


WORKING PAPERS

Mathias Staudigl

Potential games in volatile environments

February 2010

Working Paper No: 1002



DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

All our working papers are available at: <http://mailbox.univie.ac.at/papers.econ>

Potential games in volatile environments*

Mathias Staudigl[†]

February 4, 2010

Abstract

This paper studies the co-evolution of networks and play in the context of finite population potential games. Action revision, link creation and link destruction are combined in a continuous-time Markov process. I derive the unique invariant distribution of this process in closed form, as well as the marginal distribution over action profiles and the conditional distribution over networks. It is shown that the equilibrium interaction topology is an inhomogeneous random graph. Furthermore, a characterization of the set of stochastically stable states is provided, generalizing existing results to models with endogenous interaction structures.

Keywords: Markov process, Potential games, Stochastic stability, Network co-evolution, Random graphs

JEL Classification Numbers: C73, D83, D85

1 Introduction

The analysis of social networks has recently gained interest in various fields in the sciences and social sciences. By now there is a rich literature

*An earlier version circulated under the title “Efficiency in coordination games played in a volatile environment”. I would like to thank two anonymous referees and the audience of the EEA-ESEM 2009 in Barcelona which helped to improve the paper considerably. Special thanks to my supervisors Immanuel Bomze and Manfred Nermuth for helpful discussions, as well as Christina Pawlowitsch and Stefano DeMichelis for important recommendations. The usual disclaimer applies.

[†]Department of Economics, University of Vienna, Hohenstaufengasse 9, A-1010 Vienna (Austria). E-mail: mathias.staudigl@univie.ac.at

on social networks in economics; the textbooks by Jackson (2008) and Vega-Redondo (2007) give a concise overview on this emerging field. Recently, tools from evolutionary game theory have been used to study the co-evolution of networks and play. Models in this vein are Jackson and Watts (2002), Goyal and Vega-Redondo (2005), and Hojman and Szeidl (2006). Another type of model, which is more in the tradition of statistical physics, puts more weight on modelling the evolution of the network, without paying too much attention to the role of strategic interactions. Prominent examples are Ehrhardt et al. (2006; 2008a;b). This paper aims to combine these two streams of literature in a very simplistic model. I present a stochastic co-evolutionary model which includes three sub-processes: action adjustment, link creation, or link destruction. These three sub-processes are combined into one continuous-time Markov process called a co-evolutionary model with noise. For positive noise levels the process is ergodic. For the class of potential games (Monderer and Shapley, 1996) many fundamental characteristics of the system are explicitly computable. Key to all the results in this paper is the closed-form expression of the invariant distribution. This probability distribution describes the long-run behavior of the system in two complementary ways. First, it gives us complete information on the joint probability distribution over action profiles and networks which governs the “equilibrium” of the stochastic dynamics. Second, by virtue of ergodicity, it gives us complete information which states are more frequently observed over time compared to others. From the invariant distribution one can deduce the conditional probability distribution over networks for a fixed profile of actions. In the parlance of random graph theory this gives us the ensemble of random graphs. The interesting result is that the model generates so-called *inhomogeneous random graphs*. Inhomogeneous random graphs are a straightforward extension of the classical random graph model proposed by Erdős and Rényi (1960), where the probability with which two vertices are linked depend in some way on

the characteristics of the vertices. Söderberg (2002) and Bollobás et al. (2007) are models in this direction. These papers fix the edge success probability at the outset. On the contrary, in the present model the edge success probability is derived from the long-run equilibrium of the system, hence is explained *endogenously*. To the best of my knowledge, this is a new result, which opens the way to interesting linkages between evolutionary game theory and the theory of random graph dynamics. This relationship is further explored in the companion article Staudigl (2009b). Next, I provide an expression for the marginal distribution over action profiles in the population. This measure is interesting if one is not interested in the effects of the interaction structure. Finally, we explore the well-known relationship between potential maximizers and stochastic stability (for early work in this direction see for instance Blume, 1993, Young, 1998, ch. 6). A fairly general argument is provided, showing that as noise vanishes the invariant distribution concentrates on the set of potential maximizers. At first sight, this might not be a too surprising result. However, former models were only concerned with fixed interaction structures, so the conclusion of our theorem extends the previous ones. Moreover, the argument presented in this paper is much more general than the proofs in the just mentioned literature. This technique allows to study the low-noise behavior of the invariant distribution also in more complicated models, as for instance Staudigl (2009a). Since the class of potential games is rather narrow, I also sketch briefly how the results obtained extend if the potential game assumption is dropped. In Staudigl (2009b) a rather general class of co-evolutionary dynamics is presented, and I refer to this paper for further details. However, many games arising in economic applications have this special structure. The most prominent class of potential games are congestion games (Rosenthal, 1973). They also arise in Cournot oligopoly models with linear inverse demand functions (Monderer and Shapley, 1996). Recently, Sandholm (2007) studies a mechanism design problem where the planner can construct a pricing

scheme, so that the transformed game is a potential game, which leads, in his model, to the long-run selection of socially efficient outcomes. Ui (2000) has shown interesting interconnections between the Shapley value and potential functions, and Morris and Ui (2005) use potential methods to study equilibria which are robust to incomplete information.

Closest to the present work is a recent paper by Ehrhardt et al. (2008b), who study a similar dynamic process. Their link formation mechanism is designed in such a way that only players who play the same action form a link. This is interpreted as a pure homophily based linking process. They also characterize the induced ensemble of random graphs, and find that the network consists of disjoint components, each following the distribution of an Erdős-Rényi ensemble. This paper extends their result by allowing for much more general behavioral rules, both in the action adjustment and the link creation process, which results in a richer interaction structure.

The rest of the paper is organized as follows. In Section 2, the model framework is explained in detail. In sections 3 and 4, I derive the asymptotic characteristics of the model. Sections 5 and 6 present an analysis of the joint distribution of action profiles and social networks as well as the induced marginal distributions. Section 7 characterizes stochastically stable states. Section 8 sketches a general class of stochastic processes on the co-evolution of networks and play. A, found at the end of the paper, collects lengthy and technical proofs.

2 The model

Consider a finite population of individuals $i, j, k \in \mathcal{I} = \{1, 2, \dots, N\}$, members of which are called players or agents. Each player can choose one out of q different pure actions from the set $\mathcal{A} = \{a_1, a_2, \dots, a_q\}$. I will also say “playing action r ” with the understanding that this is action a_r . An *action profile* (configuration) is a tuple $\alpha = (\alpha^i)_{i \in \mathcal{I}} \in \mathcal{A}^{\mathcal{I}}$. When

individual i meets individual j , they engage in a 2-player game defined by the payoff function $u : \mathcal{A}^2 \rightarrow \mathbb{R}$. We assume that this function is symmetric in the following sense:

Assumption 1.

$$(\forall a, a' \in \mathcal{A}) : u(a, a') = u(a', a) \quad (2.1)$$

Games with this special property are known as (exact) *potential games* (Monderer and Shapley, 1996). This defines the *base game* $\Gamma^b := (\mathcal{A}, u)$.

The interaction structure is modeled as an undirected graph (network). Let $\mathcal{I}^{(2)}$ denote the set of unordered pairs of players. There are $N(N - 1)/2$ such pairs. A graph is a pair $G = (\mathcal{I}, \mathcal{E})$, where we interpret \mathcal{I} as the set of vertices (nodes) and $\mathcal{E} = \mathcal{E}(G) \subseteq \mathcal{I}^{(2)}$ the set of *edges* (links). An edge is an unordered pair of players $(i, j) \equiv (j, i)$ with the interpretation that if $(i, j) \in \mathcal{E}$, then players i and j play against each other. If $\mathcal{E} = \mathcal{I}^{(2)}$ we obtain the *complete graph* on \mathcal{I} , denoted by G^c . In this graph each individual is connected to everybody else and we obtain the standard global matching model. If $\mathcal{E} = \emptyset$ then we speak of the *empty graph* G^e . A graph $G' = (\mathcal{I}', \mathcal{E}')$ is a *subgraph* of $G = (\mathcal{I}, \mathcal{E})$ if $\mathcal{I}' \subseteq \mathcal{I}$ and $\mathcal{E}' \subseteq \mathcal{E}$. For two disjoint subsets of players $\mathcal{V}, \mathcal{V}' \subset \mathcal{I}$ denote the set of edges that join players from \mathcal{V} with players belonging to \mathcal{V}' (and vice-versa) as $\mathcal{E}(\mathcal{V}, \mathcal{V}')$. All graphs on \mathcal{I} differ only in terms of their edge set \mathcal{E} . Let $\mathcal{G}[\mathcal{I}]$ denote the set of graphs that can be formed on the vertex set \mathcal{I} . It is often more convenient to work with networks via the function $g : \mathcal{I}^{(2)} \times \mathcal{G}[\mathcal{I}] \rightarrow \{0, 1\}$, assigning to each pair $(i, j) \in \mathcal{I}^{(2)}$ the value $g((i, j), G) \equiv g((j, i), G) \equiv g_j^i(G) \in \{0, 1\}$. If $g_j^i(G) = 1$ then players i and j are linked under the graph G and play against each other. Thus, we have the identity $\mathcal{E}(G) = \{(i, j) \in \mathcal{I}^{(2)} \mid g_j^i(G) = 1\}$ for all graphs $G \in \mathcal{G}[\mathcal{I}]$. It follows that every graph $G \in \mathcal{G}[\mathcal{I}]$ can be identified through the *realization* of links $g(G) = (g_j^i(G))_{1 \leq i < j \leq N} \in \{0, 1\}^{\mathcal{I}^{(2)}}$. In view of this equivalence, we will identify the space $\mathcal{G}[\mathcal{I}] \equiv \mathcal{G}$ as the set of all possible edge realizations $\{0, 1\}^{\mathcal{I}^{(2)}}$, members of which are vectors $g = (g_j^i)_{1 \leq i \leq j \leq N}$. The number of edges of the graph g is $e(g) := \sum_{i=1}^N \sum_{j>i} g_j^i$.

