}{K_\alpha + 1} + \mathcal{O}(1/n)\right] \times n + (u-2)$$

On the other hand, by Lemmas 1 and 3, for any $s' \in S^* \backslash S^\alpha$, the cost of any s'-tree h' is bounded below as follows:

$$c(h') \geq \left[\frac{H_\beta}{H_\beta + 1} + \mathcal{O}(1/n) \right] \times n + (u - 2).$$

where we rely on the fact that $K_\beta > H_\beta$. Thus, in view of (32), we may conclude that, for large enough n, $c(h) < c(h')$, which is enough to establish the desired conclusion. This completes the proof of Part (b).

□

7 Appendix B

Proof of Proposition 4.1: We first show that every player in a component chooses the same action in equilibrium. Suppose player i chooses α while player j chooses β and they both belong to the same component. Let there be be k players in this component with k_α players choosing action α and k_β $(= k - k_\alpha)$ players choosing action β. The payoff to player i from action α is given by

$$(k_\alpha - 1)d + k_\beta e \tag{33}$$

The payoff to player i from action β is given by

$$(k_\alpha - 1)f + k_\beta b \tag{34}$$

Since, in equilibrium, player i prefers action α it follows that $(k_\alpha - 1)(d - f) \geq k_\beta(b - e)$. Similar calculations show that since, in equilibrium, player j prefers action β it must be true that $(k_\beta - 1)(b - e) \geq k_\alpha(d - f)$. Given that $d > f$ and $b > e$, this generates a contradiction. Thus every player in the same component must choose the same action, in equilibrium.

We now show that there can be at most two non-singleton components in an equilibrium network. Consider a non-empty network. It follows that there is at least one non-singleton component in this network. Let $g' \subset g$ be such a component. From above we know that every player in this component chooses the same action. Let this action be α, without loss of generality. We now show that every other player who chooses chooses α also belongs to this component. Suppose not. Let $i \in g'$ and $j \in g''$, where g' is as above and g'' is some other component of the network g. Suppose that $a_i = a_j = \alpha$. Also let k' and k'' be, respectively, the cardinality of the two components and let $k' \geq k''$. Recall that $k' \geq 2$, so there is some agent in g' who has formed a link. Suppose without loss of generality that this player is i. It follows from the hypothesis of Nash equilibrium that the return from a link must exceed the cost. The maximum possible payoff in the component g' is given by $(k' - 1)d$. It follows therefore that $(k' - 1)d \geq c$. Note however that if player j forms a link with some player in the component g' then his additional payoff is given by $k'd > (k' - 1)d \geq c$. This contradicts the optimality of the strategy of j. A similar argument holds if the players in the component choose β. This argument proves that there can be at most one non-singleton component comprising players choosing an action. Since there are only two actions in the coordination game it follows that there can be at most two non-singleton components in any equilibrium.

We next show that if n is large then there can be at most one non-singleton component in equilibrium. Consider an equilibrium network g. Suppose that there are two non-singleton components in g, g' and g'', with respective cardinalities k' and k''. Suppose without loss of generality that $k' \geq k''$. Then, it follows that $k' \geq n/2$. The payoffs to a player $j \in g''$ from forming a link with a player $i \in g'$ are bounded below by $(n/2)e$. Therefore, for large enough n, it is worthwhile for player i to form such a link. This contradicts the hypothesis that g is part of a Nash equilibrium. This proves that an equilibrium network is either connected or empty.

We finally consider the issue of minimal connectedness. Suppose that g is an equilibrium network and it is not minimally connected. This means that there is a link $g_{i,j} = 1$ such that $g - g_{i,j}$ is also connected. But then player i can delete the link $g_{i,j}$, and still be part of a connected society. Given the payoffs specified in (7), this implies that player i can strictly increase his payoffs, contradicting the hypothesis that g is an equilibrium network. This completes the proof.

$\square$

Proof of Proposition 4.2: From Proposition 4.1, we know that an equilibrium network is either empty or minimally connected. Consider a minimally connected equilibrium network g. Suppose that player i has a link with player j in this network, i.e. $g_{i,j} = 1$. We show that in a strict Nash equilibrium, this implies that player j does not have a link with any other player, i.e., $\overline{g}_{j,k} = 0$ for all $k \neq i$.

Suppose there is some player k such that $\overline{g}_{j,k} = 1$. In this case, individual i can simply interchange his link with j for a link with k and get the same payoffs. Thus, the strategy of forming a link with j is *not* a strict best response. Hence g is not a strict Nash network.

The above argument also implies that, since g is connected, player i must be linked to every other player directly. The resulting network is therefore a star. Moreover, it also follows that this link must be formed by player i himself. For otherwise, if there is a player k such that $g_{k,i} = 1$, then this player is again indifferent between the link with i and some other agent in the star. This implies that the star must be center-sponsored and completes the proof.

$\square$

Proof of Proposition 4.3: First, consider case (a). We know from Proposition 4.1 that every player in a component chooses the same action. We also know that there are only two possible equilbrium arhitectures, $g \in G^{cs}$ and g^e. Clearly, the empty network cannot be part of a strict Nash equilibrium (see also arguments for part (c) below). Thus the only candidates for strict Nash equilibrium are $s \in G^{cs} \times \{(\alpha, \alpha, ..., \alpha)\}$ or $s \in G^{cs} \times \{(\beta, \beta, ..., \beta)\}$. It is easily checked that any of those are indeed strict Nash equilibria.

