Int J Game Theory (2013) 42:179–210
DOI 10.1007/s00182-012-0331-0

Co-evolutionary dynamics and Bayesian interaction games

Mathias Staudigl

Accepted: 14 April 2012 / Published online: 25 April 2012

Abstract Recently there has been a growing interest in evolutionary models of play with endogenous interaction structure. We call such processes co-evolutionary dynamics of networks and play. We study a co-evolutionary process of networks and play in settings where players have diverse preferences. In the class of potential games we provide a closed-form solution for the unique invariant distribution of this process. Based on this result we derive various asymptotic statistics generated by the co-evolutionary process. We give a complete characterization of the random graph model, and stochastically stable states in the small noise limit. Thereby we can select among action profiles and networks which appear jointly with non-vanishing frequency in the limit of small noise in the population. We further study stochastic stability in the limit of large player populations.

Keywords Potential games · Network evolution · Heterogeneous populations · Inhomogeneous random graphs · Large deviations

1 Introduction

A co-evolutionary process of networks and play (Staudigl 2010b) is a stochastic dynamic process, where the interaction structure of the players evolves over time as a function of their chosen actions, which are themselves dynamic variables. These processes combine elements from evolutionary game theory with dynamic random graph theory. As such, any co-evolutionary model consists of a specification what

M. Staudigl (✉)
Institute of Mathematical Economics (IMW), Bielefeld University, Postfach 100131,
33501 Bielefeld, Germany
e-mail: mathias.staudigl@uni-bielefeld.de

the players do in the event of *action adjustments* and *network adjustments*.[1] These two events combined gives rise to a stochastic process on the product set of action profiles and networks. Studying these stochastic processes is in general complicated. However, Staudigl (2011) presents a model where analytical results can be obtained. In this previous paper we have focused on potential games (Hofbauer and Sigmund 1988; Monderer and Shapley 1996), combined with evolutionary game dynamics of the logit form.[2] The present paper provides an extension of this model along various lines.

1.1 Outline of the model and main results

First, we present an evolutionary process where players have heterogeneous preferences. Since (exact) potential games require a strong form of symmetry in game–theoretic models, this provides a significant extension to the setting of Staudigl (2011). We introduce payoff heterogeneity, by allowing players to have an idiosyncratic utility term, which is randomly determined before the co-evolutionary process starts. In a certain sense, this model can be interpreted as a Bayesian population game in the sense of Ely and Sandholm (2005). However, our population game is formulated on the individual player level, while a Bayesian population game is an aggregated version thereof. We will make this connection precise in Sect. 4.2 of this paper.[3] Second, our network formation dynamics is dependent of the size of the population. This is achieved by allowing the rates at which players create the network to depend on the population size. The idea is that, in a large population, not all players will be able to connect to everybody else, but only to a small subset of the population. Given these extensions, the present paper presents some new results, which are outside the framework of Staudigl (2011). We sketch these results briefly, before we start with the formal model.

As in Staudigl (2011), our analysis of the co-evolutionary process is focused on the *long-run* (or ergodic) properties of the stochastic process. Hence, our starting point is the characterization of the invariant distribution of the process, which is unique thanks to the presence of *behavioral noise* (or simply *noise*). Building on a theorem proved in Staudigl (2011), we provide a closed-form expression for this invariant probability measure. However, in the present model the invariant measure has to be interpreted as a *random element* of a set of invariant measures. This is because the realized type

[1] Related co-evolutionary models are Jackson and Watts (2002) and Goyal and Vega-Redondo (2005). The textbook by Vega-Redondo (2007) provides some alternative models in this direction. A recent survey of random graph models can be found in Durrett (2007).

[2] The logit dynamics is a very attractive alternative to the best-response with mutations model pioneered by Kandori et al. (1993) and Young (1993). The seminal reference in game theory seems to be Blume (1993), although this decision model has a long tradition in the discrete-choice literature. See, Anderson et al. (1992) for a discussion and references.

[3] Allowing for payoff heterogeneity allows us to define co-evolutionary processes for concrete economic phenomena. In an implementation problem Sandholm (2007) studies such games, assuming that players care about the population average behavior of the opponents. Our model is formulated in different lines, but we extend it to general random interaction structures. Young (2003) considers the diffusion of a new technology in exactly the class of games we study here, but he assumes a given fixed interaction structure.

profile affects the shape of the invariant distribution, so that every possible realization of the type profile induces a new invariant distribution.