A *population state* is the pair $\omega = (\alpha, g) \in \Omega \equiv \mathcal{A}^{\mathcal{I}} \times \mathcal{G}$. It contains an action profile and a network. Let $\alpha_i^{a_v} := (\alpha^1, \dots, \alpha^{i-1}, a_v, \alpha^i, \dots, \alpha^N)$. Let $g \oplus (i, j)$ denote the network that we obtain if the (previously non-existing) edge connecting players i and j is created, and $g \ominus (i, j)$ be the network resulting from the deletion of the edge connecting players i and j .

Given a population state ω , define for every player $i \in \mathcal{I}$ the (open) *interaction neighborhood*

$$\mathcal{N}^i(\omega) = \bigcup_{r=1}^q \{j \in \mathcal{I} | g_j^i = 1 \ \& \ \alpha^j = a_r\}.$$

The set $\mathcal{N}^i \cup \{i\} \equiv \bar{\mathcal{N}}^i$ defines the closed interaction neighborhood of a player. There are $\kappa_r^i(\omega) := |\{j \in \mathcal{I} | g_j^i = 1 \ \& \ \alpha^j = a_r\}|$ r -players against which player i has to play. The total number of games in which player i is involved is given by his *degree* $\kappa^i(\omega) = \sum_{r=1}^q \kappa_r^i(\omega)$. From all these interactions, player i receives the total payoff

$$\pi(\alpha, g) \equiv \pi^i(\omega) := \sum_{j \in \mathcal{N}^i(\omega)} u(\alpha^i, \alpha^j) = \sum_{r=1}^q u(\alpha^i, a_r) \kappa_r^i(\omega). \quad (2.2)$$

In analogy with standard population games, I will call the collection of payoff functions $\pi = (\pi^i)_{i \in \mathcal{I}}$ the *structured population game*.

3 Co-evolution with noise

Consider the family of perturbed Markov processes

$$\mathcal{M}^\beta = (\Omega, \mathcal{F}, \mathbb{P}, (Y^\beta(t))_{t \geq 0}, \beta \in \mathbb{R}_+),$$

where Ω is the finite state space of pairs $\omega = (\alpha, g)$, \mathcal{F} a suitably chosen σ -algebra (e.g. 2^Ω), $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ a probability measure, and $(Y^\beta(t))_{t \geq 0}$ a family of Ω -valued random variables indexed by a noise parameter $\beta \geq 0$ and a continuous time parameter t . \mathcal{M}^β will define a co-evolutionary

model with noise. The time evolution of this process can be studied by its infinitesimal generator. Define the operator $\eta^\beta := [\eta^\beta(\omega \rightarrow \omega')]_{\omega, \omega' \in \Omega}$ whose components are mappings $\eta^\beta : \Omega \times \Omega \rightarrow \mathbb{R}$ satisfying $0 \leq \eta^\beta(\omega \rightarrow \hat{\omega}) < \infty$ for all $\hat{\omega} \neq \omega$, and $\sum_{\hat{\omega}} \eta^\beta(\omega \rightarrow \hat{\omega}) = 0$ for all $\omega \in \Omega$. The value $\eta^\beta(\omega \rightarrow \hat{\omega})$ is interpreted as the *rate* with which the process moves from state ω to some other state $\hat{\omega}$.¹ The generator is defined by the following sub-processes.

Action update: The way how players update their actions is modeled as in Blume (2003) or Hofbauer and Sandholm (2007). Players are endowed with independent Poisson alarm clocks, ringing at the common rate $\nu > 0$. The total rate of this subprocess is thus $N\nu$. Conditional on the event of a revision opportunity, player i receives the chance to adjust his action with probability $1/N$. When player i gets a revision opportunity he calculates the current expected payoff of all of his pure actions, given the set of neighbors, but his computations are perturbed by some random shock $\varepsilon^i = (\varepsilon_a^i)_{a \in \mathcal{A}}$. Assume that these perturbations are i.i.d. type 1 extreme value distributed,² and that i selects action $a_r \in \mathcal{A}$ with probability

$$b^i(a_r|\omega) := \mathbb{P} \left(a_r \in \arg \max_{a_v \in \mathcal{A}} (\pi^i(\alpha_i^{a_v}, g) + \varepsilon_{a_v}^i) | \omega \right). \quad (3.1)$$

Computing this probability explicitly leads to

$$b^{i,\beta}(a_r|\omega) = \frac{\exp(\pi^i(\alpha_i^{a_r}, g) / \beta)}{\sum_{v=1}^q \exp(\pi^i(\alpha_i^{a_v}, g) / \beta)}. \quad (3.2)$$

¹In a very small time interval $[t, t+h)$, the probability that the process moves from ω to $\hat{\omega}$ is then approximately $\eta^\beta(\omega \rightarrow \hat{\omega})h$.

²This formulation of stochastically perturbed payoffs has a very long tradition in the theory of discrete choice, see e.g. Anderson et al. (1992). For a more recent treatise and alternative interpretation see van Damme and Weibull (2002). The cumulative distribution function of a doubly exponential distributed random variable with mean 0 and variance $\frac{\beta^2 \pi^2}{6}$ is $F(x) = \exp[-\exp(-x/\beta - \gamma)]$. Beside its importance in theoretical economics, it has also been used in experimental studies, see for instance McKelvey and Palfrey (1995), where it is known as the “quantal response function”.

The transition $\omega = (\alpha, g) \rightarrow \hat{\omega} = (\alpha_i^{ar}, g) \neq \omega$ proceeds therefore at a rate

$$\eta^\beta(\omega \rightarrow \hat{\omega}) = v b^i(a_r | \omega). \quad (3.3)$$

Link creation: Here ideas of the *stochastic-actor model*, developed in Snijders (2001), are used. The key-ingredients of this model are a *rate function*, governing the pace at which individuals update their connections, and an *objective function*, capturing the preferences of the individuals concerning link creation. For the rate function, I make the following assumption:

Assumption 2. *The rate functions of individuals take the form*

$$(\forall i \in \mathcal{I})(\forall \omega \in \Omega) : \lambda^{i,\beta}(\omega) = \sum_{k \notin \mathcal{N}^i(\omega)} \exp(u(\alpha^i, \alpha^k) / \beta). \quad (3.4)$$

This formulation reflects the intuitive idea that players, who expect a large profit from interactions with currently unknown players, should be relatively fast in creating their network. Let

$$\bar{\lambda}^\beta(\omega) := \sum_{i \in \mathcal{I}} \lambda^{i,\beta}(\omega) = 2 \sum_{i,j > i} \exp(u(\alpha^i, \alpha^j) / \beta) (1 - g_j^i),$$

so that the conditional probability that player i receives a link creation opportunity is simply $\lambda^{i,\beta}(\omega) / \bar{\lambda}^\beta(\omega)$. Conditional on this event, player i screens the set of unknown players (i.e. those player who are not neighbors yet) and picks one player from this set who yields the highest per-interaction payoff, perturbed by a noisy signal $\zeta^i = (\zeta_k^i)_{k \notin \mathcal{N}^i(\omega)}$. Hence, the conditional probability that i selects j for a linking partner is

$$w_j^i(\omega) := \mathbb{P} \left(u(\alpha^i, \alpha^j) + \zeta_j^i \geq u(\alpha^i, \alpha^k) + \zeta_k^i \quad \forall k \notin \mathcal{N}^i | \omega \right). \quad (3.5)$$

If we assume that the random perturbation follows the same distributional law as in the action adjustment process one obtains the logit formula

$$(\forall i \in \mathcal{I})(\forall j \notin \mathcal{N}^i(\omega)) : w_j^{i,\beta}(\omega) = \frac{\exp(u(\alpha^i, \alpha^j) / \beta)}{\sum_{k \notin \mathcal{N}^i(\omega)} \exp(u(\alpha^i, \alpha^k) / \beta)}. \quad (3.6)$$

For general link creation probabilities (3.5), the rate of transiting from state $\omega = (\alpha, g)$ to state $\hat{\omega} = (\alpha, g \oplus (i, j))$ is

$$\eta^\beta(\omega \rightarrow \hat{\omega}) = \lambda^i(\omega)w_j^i(\omega) + \lambda^j(\omega)w_i^j(\omega). \quad (3.7)$$

Using Assumption 2 and (3.6) gives us

$$\eta^\beta(\omega \rightarrow \hat{\omega}) = 2 \exp(u(\alpha^i, \alpha^j) / \beta). \quad (3.8)$$