Consider case (b) next. Again, the empty network is not sustainable by a strict Nash equilibrium. Then the only candidates are $s \in G^{cs} \times \{(\alpha, \alpha, ..., \alpha)\}$ or $s \in G^{cs} \times \{(\beta, \beta, ..., \beta)\}$. It is immediate to see that none of the latter is sustainable as an equilibrium since $c > b$, which implies that the central player does not have an incentive to form a link with isolated players. Thus the only remaining candidates are the former states, which are easily checked to be strict Nash equilibria.

Finally, consider case (c). If $c > d$, then the center-sponsored star cannot be an equilibrium network. Thus, the only candidate for a strict equilibrium network is the empty one. However, if a network is empty, the choice of actions is irrelevant. This means that there is no strict Nash equilibrium in this case. The proof is complete.

$\square$

References

[1] Allen, F. and D. Gale (1998), Financial Contagion, *Journal of Political Economy,* forthcoming.

[2] Anderlini, L. and A. Ianni (1996), Path-dependence and Learning from Neighbors, *Games and Economic Behavior*, 13, 2, 141-178.

[3] Bala, V. and S. Goyal (1999), A Non-Cooperative Model of Network Formation, *Econometrica*, forthcoming.

[4] Bala V. and S. Goyal (1998), Learning from Neighbours, *Review of Economic Studies* **65**, 595-621.

[5] Bergin, J. and B. Lipman, (1996), Evolution with State-Dependent Mutations, *Econometrica*, 64, 943-957.

[6] Bhaskar, V. and F. Vega-Redondo (1998), Migration and the evolution of conventions, WP *A Discusión*, Universidad de Alicante.

[7] Blume, L. (1993), Statistical Mechanics of Strategic Interaction, *Games and Economic Behavior* **4**, 387-424.

[8] Canning, D. (1992), Average Learning in Learning Models. *Journal of Economic Theory* **57**, 442-472.

[9] Carlson, H. and and E. van Damme (1993), Global Games and Equilibrium Selection, *Econometrica* **61**, 5, 989-1019.

[10] Chwe, M. (1998), Structure and Strategy in Collective Action, mimeo, University of Chicago.

[11] Coleman, J. (1966), *Medical Innovation: A Diffusion Study*, Second Edition, Bobbs-Merrill, New York.

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

[13] Ellison, G. (1993), Learning, Local Interaction, and Coordination. *Econometrica* **61**, 1047–1071.

[14] Ellison, G. and D. Fudenberg (1993), Rules of Thumb for Social Learning, *Journal of Political Economy* **101**, 612-644.

[15] Ely, J. (1996), Local Conventions, mimeo Northwestern University.

[16] Freĭdlin, M.I. and A. D. Venttsel′ (1979), *Fluktuatsii v Dinamicheskikh Sistemakh Pod Deĭstviem Malykh Sluchaĭnykh Vozmushcheniĭ.* Nauka, Moscow.

[17] Freidlin, M. I. and A. D. Wentzell (1984), *Random Perturbations of Dynamical Systems.* Translation of Freĭdlin and Venttsel′ (1979) by J. Szücs. Springer-Verlag, New York.

[18] Galesloot, B. and S. Goyal (1997), Costs of Flexibility and Equilibrium Selection, *Journal of Mathematical Economics* **28**, 249-264.

[19] Goyal, S. (1993), Sustainable Communication Networks, Tinbergen Institute Discussion Paper 93-250.

[20] Goyal, S. and M. Janssen (1997), Non-Exclusive Conventions and Social Coordination, *Journal of Economic Theory* **77**, 34-57.

[21] Granovetter, M. (1974), *Getting a Job: A Study of Contacts and Careers*, Harvard University Press, Cambridge MA.

[22] Granovetter, M (1985), Economic Action and Social Structure: The Problem of Embeddedness, *American Journal of Sociology* **3**, 481-510.

[23] Haag, M. and R. Lagunoff (1999), Social Norms, Local Interaction, and Neighborhood Planning, *mimeo*, Georgetown University.

[24] Harsanyi, J.C. and R. Selten (1988), *A General Theory of Equilibrium Selection*, Cambridge, Mass., MIT Press.

[25] Jackson, M. and A. Wolinsky (1996), A Strategic Model of Economic and Social Networks, *Journal of Economic Theory* **71**, 1, 44-74.

[26] Jackson, M. and A. Watts (1999), On the formation of interaction networks in social coordination games, mimeo, Caltech.

[27] Kandori, M., and R. Rob, (1995), Evolution of equilibria in the long run: A general theory and applications, *Journal of Economic Theory* **65**, 383-414.

[28] Kirman, A. (1997), The Economy as an Evolving Network, *Journal of Evolutionary Economics* **7**, 339-353.

[29] Kandori, M. and G. J. Mailath and R. Rob (1993), Learning, Mutation, and Long Run Equilibria in Games. *Econometrica* **61**, 29–56.

[30] Mailath, G. Samuelson, L. and Shaked, A., (1994), Evolution and Endogenous Interactions, mimeo., Social Systems Research Institute, University of Wisconsin.

[31] Morris, S. (1997), Contagion, mimeo Yale University.

[32] Nouweland, A. van den (1993), *Games and networks in Economic Situations*, unpublished Ph.D Dissertation, Tilburg University.

[33] Oechssler, J (1997), Decentralization and the coordination problem, *Journal of Economic Behavior and Organization* **32**, 119-135.

[34] Robson, A. and F. Vega-Redondo (1996), Efficient Equilibrium Selection in Evolutionary Games with Random Matching, *Journal of Economic Theory* **70**, 65-92.

[35] Samuelson, L. (1994), Stochastic stability in games with alternative best replies, *Journal of Economic Theory* **64**, 35-65.

[36] Young, H. P. (1993), The Evolution of Conventions. *Econometrica* **61**, 57–84.