Next we prove a general representation theorem, stating that a co-evolutionary process of networks and play generates so-called inhomogeneous, or generalized random graphs. This type of random graph has been studied in physics and mathematics intensively (see, e.g. Söderberg 2002 and Bollobás et al. 2007).[4] This result is an extension of Staudigl (2011), where player heterogeneity was ignored. We establish this fact by proving a rather general characterization theorem, to be found in Appendix A of the present paper. This result establishes an interesting and new connection with random graph theory and evolutionary game dynamics.

We then proceed by studying more in detail the role of the noise level, and the population size, on the structure of the invariant distribution. This kind of analysis of weak limits of the invariant measure is well established in evolutionary game theory as *stochastic stability analysis*.[5] Our stochastic stability analysis is divided into two categories:[6] First, we investigate stochastic stability in the *small noise limit*. In this analysis we want to characterize the support of the invariant distribution as the behavioral noise parameter vanishes. A population state contained in the support of this limiting measure is an action profile and a network, which is observed with a non-vanishing relative frequency in the limit of small noise. Under mild assumptions on the transition rates, stochastically stable states in the small noise limit coincide with the set of potential maximizers of a suitably defined potential function. Although this result is similar to the findings in Staudigl (2011), it provides an extension, as the potential function has to take care of the player heterogeneity in the population.

The dependence on the size of the population allows us to perform a stochastic stability analysis in the *large population limit*. A complication arises in this limit since the state space varies with the number of players. Thus, if we take this limit seriously we would have to define (i) a model of limit graphs, and (ii) a definition of limit action profiles describing the population with countably many players. In principle the characterization of the limit properties of the random graph model generated by the co-evolutionary process can be done, using the characterization results obtained for the finite model, and combining results from the literature on random graphs (see Remark 3). We will not pursue this analysis in this paper explicitly, but refer to the relevant literature. In this paper we are much more interested in characterizing the probability law of *aggregate behavior*, captured by the empirical distribution of actions and types among the players. This setting provides the

4 In the field of social and economic networks this random graph has been used by Golub and Jackson (2010) (who call it a multi-type random network) to study the influence of homophily in models of social learning.

5 This literature originated in the seminal papers by Kandori et al. (1993) and Young (1993). Blume (1993) is the seminal reference for stochastic stability in games with a fixed local interaction structure. The recent textbook Sandholm (2010b) provides a state of the art description of this technique.

6 We follow the terminology used by Sandholm (2010a).

finite-population analogue of the the Bayesian population games introduced by Ely and Sandholm (2005). We study in detail the long-run probability law of action-type distributions, which is derived by a careful aggregation procedure from the invariant distribution of the co-evolutionary process. In the limit of large player populations we want to determine the limit of this sequence of laws, in the sense of weak convergence of probability measures. Instead of characterizing the weak limit of this sequence of laws explicitly, we use the theory of large deviations to give a characterization of the support of the limiting measure. In the language of large deviations theory, this requires the identification of a *rate function*, measuring the exponential rate of decay of the stationary distribution weights. We provide such a characterization, and describe the set of maximizers of the rate function, which corresponds to the set of action-type distributions on which the limiting measure tends to concentrate.

The organization of the paper is as follows: In Sect. 2 we describe the game theoretic model and the dynamic evolutionary process of networks and play. Sect. 3 contains the characterization of the invariant distribution of the co-evolutionary process, and we provide the representation theorem for the random graph model. Sect. 4 contains the stochastic stability analysis. Sect. 4.1 is devoted to stochastic stability in the small noise limit, while Sect. 4.2 contains the large population limit results. Some technical details and proofs are provided in Appendices A and B.

2 The model

2.1 Interaction games

An interaction game consists of a finite number of players $[N] = \{1, 2, \ldots, N\}$ and a finite set of actions A_i for each player $i \in [N]$. We assume that all players make choices from the same common action set $A_i \equiv A = \{1, \ldots, n\}$. We denote by a_i the action of player i. Types of the players are contained in the finite set $\Theta := \{\theta_1, \ldots, \theta_K\}$, where each element of this set is a function from A to $\mathbb{R}$. Thus, each type can be identified with a vector $\boldsymbol{\theta}_k \in \mathbb{R}^n$, so that a probability distribution over types is an element of the unit simplex spanned by the K vectors $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K$.[7] The type of a player is a realization of a Θ-valued random variable $\tilde{\tau}_i$. A realization is denoted by τ_i. An action profile is a list $\mathbf{a} := (a_1, \ldots, a_N)$, and a (realized) type profile is a list $\boldsymbol{\tau} := (\tau_1, \ldots, \tau_N)$. For a given number of players N, we call $\mathcal{G}[N]$ the set of undirected simple graphs (networks) on the vertex set $[N]$. We identify a network with a list of $\{0, 1\}$-valued functions denoted by $\mathbf{g} = (g_{ij})_{(i,j)\in[N]^2}$. If $g_{ij} = 1$ then there is an edge connecting