Link destruction: To make the dynamic interesting, we need a process that counteracts the creation of links. Following recent papers by Ehrhardt et al. (2006; 2008b), I assume that there exists an *exogenous* random shock removing any of these links. This unguided drift term models the phenomenon of *environmental volatility*, and is a key ingredient of the model. It captures the idea that connections are not everlasting, but as time goes by and players change their behavior, the profitability of links will also change, making some connections obsolete. The rate at which the link (i, j) disappears is given by $\xi > 0$.³ Hence, in a very small time interval $[t, t + h)$ the *probability of survival* of a currently existing edge (i, j) is $\xi h + o(h)$. The expected life time of an edge is $1/\xi$. Hence, starting from $\omega = (\alpha, g)$, the transition rate to $\hat{\omega} = (\alpha, g \ominus (i, j))$ is

$$\eta^\beta(\omega \rightarrow \hat{\omega}) = \xi. \quad (3.9)$$

The last case we have to consider is a “phantom switch”, i.e. the transition rate $\eta^\beta(\omega \rightarrow \omega)$. Define the rate of such an event as

$$\eta^\beta(\omega \rightarrow \omega) = -\Lambda^{(\beta, \xi)}(\omega), \quad (3.10)$$

³The assumption of constant link decay rates is less restrictive as it may seem. Since link creation probabilities are payoff driven, players will be more likely to establish links which are associated with higher per-interaction payoff. Hence, if a highly valuable link disappears, *ceteris paribus*, there is a relatively high probability that it will be re-established in future periods. Extending to heterogeneous link destruction rates is straightforward. See Staudigl (2009a).

where

$$\Lambda^{(\beta, \xi)}(\omega) := \nu \sum_{i=1}^N \sum_{a \in \mathcal{A} \setminus \{\alpha^i\}} b^{i, \beta}(a|\omega) + \xi e(g) + \bar{\lambda}^\beta(\omega). \quad (3.11)$$

To summarize, the infinitesimal generator of the co-evolutionary model with noise \mathcal{M}^β is defined as

$$\eta^\beta(\omega \rightarrow \hat{\omega}) = \begin{cases} \nu b^{i, \beta}(a|\omega) & \text{if } \hat{\omega} = (\alpha_i^a, g) \neq \omega, \\ 2 \exp(u(\alpha^i, \alpha^j)/\beta) & \text{if } \hat{\omega} = (\alpha, g \oplus (i, j)), \\ \xi & \text{if } \hat{\omega} = (\alpha, g \oplus (i, j)), \\ -\Lambda^{(\beta, \xi)}(\omega) & \text{if } \hat{\omega} = \omega, \\ 0 & \text{otherwise.} \end{cases} \quad (3.12)$$

It is easily verified that $\sum_{\hat{\omega} \in \Omega} \eta^\beta(\omega \rightarrow \hat{\omega}) = 0$ for all $\omega \in \Omega$. For $\beta > 0$ we observe that $\eta^\beta(\omega \rightarrow \omega') > 0$ for $\omega \neq \omega'$, meaning that there can be no single state that is absorbing. Irreducibility of the generator follows from this easily. Furthermore, in view of the finiteness of the state space, positive recurrence of the process follows. Hence, the Markov process is ergodic.

4 The invariant distribution

By ergodicity, the co-evolutionary model with noise admits a unique invariant distribution $\mu^{(\beta, \xi)} = (\mu^{(\beta, \xi)}(\omega))_{\omega \in \Omega}$. In terms of the generator η^β , this probability distribution satisfies the global balance equation $\mu^{(\beta, \xi)} \eta^\beta = \mathbf{0}$. Determining this probability vector is facilitated in the special class of reversible Markov processes. Given the model \mathcal{M}^β with generator η^β , we can define for a given $T > 0$ its *time reversal* as the process $(\hat{Y}^\beta(t))_{0 \leq t \leq T}$, with $\hat{Y}^\beta(t) = Y^\beta(T - t)$. A Markov process $(Y^\beta(t))_{t \geq 0}$ is said to be *reversible*, if its time reversal has the same distribution as the original process (see Stroock, 2005, ch. 5). In our case, reversibility will appear as an equilibrium phenomenon (i.e. $(Y(t)^\beta)_{t \geq 0}$ is *reversible in equilibrium*). The *detailed balance condition*, relative to the infinitesimal gener-

ator η^β , gives a sufficient condition for $\mu^{(\beta, \bar{\xi})}$ being an invariant distribution. The measure $\mu^{(\beta, \bar{\xi})}$ is in detailed balance with the generator η^β if

$$(\forall \omega, \hat{\omega} \in \Omega) : \mu^{(\beta, \bar{\xi})}(\omega) \eta^\beta(\omega \rightarrow \hat{\omega}) = \mu^{(\beta, \bar{\xi})}(\hat{\omega}) \eta^\beta(\hat{\omega} \rightarrow \omega). \quad (4.1)$$

A probability distribution satisfying the detailed balance condition (4.1) must be an invariant distribution. Conversely, a probability distribution satisfying condition (4.1) implies reversibility of the corresponding Markov process.

Theorem 4.1. *Given $(\beta, \bar{\xi}) \gg (0, 0)$, the unique invariant distribution of the co-evolutionary model with noise \mathcal{M}^β equals*

$$(\forall \omega \in \Omega) : \mu^{(\beta, \bar{\xi})}(\omega) = \frac{1}{Z^{(\beta, \bar{\xi})}} \prod_{i=1}^N \prod_{j>i} \left[\frac{2}{\bar{\xi}} \exp\left(\frac{u(\alpha^i, \alpha^j)}{\beta}\right) \right]^{\bar{\xi}^i}, \quad (4.2)$$

where $Z^{(\beta, \bar{\xi})} \equiv \sum_{\omega \in \Omega} \prod_{i=1}^N \prod_{j>i} \left[\frac{2}{\bar{\xi}} \exp\left(\frac{u(\alpha^i, \alpha^j)}{\beta}\right) \right]^{\bar{\xi}^i}$ is the the partition function.

Proof. See Appendix A. □

A consequence of ergodicity is the convergence of long-run averages of sample paths to the invariant distribution. Formally, this means

$$\mathbb{P} \left(\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mathbb{1}_{\{Y^\beta(s) = \omega\}} \, ds = \mu^{(\beta, \bar{\xi})}(\omega) \right) = 1,$$

where $\mathbb{1}_A$ is the indicator function of a measurable set $A \subseteq \Omega$, and for any integrable function $f : \Omega \rightarrow \mathbb{R}$

$$\mathbb{P} \left(\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(Y^\beta(s)) \, ds = \mathbb{E}_{\mu^{(\beta, \bar{\xi})}}[f] \right) = 1,$$

where $\mathbb{E}_{\mu^{(\beta, \bar{\xi})}}[f] = \sum_{\omega \in \Omega} f(\omega) \mu^{(\beta, \bar{\xi})}(\omega)$ is the expected value of the function f under the invariant distribution $\mu^{(\beta, \bar{\xi})}$.

Observe that for $\beta > 0$ $\mu^{(\beta, \bar{\xi})}$ is a full support distribution on Ω . Thus, the only thing one may be able to deduce from it is to classify a subset of

states which receive more mass than others. The subsequent chapters are devoted to this exercise.

Define an aggregate utility index as the sum of individual utilities,

$$(\forall \omega \in \Omega) : U(\omega) = \sum_{i=1}^N \pi^i(\omega) = 2 \sum_{i=1}^N \sum_{j>i} u(\alpha^i, \alpha^j) g_j^i \quad (4.3)$$

Efficiency, in terms of this index, is a state in the argmax set of (4.3). Lemma 4.1 shows that one can construct from eq. (4.3) a real-valued function, which captures the effects of individual utilities due to a single change in the state variable ω . In game theory such a function is known as an *exact potential* (Monderer and Shapley, 1996). Since the state variable encompasses the connections among the players, but these are not part of the strategy of a single player, a potential function for $\pi = (\pi^i)_{i \in \mathcal{I}}$ is not a potential function in its game-theoretic sense. However, it fulfills the same role in the dynamic analysis to come as a conventional potential function in the sense of Monderer and Shapley (1996), and so we will still call such a function a potential function for the structured population game, having in mind that this does not conform with its established use in game theory.

Lemma 4.1. *The structured population game $(\pi^i)_{i \in \mathcal{I}}$ is a potential game with exact potential function*

$$(\forall \omega \in \Omega) : P(\omega) = \frac{1}{2} \sum_{i=1}^N \pi^i(\omega) = \sum_{i=1}^N \sum_{j>i} u(\alpha^i, \alpha^j) g_j^i. \quad (4.4)$$

Proof. We have to show that

$$\begin{aligned} P(\alpha_i^{a_v}, g) - P(\alpha, g) &= \pi^i(\alpha_i^{a_v}, g) - \pi^i(\alpha, g), \text{ and} \\ P(\alpha, g \oplus (i, j)) - P(\alpha, g) &= u(\alpha^i, \alpha^j) \end{aligned}$$

Let us start with the event of a link creation between players i and j . The destruction of such a link has the same consequences. A direct computation shows that

$$P(\alpha, g \oplus (i, j)) - P(\alpha, g) = \frac{1}{2}(u(\alpha^i, \alpha^j) + u(\alpha^j, \alpha^i)) = u(\alpha^i, \alpha^j),$$

by symmetry of the payoff function u . Now, concerning a change in action of player i , we now have to take care of the environment of this player. All players in the set $\mathcal{I} \setminus \bar{\mathcal{N}}^i(\omega)$ are not affected by the change in player i 's action. Fix the state ω and suppose player i switches to action a_v . The new state is therefore $\omega_i^{a_v}$, and we can write

$$\begin{aligned}
U(\alpha_i^{a_v}, g) &= \sum_{k=1}^N \sum_{j: g_j^k=1} u(\alpha^k, \alpha^j) \\
&= U(\alpha, g) + \sum_{j: g_j^i=1} [u(a_v, \alpha^j) - u(\alpha^i, \alpha^j)] + \sum_{\ell: g_\ell^i=1} [u(\alpha^\ell, a_v) - u(\alpha^\ell, \alpha^i)] \\
&= U(\alpha, g) + 2 \sum_{j: g_j^i=1} [u(a_v, \alpha^j) - u(\alpha^i, \alpha^j)]
\end{aligned}$$

□

Now consider the function $H : \Omega \times \mathbb{R}_+ \times \mathbb{R}_{++} \rightarrow \mathbb{R}$, defined as

$$H(\omega, \beta, \zeta) := P(\omega) + \beta e(\omega) \log \left(\frac{2}{\zeta} \right) \quad (4.5)$$

This function acts as a *graph Hamiltonian* for the invariant distribution.⁴ One sees that there are two components combined in the graph Hamiltonian. The first component is the potential function of the game, which measures (up to a linear scaling) the aggregate utility of the population. The second part is a size measure of the interaction graph, weighted by the volatility parameter ζ . If $\zeta > 2$ then too large graphs (measured by the number of edges) lead to a reduction in the value of the Hamiltonian. This effect is in turn weighted by the noise level β . Proposition 4.1 shows that it contains all the information one needs to determine the invariant distribution of the process \mathcal{M}^β . Its proof is straightforward and therefore omitted.

⁴For a general discussion of this concept see Park and Newman (2004). In statistical mechanics a Hamiltonian is, roughly, a measure of the energy of a system. In the simplest case it is the sum of the potential energy and kinetic energy. This description fits also perfectly to the form of the Hamiltonian (4.5).

Proposition 4.1. *The stationary distribution of the co-evolutionary model with noise \mathcal{M}^β is the Gibbs measure*

$$\mu^{(\beta, \zeta)}(\omega) = \frac{e^{\frac{1}{\beta}H(\omega, \beta, \zeta)}}{\sum_{\hat{\omega} \in \Omega} e^{\frac{1}{\beta}H(\hat{\omega}, \beta, \zeta)}}. \quad (4.6)$$

From the definition of the Hamiltonian (4.5), one can see that a large value of β , combined with $\zeta > 2$, implies that too large graphs will not receive too much weight in the long run. A small value of β means in turn that, for any given volatility level ζ , the penalty of densely connected societies has a small influence on the invariant distribution. It is exactly this trade-off between β and volatility ζ which makes the form of the invariant distribution interesting. High environmental volatility, accompanied with moderate noise will lead to a sparsely connected society.

5 The ensemble of random graphs

Given $\omega = (\alpha, g) \in \Omega$, define the set of r -players as $\mathcal{I}_r(\omega) := \{i \in \mathcal{I} | \alpha^i = a_r\}$. Sets of this form will be called action classes. Every state assigns each player to a single action class. Hence, the family $\{\mathcal{I}_r\}_{1 \leq r \leq q}$ defines a partition on the set \mathcal{I} . Fix a partition $\mathcal{I} \equiv \{\mathcal{I}_r\}_{1 \leq r \leq q}$ and define the subspace

$$\Omega(\mathcal{I}) := \{\omega \in \Omega | \mathcal{I}_r(\omega) = \mathcal{I}_r, 1 \leq r \leq q\}.$$

We say that state ω agrees with the action partition \mathcal{I} , if it is contained in $\Omega(\mathcal{I})$. Note that the definition of the set $\Omega(\mathcal{I})$ does not say anything about network structures. Once we condition on an action partition, we fix a strategy configuration $\alpha \in \mathcal{A}^{\mathcal{I}}$, but allow for all potential networks. In other words, $\mu^{(\beta, \zeta)}(\omega | \mathcal{I}) \equiv \mu^{(\beta, \zeta)}(g | \alpha)$.

Given a partition \mathcal{I} , the product operator $\prod_{i=1}^N \prod_{j>i}$ has the same meaning as the product operator $\prod_{r=1}^q \prod_{i \in \mathcal{I}_r(\omega)} \prod_{v \geq r} \prod_{j \in \mathcal{I}_v(\omega); j > i}$. This implies that we are able to re-formulate the stationary distribution in terms of

action classes, so that for all $\omega \in \Omega$

$$\mu^{(\beta, \xi)}(\omega | \mathcal{I}) \propto \prod_{r=1}^q \prod_{i \in \mathcal{I}_r} \left\{ \prod_{v \geq r} \prod_{j \in \mathcal{I}_v, j > i} \left[\frac{2}{\xi} \exp\left(\frac{u(a_r, a_v)}{\beta}\right) \right]^{g_j^i} \right\} \mathbf{1}_{\{\omega \in \Omega(\mathcal{I})\}}. \quad (5.1)$$

For proper normalization of this measure one has to compute the total mass received by the set $\Omega(\mathcal{I})$, which is

$$\mu^{(\beta, \xi)}(\Omega(\mathcal{I})) = \sum_{g \in \mathcal{G}} \mu^{(\beta, \xi)}(\alpha, g).$$

Let $e_{r|v}(\omega) := \sum_{i \in \mathcal{I}_r(\omega)} \sum_{j \in \mathcal{I}_v(\omega), j > i} g_j^i$, denote the number of edges connecting r -players with v -players at state ω , and define for all $(i, j) \in \mathcal{I}^{(2)}$

$$p_{i,j}^{(\beta, \xi)}(\omega) := \frac{2 \exp(u(\alpha^i, \alpha^j) / \beta)}{2 \exp(u(\alpha^i, \alpha^j) / \beta) + \xi}, \quad \theta_{i,j}^{(\beta, \xi)}(\omega) := \log \left(\frac{p_{i,j}^{(\beta, \xi)}(\omega)}{1 - p_{i,j}^{(\beta, \xi)}(\omega)} \right).$$

Setting $\theta^{(\beta, \xi)}(\omega) := (\theta_{i,j}^{(\beta, \xi)}(\omega))_{(i,j) \in \mathcal{I}^{(2)}}$, we get

$$\begin{aligned} \frac{1}{\beta} H(\omega; \beta, \xi) &= \frac{1}{\beta} \sum_{i=1, j>i}^N u(\alpha^i, \alpha^j) g_j^i + \log(2/\xi) \sum_{i=1, j>i}^N g_j^i \\ &= \sum_{i=1, j>i}^N \left[\log(\exp(u(\alpha^i, \alpha^j) / \beta) + \log(2/\xi)) \right] g_j^i \\ &= \sum_{i=1, j>i}^N \theta_{i,j}^{(\beta, \xi)}(\omega) g_j^i =: h[\omega, \theta^{(\beta, \xi)}(\omega)] \end{aligned}$$

Given an action partition \mathcal{I} , consider the subgraph $G_{r|v} := (\mathcal{I}_r \cup \mathcal{I}_v, \mathcal{E}_{r|v})$, where $\mathcal{E}_{r|v} = \mathcal{E}(\mathcal{I}_r, \mathcal{I}_v)$. For all $(i, j) \in [\mathcal{I}_r \cup \mathcal{I}_v]^{(2)}$, the numbers $p_{i,j}^{(\beta, \xi)}, \theta_{i,j}^{(\beta, \xi)}$ are constant, so that we may write $p_{r|v}^{(\beta, \xi)}$ and $\theta_{r|v}^{(\beta, \xi)}$. Lemma A.1 in Appendix A shows that

$$\mu^{(\beta, \xi)}(\Omega(\mathcal{I})) \propto \prod_{r=1}^q \prod_{v \geq r} \left(1 - p_{r|v}^{(\beta, \xi)} \right)^{-\frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}}} \quad (5.2)$$

where $\delta_{x,y} = 1$ if, and only if, $x = y$, and 0 otherwise. The main result of this section is then the following result.