[7] To make this set sufficiently rich, we assume that the collection of types $\{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K\}$ is an affinely independent set, i.e. for every collection of scalars $\lambda_1, \ldots, \lambda_K \in \mathbb{R}$ not all zero, we have

$$\sum_{k=1}^{K} \lambda_k \boldsymbol{\theta}_k = 0 \text{ and } \sum_{k=1}^{K} \lambda_k = 0 \Rightarrow \lambda_k = 0 \; \forall k = 1, 2, \ldots, K.$$

As is well-known, this assumption guarantees that the simplex $\Delta(\Theta)$ is a set of $K - 1$ dimensions, and just requires that the elements of the type space are identifiable as distinct points.

vertex i with j, and vice versa. Finally, we denote by $\mathbf{g} \oplus (i, j)$ the graph obtained from $\mathbf{g}$ by adding the edge between vertex i and j, and $\mathbf{g} \ominus (i, j)$ is the graph obtained from $\mathbf{g}$ by deleting the edge (i, j).

We consider games in which the utility function of the players consists of two parts. The first component is a *common utility* term, which one may think of as the externalities the players exert on each other. The second component is an *idiosyncratic payoff* term which depends on the player's own choice, but varies from player to player in a random way. Given an action profile $\mathbf{a}$ and a profile of types $\boldsymbol{\tau}$, the (ex-post) payoff of player i is assumed to be

$$U_i(\mathbf{a}, \mathbf{g}, \tau_i) = \sum_{j \neq i} g_{ij} v(a_i, a_j) + \tau_i(a_i).$$

In many applications it is conceivable that the common utility displays a symmetry property of the form

$$v(a, b) = v(b, a) \quad \forall a, b \in A.$$

Interaction games with such a partnership structure capture situations where all agents have the same reward function, and the payoff function of every player is the sum of all per-interaction rewards. However, having the partnership structure does not imply that all agents earn the same payoff in the interaction game since the interaction model will in general prescribe different interactions to different players.

An interaction game is played as follows: First we fix the number of players $N \in \mathbb{N}$.[8] Then each player learns his own type τ_i independently of any other player. The independent probability that player i will be of type θ_k is exogenously fixed, and given by some scalar $q_k \in (0, 1)$. The probability vector $\mathbf{q} := (q_1, \ldots, q_k) \in \operatorname{int} \Delta(\Theta)$ is called the *common prior*.[9] It is regarded as a parameter of the model. Given the realized type profile $\boldsymbol{\tau} \in \Theta^N$, players choose their actions and form the interaction network according to a *co-evolutionary process of networks and play*, as described below.

2.2 Co-evolution with noise

Following Staudigl (2011), we introduce a co-evolutionary process as a time-homogeneous Markov jump process $\{X_N^{\beta, \boldsymbol{\tau}}(t)\}_{t \geq 0}$, taking values on the finite state space $\mathcal{X}^N \equiv A^N \times \mathcal{G}[N]$.[10] An element of this space is denoted by $\mathbf{x} = (\mathbf{a}, \mathbf{g})$ and is called a *population state*. The process must be specified for the events of action revision, link creation and link destruction. The process of action revision is a

[8] The case $N = 1$ is trivial, but technically allowed.

[9] The set $\operatorname{int} \Delta(\Theta)$ are those elements of $\Delta(\Theta)$ with full support. The assumption that the common prior lies in this set is without loss of generality.

[10] The construction of such a process is standard, so that we omit technical details. For a thorough account of Markov processes in continuous time see, e.g., Ethier and Kurtz (1986).

random process defined for a *given* configuration of the network. It is the standard component of an evolutionary model, where agents receive randomly revision opportunities, and update their actions given the actions of their opponents and the network of the previous period. The other two processes model the evolution of the network. For these dynamics we take the action profile of the agents of the previous period as given and let the links in the network be created or destroyed. As is natural in the context of jump processes, we assume that in each updating step only one player may revise his action, or one link will be created or destroyed. The rates at which these events take place are described by the generator $\eta_N^{\beta,\boldsymbol{\tau}} = [\eta_N^{\beta,\boldsymbol{\tau}}(\mathbf{x},\mathbf{x}')]_{\mathbf{x},\mathbf{x}'\in\mathcal{X}^N}$. The initial population state $X_N^{\beta,\boldsymbol{\tau}}(0)$ is chosen according to some arbitrary probability law on the set $\mathcal{X}^N$. We now introduce the generator of the process.