Theorem 5.1 (The Erdős-Rényi Decomposition). *Fix an action partition \mathcal{I} and $(\beta, \xi) \gg (0, 0)$.*

(a) *The measure (5.1) is the conditional distribution over graphs $g \in \mathcal{G}$ and factorizes to*

$$\mu^{(\beta, \xi)}(\omega | \mathcal{I}) = \prod_{r=1, v \geq r} \left[p_{r|v}^{(\beta, \xi)} \right]^{e_{r|v}(\omega)} \left[1 - p_{r|v}^{(\beta, \xi)} \right]^{\frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}} - e_{r|v}(\omega)}. \quad (5.3)$$

(b) *The statistical ensemble of subgraphs $\mathcal{G}[\mathcal{I}_r \cup \mathcal{I}_v]$ is an Erdős-Rényi graph with edge success probability*

$$p_{r|v}^{(\beta, \xi)} = \frac{2 \exp(u(a_r, a_v) / \beta)}{2 \exp(u(a_r, a_v) / \beta) + \xi}.$$

Proof. See Appendix A. □

Part (a) of the Theorem shows that the equilibrium ensemble of graphs boils down to an *inhomogeneous random graph* (Söderberg, 2002, Bollobás et al., 2007). For an arbitrary action profile eq. (5.3) gives us complete information about the probability with which an r -strategist interacts with players from other action classes. Thus, if one wants to make a probabilistic prediction about the interaction pattern between r -players and v -players, all one has to do is to look at the factor

$$\left[p_{r|v}^{(\beta, \xi)} \right]^{e_{r|v}(\omega)} \left[1 - p_{r|v}^{(\beta, \xi)} \right]^{\frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}} - e_{r|v}(\omega)}$$

what is exactly the probability measure of the random graph model of Erdős and Rényi (1960). Since this the interactions among r and v -players follow a Bernoulli distribution, the expected number of interactions, given the profile α , is determined by the formula

$$\mathbb{E}_{\mu^{(\beta, \xi)}}[e_{r|v} | \mathcal{I}] = \frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}} p_{r|v}^{(\beta, \xi)}.$$

For the covariances we see that

$$\text{Cov}_{\mu^{(\beta, \xi)}}[e_{r|v}, e_{r|l} | \mathcal{I}] = \begin{cases} 0 & \text{if } v \neq l, \\ \frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}} p_{r|v}^{(\beta, \xi)} (1 - p_{r|v}^{(\beta, \xi)}) & \text{if } v = l. \end{cases}$$

The fact that the total graph can be regarded as a collection of independent Erdős-Rényi graphs (with different edge success probabilities) makes it possible to derive a probability distribution for the degree of a randomly selected individual $i \in \mathcal{I}_r$. Since $\kappa^i = \sum_{v=1}^q \kappa_v^i$, we first have to determine the distribution of the random variables $\kappa_v^i, 1 \leq v \leq q$. Theorem 5.1 tells us that κ_v^i has a Binomial distribution with parameters $(|\mathcal{I}_v| - \delta_{r,v}, p_{r|v}^{(\beta, \xi)})$ (see e.g. Bollobás, 1998).

Proposition 5.1. *Given an action partition \mathcal{I} pick a player $i \in \mathcal{I}_r$ and let $n_v := |\mathcal{I}_v|, 1 \leq v \leq q$. The degree of player i is distributed according to the mass function*

$$f_{r, \kappa}(k | \mathcal{I}) := \frac{1}{R^{(\beta, \xi)}(\mathcal{I})} \sum_{k_1 + \dots + k_q = k} \frac{k!}{k_1! \dots k_q!} \prod_{v=1}^q \left[f_{r|v}^{(\beta, \xi)}(k_v) \right]^{k_v}, \quad (5.4)$$

$$f_{r|v}^{(\beta, \xi)}(k_v) := \binom{n_v - \delta_{r,v}}{k_v}^{1/k_v} \left(\frac{p_{r|v}^{(\beta, \xi)}}{1 - p_{r|v}^{(\beta, \xi)}} \right). \quad (5.5)$$

where $R^{(\beta, \xi)}(\mathcal{I})$ is the normalizing factor.

Proof. See Appendix A. □

Observe that for the degree distribution it suffices to know the number of players in the various action classes, not their identity. Hence, all action partitions \mathcal{I} that put the same number of players into the various classes are equivalent in terms of the connectivity structure of the network. Thus, instead of looking at a specific action partition \mathcal{I} , it is sufficient to work with less information contained in a tuple $n = (n_1, \dots, n_q)$ such that $n_v = |\mathcal{I}_v|$ for all v and $\sum_{v=1}^q n_v = N$.

Example 1. Consider the coordination game

| | | |
|-------|--------|--------|
| | a_1 | a_2 |
| a_1 | (3, 3) | (0, 0) |
| a_2 | (0, 0) | (1, 1) |

We will examine the degree distribution for 1-players under various parameter constellations (β, ξ) for the frequency vector $n = (80, 20)$. Figure 1 shows the degree distribution for a typical 1-player under the parameter constellation $(\beta, \xi) = (0.5, 70)$. The mean degree of 1-players is seen to be 78. However, we




Figure 1: Degree distributions for 1-players under various parameter constellations. The triple at the top of each plot is (β, ξ, \bar{k}) , i.e. the noise and volatility rate and the resulting average degree for this action class. The point marks the position of the mean of this distribution.

cannot say to which action class most of this links lead to since we only look at the distribution of the total degree κ . Applying Theorem 5.1, we get complete information about the inter-group connectivity pattern by inspecting the two numbers

$$p_{1|1} = \frac{2 \exp(3/\beta)}{2 \exp(3/\beta) + \xi}, \quad p_{1|2} = \frac{2}{2 + \xi}$$

Note that $p_{1|1} \rightarrow 1$ as $\beta \rightarrow 0$, implying that in this limit only links within the same action class exist with probability 1. Consequently, for small noise levels the majority of the 78 neighbors will be 1-players as well. For larger levels of noise (the right figure with $\beta = 1.5$) we observe a drastically smaller average

degree. This implies that the effect of the parameter values β and ξ goes into the same direction. Increasing β with constant ξ will have qualitatively the same effect as increasing ξ with constant β .

6 An invariant distribution over action profiles

Having derived a probability distribution on the set of networks, we will now derive a probability distribution on the set of action frequency vectors $n = (n_1, \dots, n_q)$. Let $\mathcal{D} := \{n \in \mathbb{N}^q \mid \sum_{r=1}^q n_r = N\}$ denote the set of admissible action frequency vectors and define the correspondence $\Psi : \mathcal{D} \rightarrow 2^\Omega$ as $\Psi(n) = \{\omega \in \Omega \mid (\forall r = 1, 2, \dots, q) : |\mathcal{I}_r(\omega)| = n_r\}$.

Proposition 6.1. *The invariant distribution over action frequency vectors $n \in \mathcal{D}$ is given by the mapping $\rho^{(\beta, \xi)} = \mu^{(\beta, \xi)} \circ \Psi : \mathcal{D} \rightarrow [0, 1]$, defined as*

$$\rho^{(\beta, \xi)}(n) := \mathcal{K}^{-1} \frac{N!}{\prod_{r=1}^q n_r!} \prod_{r=1}^q \left[z_r^{(\beta, \xi)}(n) \right]^{n_r}, \quad (6.1)$$

where

$$z_r^{(\beta, \xi)}(n) := \prod_{v \geq r} \left[1 + \frac{2}{\xi} \exp \left(\frac{u(a_r, a_v)}{\beta} \right) \right]^{\frac{n_v - \delta_{r,v}}{1 + \delta_{r,v}}}, \quad 1 \leq r \leq q \quad (6.2)$$

$$\mathcal{K} = \sum_{n \in \mathcal{D}} \frac{N!}{\prod_{r=1}^q n_r!} \prod_{r=1}^q \prod_{v \geq r} \left[1 + \frac{2}{\xi} \exp \left(\frac{u(a_r, a_v)}{\beta} \right) \right]^{\frac{n_r(n_v - \delta_{r,v})}{1 + \delta_{r,v}}}. \quad (6.3)$$

Proof. The proof starts from the distribution over action classes \mathcal{I} (5.2). The rest is a simple combinatorial exercise. The population consists of N distinct elements. There are q different boxes over which we want to distribute the N elements, and in each box $r = 1, \dots, q$ there should be n_r elements at the end of the day, and all N elements must be in one box, so that $\sum_{r=1}^q n_r = N$ holds. There are $\frac{N!}{n_1! \dots n_q!}$ different ways of solving this allocation problem. Counting all states ω that agree with a given action class size profile n leads to a probability distribution having the form

$$\mu^{(\beta, \xi)}(\Psi(n)) \propto \frac{N!}{\prod_{r=1}^q n_r!} \prod_{r=1}^q \prod_{v \geq r} \left[1 + \frac{2}{\xi} \exp \left(\frac{u(a_r, a_v)}{\beta} \right) \right]^{\frac{n_r(n_v - \delta_{r,v})}{1 + \delta_{r,v}}}. \quad (6.4)$$

Using the respective definitions of the maps $\rho^{(\beta, \xi)}$ and $z_r^{(\beta, \xi)}(n)$ yields the desired result. \square

7 Stochastic stability

Stochastic game dynamics have become important due to their power concerning equilibrium selection. The concept of stochastic stability, introduced by Foster and Young (1990), Young (1993) and Kandori et al. (1993) into game theory, gives a selection criterion based on the underlying dynamic process.

Definition 1. *A state $\omega \in \Omega$ is a stochastically stable state if*

$$\lim_{\beta \rightarrow 0} \mu^{(\beta, \xi)}(\omega) > 0.$$

The set of stochastically stable states is denoted as Ω^ .*

It has been shown by Blume (1993; 1997) and Young (1998) that the logit dynamics concentrates on the set of potential maximizers as the noise level goes to zero. However, their results are not directly applicable in the current context, since the graph is itself part of the state variable.

7.1 Selection of Potential maximizers

The following Theorem, the proof of which is based on the general discussion in Catoni (1999), is the main result of this section.

Theorem 7.1. *The Gibbs distribution (4.6) concentrates on the set $\mathcal{P} := \arg \max_{\omega \in \Omega} P(\omega)$ as $\beta \rightarrow 0$.*

Proof. See Appendix A. \square

This shows that in the limit of vanishingly small noise the process will spend almost all of its time in the vicinity of potential maximizers. In view of the relation between the potential function and aggregate utility,

this gives an efficiency result for long run behavior. Furthermore, in view of the ergodic theorem, which has been mentioned briefly in Section 4, we know that long run averages of the potential function converge to the expected value under the invariant distribution $\mu^{(\beta, \xi)}$. Since this expected value converges to $\max_{\omega \in \Omega} P(\omega)$ as $\beta \rightarrow 0$, we get the following corollary.

Corollary 7.1. *Let $\mathcal{U} := \arg \max_{\omega \in \Omega} U(\omega) = \mathcal{P}$. Then*

$$\lim_{\beta \rightarrow 0} \mu^{(\beta, \xi)}(\mathcal{U}) = 1 \quad (7.1)$$

Almost surely therefore the process arrives at states where social welfare is maximized.

7.2 Efficiency in pure coordination games

Consider the class of games with payoff function $u(a, a') := \phi(a, a') - c$, that satisfies condition (2.1), as well as

$$\begin{aligned} (\forall r = 1, 2, \dots, q) : \max_{1 \leq v \leq q} \phi(a_v, a_r) &= \phi(a_r, a_r), \\ \phi(a_1, a_1) &\leq \phi(a_2, a_2) \leq \dots \leq \phi(a_q, a_q). \end{aligned} \quad (7.2)$$

The first condition states that matching the action chosen by the opponent is always a best reply. The second condition imposes an ordering on the payoffs of actions, where a_q denotes the payoff dominant action. From the symmetry of the payoff function, eq. (2.1), it follows that there are q strict Nash equilibria in the base game where the two players choose the same action. The constant $c \geq 0$ has no strategic effect, and can be interpreted as the costs of a link.⁵ To keep notation simple, suppose that all strict Nash equilibria have different payoffs. Let $g^e = g(G^e)$, $g^c = g(G^c)$ denote the empty and the complete graph, respectively.

⁵Jackson and Watts (2002), Goyal and Vega-Redondo (2005) consider symmetric 2×2 coordination games, which are potential games, having this payoff structure.

Proposition 7.1. *Let $P : \Omega \rightarrow \mathbb{R}$ be the potential function (4.4), and suppose that the payoff function of the base game G^b satisfies eq. (7.2). Then*

$$\mathcal{P} = \begin{cases} \{(a_q, \dots, a_q)\} \times \{g^e\} & , \text{if } u(a_q, a_q) > 0 \\ (\mathcal{A}^{\mathcal{I}} \times \{g^e\}) \cup \{\omega \in \Omega \mid \alpha = (a_q, \dots, a_q)\} & , \text{if } u(a_q, a_q) = 0 \\ \mathcal{A}^{\mathcal{I}} \times \{g^e\} & , \text{if } u(a_q, a_q) < 0 \end{cases}$$

Proof. It is straightforward to see that the potential function (4.4) can be written as

$$P(\omega) = \sum_{r=1}^q \sum_{v \geq r} u(a_r, a_v) \left(\sum_{i \in \mathcal{I}_r(\omega)} \sum_{j \in \mathcal{I}_v(\omega); j > i} g_j^i \right).$$

From this one can immediately see the validity of the claim for the high-cost scenario $u(a_q, a_q) < 0$.

Now consider the case where $u(a_q, a_q) = 0$. Clearly $P(\omega) \leq 0$ for all $\omega \in \Omega$, with equality only at the states that are in the set described in the text of the Proposition.

Finally, consider the case $u(a_q, a_q) > 0$. Since this is the largest payoff obtainable from the base game, and the potential function is linear in the links, the claim follows. This is also the unique maximizer of the potential function. \square

Corollary 7.2. *Consider the co-evolutionary model \mathcal{M}^b , with base game from the class of pure-coordination games (7.2). Then $\Omega^* = \mathcal{P}$.*

8 A general class of stochastic co-evolutionary dynamics

The model presented so far relied on the assumptions that the base game has an exact potential, and the rate functions of the individual players have the particular form (3.4). These assumptions make the model very tractable, and we were able to deduce many fundamental characteristics

of the long-run behavior of the system. On the other hand, one may say that these assumptions are too strict. Let me shortly discuss how the model can be extended to a rather general class of co-evolutionary models with noise. For a detailed discussion I refer to the companion paper Staudigl (2009b). There a rather general class of perturbed time-homogeneous Markov chains, similar to \mathcal{M}^β , is presented, where players may have heterogeneous preferences in the base game, but choose from a common action set.⁶ A general characterization of the invariant distribution of such models is provided, as well as an algorithm which identifies stochastically stable states, based on tree-constructions in the spirit of Freidlin and Wentzell (1998). The present model fits into this general framework, and let me just sketch what the long run behavior of this model would be, if one drops Assumptions 1 and 2. Instead of (3.4), assume that the players' rate function equals $\lambda^i(\omega) = \lambda \mathbb{1}_{\{\kappa^i(\omega) < N-1\}}$, and λ is a positive constant. For sake of illustration suppose the base game is a symmetric 2×2 coordination game with one Pareto efficient equilibrium (a_1, a_1) , and one risk-dominant equilibrium (a_2, a_2) . The specific payoffs are not important.⁷ I claim that these small alterations of the model lead to a *non-selection* result. Any pair of players, which use the same action, may be connected in the long-run equilibrium; putting it differently, as β goes to zero we do not obtain a point prediction as in Section 7, but the limit distribution will (in general) put positive weight on a proper subset of Ω . The heuristic explanation of this “negative” result is the following.

- Since the rate function of players is uncoupled with the noise parameter, the speed of the link creation process is unaffected by the level of noise. Looking back at (3.4), we see that as β goes to 0 the link creation process becomes arbitrary fast.
- The link destruction process deletes any edge with the constant rate

⁶An extension to different action sets is possible, but does not give more insights in the model.

⁷Of course, this is still a potential game.

ξ . This process is pure drift, i.e. it is not depending on the base game, and in particular is independent of the noise level β . In the terminology of stochastic stability calculus, this implies that link destructions are zero cost events. However, it turns out that the rate-ratio λ/ξ determines the number of links the system can carry in the long run.

- The logit choice function of the action adjustment process (3.2) puts equal probability on all actions a loner may choose. However, if a player has at least one neighbor and if β goes to zero, this player will play a best response against the neighbors' behavior with probability arbitrary close to 1.
- Suppose the system is currently in a full coordination state, say the population coordinates on the efficient equilibrium (a_1, \dots, a_1) . The network will not be complete in general, but one can derive a distribution over networks, given this action configuration. If there are some loners in the current state, let them switch to a_2 , and give them a link creation opportunity. These steps can be made with zero costs. Now, by definition of the coordination game, an optimal decision in the link creation process is to connect the a_2 players. We are then already in a state where a_1 and a_2 co-exist. At this state no player has an incentive to change his action, so we will not return to the state we were coming from. If there are no loners, we can construct a sequence of link destruction, action adjustment and link creation events, all causing no costs, which leads to a state where two coordination equilibria co-exist, as follows: Destroy the links of player i . Give him an action adjustment opportunity where he chooses a_2 . Since a loner may choose any action with equal probability without making an error, this causes no costs. Do the same thing with player $j \neq i$. Then give them a link creation opportunity. Since i and j are the only agents playing a_2 , an optimal decision

in the link creation process is to create the link (i, j) . Now we are in a co-existence state and no player has an incentive to change his action.

- In the same vein we can walk through the set

$$\Omega^* = \{\omega \in \Omega \mid g_j^i = 1 \Rightarrow \alpha^i = \alpha^j\}$$

without any costs, in the sense of stochastic stability analysis. As a result, all states contained in this set are stochastically stable.

A similar result, but with admittedly sharper limit predictions, is obtained in the model of Jackson and Watts (2002). These authors add to the drift term ζ a direction, by assuming that only links where at least one player is better off after the destruction of the link, are very likely to become destroyed. For a fairly large set of parameters (such as linking costs as in Section 7) they also get a co-existence result. However, due to this directionality in the link destruction process, they get sharper limit results in the network dimension under the assumptions that the costs per link are constant. The framework presented in Staudigl (2009b) is sufficiently flexible to capture the model of Jackson and Watts (2002).

9 Conclusion

This paper presented a stylized model on the co-evolution of networks and play in the class of potential games. Assumption 2 was crucial to derive a closed-form solution of the unique invariant distribution and to obtain sharp predictions as the noise in the players' decision rules goes to zero. A general selection theorem of potential maximizers applies in this case. Without Assumption 2 the invariant distribution can still be completely characterized, but the model loses its predictive power in the low-noise limit. It seems therefore that some assumptions in this direction are needed if one wants to obtain sharp limiting predictions.

There are many possible routes for extensions. In a companion work (Staudigl, 2009a) I analyze the current model with Assumption 2, but assuming an inverse relationship in the rate function with the size of the population. The intuition is that a larger population should make it less likely that a single agent receives the chance to create a link. In the infinitely large population limit and small positive noise the generated networks do not converge to complete graphs anymore. Hence, nicer asymptotic results are obtained without losing much in analytical power. A more fundamental question is, however, which class of networks (in the sense of random graph theory) such co-evolutionary models are capable to create. One first step in this direction is the companion work Staudigl (2009b). There I propose a rather general model of co-evolutionary models with noise, which is rich enough to incorporate the just presented model, as well as the “volatility” models of Ehrhardt et al. (2008b) and Jackson and Watts (2002). A first result is that such models seem to generate, under fairly mild assumptions on the structure of the random process, so-called inhomogeneous random graphs (see e.g. the nice survey by Newman, 2003). These models are straightforward extensions of the classical Erdős-Rényi model, where the edge-success probabilities depend on the attributes of the individual vertices. It would be interesting to see how deep this connection indeed is.

A Proofs of Selected Theorems and Propositions

Proof of Theorem 4.1. Uniqueness follows from irreducibility and recurrence of η^β .

By construction of the dynamics, we know that changes occur in the process only in one “coordinate”: either a single change in the links of the network takes place, or one, and almost surely only one, player switches to another action. By statistical independence of these two processes we can treat them separately. Start with a change in the network structure. It suffices to consider the creation of a fresh link. Let $\omega = (\alpha, g), \hat{\omega} = (\alpha, g \oplus (i, j)) \in \Omega$. The rate of link creation

between players i and j is given by eq. (3.8). The rate with which one returns to the state ω is eq. (3.9). Detailed balance (4.1) demands that

$$\frac{\mu^{(\beta, \xi)}(\hat{\omega})}{\mu^{(\beta, \xi)}(\omega)} = \frac{2}{\xi} \exp(u(\alpha^i, \alpha^j) / \beta). \quad (\text{A.1})$$

It is easy to see that the measure (4.2) satisfies this condition.

Now consider the event of action adjustment. Let player k be the one who receives such an opportunity and suppose she switches to action $a_v \in \mathcal{A}$. Let $\omega, \hat{\omega} = (\alpha_k^{a_v}, g) \in \Omega$ be the states involved in this transition. The associated rate ratio is

$$\begin{aligned} \frac{\eta^\beta(\omega \rightarrow \hat{\omega})}{\eta^\beta(\hat{\omega} \rightarrow \omega)} &= \frac{\nu b^{k, \beta}(a_v | \omega)}{\nu b^{k, \beta}(\alpha^k | \omega)} \\ &= \exp \left[\frac{1}{\beta} \left(\sum_{j: g_j^k = 1} [u(a_v, \alpha^j) - u(\alpha^k, \alpha^j)] \right) \right] \end{aligned} \quad (\text{A.2})$$

Rewrite the invariant distribution as

$$\mu^{(\beta, \xi)}(\omega) \propto \prod_{i=1}^k \prod_{j>i} \left[\frac{2}{\xi} \exp(u(\alpha^i, \alpha^j) / \beta) \right]^{g_j^i} \times \prod_{i=k+1}^N \prod_{j>i} \left[\frac{2}{\xi} \exp(u(\alpha^i, \alpha^j) / \beta) \right]^{g_j^i}.$$

Note that the second term on the right-hand side does not depend on player k , and thus the change in the action of this player does have no effect on this term.

Hence, we see that the probability ratio boils down to

$$\frac{\mu^{(\beta, \xi)}(\hat{\omega})}{\mu^{(\beta, \xi)}(\omega)} = \prod_{i=1}^k \prod_{j>i} \left[\exp \left(\frac{u(\hat{\alpha}^i, \hat{\alpha}^j) - u(\alpha^i, \alpha^j)}{\beta} \right) \right]^{g_j^i}.$$

Since $\hat{\alpha}^i = \alpha^i$ for all $i \neq k$, $\hat{\alpha}^k = a_v$, and payoffs as well as the indicators g_j^i are symmetric, we see that

$$\begin{aligned} \frac{\mu^{(\beta, \xi)}(\hat{\omega})}{\mu^{(\beta, \xi)}(\omega)} &= \prod_{j=1}^N \left[\exp \left(\frac{u(a_v, \alpha^j) - u(\alpha^k, \alpha^j)}{\beta} \right) \right]^{g_j^k} \\ &= \exp \left[\frac{1}{\beta} \left(\sum_{j: g_j^k = 1} [u(a_v, \alpha^j) - u(\alpha^k, \alpha^j)] \right) \right] \end{aligned}$$

This is the rate ratio (A.2). □

Lemma A.1. Fix an action partition \mathcal{I} and let $\omega \in \Omega(\mathcal{I})$. Define

$$m^{(\beta, \xi)}(\omega|\mathcal{I}) = \prod_{r=1}^q \prod_{i \in \mathcal{I}_r} \left\{ \prod_{v \geq r} \prod_{j \in \mathcal{I}_v; j > i} \left[\frac{2}{\xi} \exp\left(\frac{u(a_r, a_v)}{\beta}\right) \right]^{\delta_j^i} \right\} \mathbf{1}_{\{\omega \in \Omega(\mathcal{I})\}} \quad (\text{A.3})$$

the mass of state ω , conditional on the event that the action partition \mathcal{I} is realized. Then the mass received by the set $\Omega(\mathcal{I})$ in the long run is given by

$$m^{(\beta, \xi)}(\Omega(\mathcal{I})) = \prod_{r=1, v \geq r}^q \left(1 - p_{r|v}^{(\beta, \xi)}\right)^{-\frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}}}. \quad (\text{A.4})$$

Proof. We have to compute $\sum_{\omega \in \Omega} m^{(\beta, \xi)}(\omega|\mathcal{I})$. On $\Omega(\mathcal{I})$ the action profile is fixed, and all states differ only in the number of edges. We can write

$$m^{(\beta, \xi)}(\omega|\mathcal{I}) = \prod_{r=1, v \geq r}^q \left(\frac{p_{r|v}^{(\beta, \xi)}}{1 - p_{r|v}^{(\beta, \xi)}} \right)^{e_{r|v}(\omega)}.$$

Hence

$$\begin{aligned} m^{(\beta, \xi)}(\Omega(\mathcal{I})) &= \sum_{\omega \in \Omega(\mathcal{I})} m^{(\beta, \xi)}(\omega) \\ &= \prod_{r=1, v \geq r}^q \sum_{k=0}^{\frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}}} \binom{\frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}}}{k} \left(\frac{p_{r|v}^{(\beta, \xi)}}{1 - p_{r|v}^{(\beta, \xi)}} \right)^k \\ &= \prod_{r=1, v \geq r}^q \left(1 - p_{r|v}^{(\beta, \xi)}\right)^{-\frac{|\mathcal{I}_r|(|\mathcal{I}_v| - \delta_{r,v})}{1 + \delta_{r,v}}} \end{aligned}$$

□

Proof of Theorem 5.1. (a) Lemma A.1 shows that (5.1) is given by

$$\mu^{(\beta, \xi)}(\omega|\mathcal{I}) = \frac{m^{(\beta, \xi)}(\omega|\mathcal{I})}{m^{(\beta, \xi)}(\Omega(\mathcal{I}))}.$$

A direct calculation of this ratio gives Eq. (5.3).

(b) This follows directly from the product measure (5.3) and the definition of the Erdős-Rényi-model.

□

Proof of Proposition 5.1. For ease of notation I skip again the parameters (β, ξ) . $\kappa_1^i, \dots, \kappa_q^i$ are independent Binomially distributed random variables with respective parameters $(n_v - \delta_{r,v}, p_{r|v}), 1 \leq v \leq q$. Thus

$$\mathbb{P}(\kappa^i = k_1, \dots, \kappa_q^i = k_q | \mathcal{I}, i \in \mathcal{I}_r) = \prod_{v=1}^q \mathbb{P}(\kappa_v^i = k_v | \mathcal{I}, i \in \mathcal{I}_r),$$

where for $1 \leq v \leq q$

$$\begin{aligned} \mathbb{P}(\kappa_v^i = k_v | \mathcal{I}, i \in \mathcal{I}_r) &= \binom{n_v - \delta_{r,v}}{k_v} p_{r|v}^{k_v} (1 - p_{r|v})^{n_v - \delta_{r,v} - k_v} \\ &= [f_{r|v}(k_v)]^{k_v} (1 - p_{r|v})^{n_v - \delta_{r,v}}. \end{aligned}$$

The function $f_{r|v}(\cdot)$ has been defined in the text of the Proposition. There are $\frac{k!}{k_1! \dots k_q!}$ ways to construct a list (k_1, \dots, k_q) whose sum equals k . Hence

$$\mathbb{P}(\kappa^i = k | \mathcal{I}, i \in \mathcal{I}_r) \propto \sum_{k_1 + \dots + k_q = k} \frac{k!}{k_1! \dots k_q!} \prod_{v=1}^q \mathbb{P}(\kappa^i = k_v | \mathcal{I}, i \in \mathcal{I}_r).$$

In each of the products on the right hand side, the factor $(1 - p_{r|v})^{n_v - \delta_{r,v}}$ is a constant and so cancels out after normalization. Hence, define the normalization factor

$$R(\mathcal{I}) = \sum_{k=0}^{N-1} \mathbb{P}(\kappa^i = k | \mathcal{I}, i \in \mathcal{I}_r),$$

and call $f_{r,\kappa}(k | \mathcal{I}) := \mathbb{P}(\kappa^i = k | \mathcal{I}, i \in \mathcal{I}_r)$ to get the desired result. \square

Proof of Theorem 7.1. For any $\varepsilon > 0$ consider the set $A_\varepsilon := \{\omega \in \Omega | P(\omega) < \max_{\omega' \in \Omega} P(\omega') - \varepsilon\}$. I will show that $\lim_{\beta \rightarrow 0} \mu^{(\beta, \xi)}(A_\varepsilon) = 0$. Let $P^* := \max_{\omega' \in \Omega} P(\omega')$ the global maximum value of the potential function, and $\mathcal{P} = \arg \max_{\omega \in \Omega} P(\omega)$ the set of maximizers. Define the measure $\mu_0^\xi : \mathcal{G} \rightarrow [0, \infty], g \mapsto \mu_0^\xi(g) := (2/\xi)^{e(g)}$. Since the Hamiltonian of the Gibbs measure is additive separable in the measure μ_0^ξ and the potential function P , we get for all $\omega = (\alpha, g)$

$$\mu^{(\beta, \xi)}(\alpha, g) \propto e^{\frac{1}{\beta} H(\omega, \beta, \xi)} = e^{\frac{1}{\beta} P(\omega)} \mu_0^\xi(g).$$

The set A_ε can be written as

$$A_\varepsilon = \{\omega \in \Omega | e^{-\frac{1}{\beta} P(\omega)} > e^{-\frac{1}{\beta} (P^* - \varepsilon)}\}.$$

Markov's inequality⁸ gives us

$$\mu^{(\beta, \zeta)}(A_\varepsilon) \leq e^{\frac{1}{\beta}(P^* - \varepsilon)} \mathbb{E}_{\mu^{(\beta, \zeta)}} \left[e^{-\frac{1}{\beta}P} \right],$$

where

$$\mathbb{E}_{\mu^{(\beta, \zeta)}} \left[e^{-\frac{1}{\beta}P} \right] = \sum_{\omega=(\alpha, g) \in \Omega} \mu^{(\beta, \zeta)}(\omega) e^{-\frac{1}{\beta}P(\omega)} = \frac{1}{Z} \sum_{\omega=(\alpha, g) \in \Omega} \mu_0^\zeta(g)$$

with $Z = \sum_{\omega \in \Omega} e^{\frac{1}{\beta}H(\omega, \beta, \zeta)} \geq |\mathcal{P}| e^{\frac{1}{\beta}H^*}$, and $H^* := \min_{\omega \in \mathcal{P}} H(\omega, \beta, \zeta)$. Let $K := \min_{\omega=(\alpha, g) \in \mathcal{P}} \mu_0^\zeta(g) > 0$ be the minimum value of the graph measure on the set of potential maximizers. Thus, $H^* \leq P^* + \beta \log K$, and so $Z \geq |\mathcal{P}| K e^{\frac{1}{\beta}P^*} > 0$. Next, we compute

$$\begin{aligned} \sum_{\omega=(\alpha, g) \in \Omega} \mu_0^\zeta(g) &= \sum_{n \in \mathcal{D}} \frac{N!}{n_1! \cdots n_q!} \prod_{r=1}^q \prod_{v \geq r} \prod_{i \in \mathcal{I}_r, j \in \mathcal{I}_v: g^j \in \{0,1\}} (2/\zeta)^{\delta_j^i} \\ &= \sum_{n \in \mathcal{D}} \frac{N!}{n_1! \cdots n_q!} \prod_{r=1}^q \prod_{v \geq r} (1 + 2/\zeta)^{\frac{nr(nv - \delta_r, v)}{1 + \delta_r, v}} \\ &= q^N (1 + 2/\zeta)^{\frac{N(N-1)}{2}}. \end{aligned}$$

where in the last step we have made use of the Multinomial Theorem. It follows that

$$\begin{aligned} \mu^{(\beta, \zeta)}(A_\varepsilon) &\leq \frac{1}{Z} q^N (1 + 2/\zeta)^{\frac{N(N-1)}{2}} e^{\frac{1}{\beta}(P^* - \varepsilon)} \leq \frac{q^N (1 + 2/\zeta)^{\frac{N(N-1)}{2}}}{K |\mathcal{P}| e^{\frac{1}{\beta}P^*}} e^{\frac{1}{\beta}(P^* - \varepsilon)} \\ &= K_1 e^{-\frac{\varepsilon}{\beta}}. \end{aligned}$$

where $K_1 := \frac{q^N (1 + 2/\zeta)^{\frac{N(N-1)}{2}}}{|\mathcal{P}| K} > 0$ a factor independent of β and ε . For $\beta \rightarrow 0$ the upper bound goes to zero, establishing the result. \square

References

Anderson, S., de Palma, A., Thisse, J., 1992. Discrete choice theory of product differentiation. The MIT Press.

⁸See e.g. Grimmett and Stirzaker (2001), p.319.

- Blume, L., 1993. Statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–424.
- Blume, L., 1997. Population games. In: Arthur, W., Durlauf, S., Lane, D. (Eds.), *The Economy as an evolving complex system II*. SFI Studies in the science of complexity, Vol. XXVII, Addison-Wesley, pp. 425–462.
- Blume, L. E., 2003. How noise matters. *Games and Economic Behavior* 44, 251–271.
- Bollobás, B., 1998. *Modern Graph Theory*. Springer: New York, Graduate Texts in Mathematics 184.
- Bollobás, B., Janson, S., Riordan, O., 2007. The phase transition in inhomogeneous random graphs. *Random Structures and Algorithms* vol. 31, Issue 1, 3–122.
- Catoni, O., 1999. Simulated annealing algorithms and Markov chains with rare transitions. *Séminaire de Probabilités XXXIII*. Lecture Notes in Mathematics, Vol. 1079.
- Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2006. Phenomenological models of socio-economic network dynamics. *Physical Review E* 74, 036106.
- Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2008a. Emergence and resilience of social networks: a general theoretical framework. *Annales d'économie et de statistique* 86, 1–13.
- Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2008b. Networks emerging in a volatile world. EUI Working papers.
- Erdős, P., Rényi, A., 1960. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences* 5, 17–61.

- Foster, D., Young, H. P., 1990. Stochastic evolutionary game dynamics. *Theoretical Population Biology* 38, 219–232.
- Freidlin, M., Wentzell, A., 1998. *Random Perturbations of Dynamical Systems*, 2nd Ed. Springer Verlag, New York.
- Goyal, S., Vega-Redondo, F., 2005. Network formation and social coordination. *Games and Economic Behavior* 50, 178–207.
- Grimmett, G., Stirzaker, D., 2001. *Probability and Random Processes*. Oxford University Press, USA; 3 edition.
- Hofbauer, J., Sandholm, W. H., 2007. Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory* 132, 47–69.
- Hojman, D. A., Szeidl, A., 2006. Endogenous networks, social games, and evolution. *Games and Economic Behavior* 55, Issue 1, 112–130.
- Jackson, M. O., 2008. *Social and Economic Networks*. Princeton University Press.
- Jackson, M. O., Watts, A., 2002. On the formation of interaction networks in social coordination games. *Games and Economic Behavior* 41, 265–291.
- Kandori, M., Mailath, G., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 61, 29–56.
- McKelvey, R., Palfrey, T., 1995. Quantal response equilibria for normal form games. *Games and Economic Behavior* 10, 6–38.
- Monderer, D., Shapley, L., 1996. Potential games. *Games and Economic Behavior* 14, 124–143.
- Morris, S., Ui, T., 2005. Generalized potentials and robust sets of equilibria. *Journal of Economic Theory* 124, 45–78.

- Newman, M., 2003. The structure and function of complex networks. *SIAM Review* 45, no.2, 167–256.
- Park, J., Newman, M., 2004. The statistical mechanics of networks. *Physical Review E* 70, 066117.
- Rosenthal, R. W., 1973. On a class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2, 65–67.
- Sandholm, W., 2007. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory* 132, 367–382.
- Snijders, T. A. B., 2001. The statistical evaluation of social network dynamics. *Sociological Methodology* Vol. 31, 361–395.
- Söderberg, B., 2002. General formalism for inhomogeneous random graphs. *Physical Review E* 66, 066121– 1– 6.
- Staudigl, M., 2009a. Co-evolutionary dynamics and bayesian interaction games. Mimeo, University of Vienna.
- Staudigl, M., 2009b. On a general class of stochastic co-evolutionary dynamics. Mimeo, University of Vienna.
- Stroock, D., 2005. *An Introduction to Markov Processes*. Graduate Texts in Mathematics 230, Springer Berlin Heidelberg New York.
- Ui, T., 2000. A shapley value representation of potential games. *Games and Economic Behavior* 31, 121–135.
- van Damme, E., Weibull, J. W., 2002. Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory* 106, 296–315.
- Vega-Redondo, F., 2007. *Complex Social Networks*. Cambridge University Press, *Econometric Society Monograph series*.
- Young, H., 1993. The evolution of conventions. *Econometrica* 61 (1), 57–84.

Young, H., 1998. Individual Strategy and Social Structure - An Evolutionary Theory of Institutions. Princeton University Press.