Action adjustment: In case of a revision opportunity we assume that the agent switches to action $a \in A$ with probability determined by the log–linear response function

$$(\forall a \in A) : \ell_a^{i,\beta}(\mathbf{x}|\tau_i) = \frac{\exp\left[\beta^{-1}U_i(a,\mathbf{a}_{-i}),\mathbf{g},\tau_i)\right]}{\sum_{b\in A}\exp\left[\beta^{-1}U_i[(b,\mathbf{a}_{-i}),\mathbf{g},\tau_i]\right]}. \tag{1}$$

The *rate* of the transition $\mathbf{x} \to \mathbf{x}' = [(a,\mathbf{a}_{-i}),\mathbf{g}]$ is

$$\eta_N^{\beta,\boldsymbol{\tau}}(\mathbf{x},\mathbf{x}') = \lambda_1 \ell_a^{i,\beta}(\mathbf{x}|\tau_i),$$

where λ_1 is a non-negative constant, interpreted as the rate of the action revision process.

Link creation: A process of link creation describes the rates at which the indicator functions $(g_{ij})_{j>i}$ flip from 0 to 1. These rates are defined via an *attachment mechanism*. An *attachment mechanism* is a $n\times n$ matrix $\mathbf{C}^{\beta,N} = [c^{\beta,N}(a,b)]_{(a,b)\in A\times A}$, where the scalar $c^{\beta,N}(a,b)$ is interpreted as the *rate* that a player who is using action a meets a player who is playing action b. We make the assumption that these rates are of the form

$$c^{\beta,N}(a,b) = \frac{2}{N}\exp[v(a,b)/\beta] \quad \forall (a,b) \in A\times A. \tag{2}$$

The *rate* of a transition $\mathbf{x} = (\mathbf{a},\mathbf{g}) \to \mathbf{x}' = [(\mathbf{a},\mathbf{g}\oplus(i,j)]$ is then given by

$$\eta_N^{\beta,\boldsymbol{\tau}}(\mathbf{x},\mathbf{x}') = \lambda_2(1-g_{ij})\frac{2}{N}\exp[v(a_i,a_j)/\beta],$$

where λ_2 is the rate of the network formation process.

Link destruction: A process of link destruction describes the rates at which the indicator functions $(g_{ij})_{j>i}$ flip form 1 to 0. These rates define a *volatility mechanism*. For each pair of players (i,j) we define the rate of link destruction as

$$w_{ij}^{\beta,N}(\boldsymbol{\tau}) = \xi_{\tau_i,\tau_j}^{\beta,N} \tag{3}$$

where $\xi^{\beta,N}_{\theta_k,\theta_l} \equiv \xi^{\beta,N}_{k,l}$ is the *volatility rate* of a link between a player of type θ_k and a player of type θ_l. We identify the volatility mechanism with the $K \times K$ matrix $\Xi^{\beta,N} = \left(\xi^{\beta,N}_{k,l}\right)_{1\leq k,l\leq K}$. We assume that the following conditions are satisfied:

(SYM) $\xi^{\beta,N}_{kl} = \xi^{\beta,N}_{lk}$ for all $1 \leq k, l \leq K$;

(SNB) There exists a function $\bar{\xi}^{\beta,N} : \Theta \times \Theta \to \mathbb{R}_+ \cup \{\infty\}$ which satisfies $\lim_{\beta\to 0} \beta\bar{\xi}^{\beta,N}_{k,l} = 0$ for all $1 \leq k, l \leq K$ such that

$$0 < \underline{\xi}^N \leq \xi^{\beta,N}_{k,l} \leq c \exp(\bar{\xi}^{\beta,N}_{k,l}) \tag{4}$$

holds uniformly in $\beta \in (0, \infty)$ for some constants $\underline{\xi}^N$ and $c > 0$.

The rate of the transition $\mathbf{x} \to \mathbf{x}' = (\mathbf{a}, \mathbf{g} \ominus (i, j))$ is then given by

$$\eta^{\beta,\boldsymbol{\tau}}_N(\mathbf{x}, \mathbf{x}') = \lambda_2 g_{ij} w^{\beta,N}_{ij}(\boldsymbol{\tau}).$$

For positive β the Markov process $\{X^{\beta,\boldsymbol{\tau}}_N(t)\}_{t\geq 0}$ is easily seen to be ergodic. Hence, it possesses a unique invariant distribution $\mu^{\beta,\boldsymbol{\tau}}_N$, whose full characterization will be given in the next section. We close this section with a few remarks.

Remark 1 – For several properties of the model we have analytical answers without assuming the specific form for the attachment mechanism (2). In Staudigl (2010c) we present a much more general family of stochastic processes. However, the closed form of the stationary distribution relies heavily on this assumption. Since we make use of this closed-form in our equilibrium selection exercise we present here the more restrictive version.

– We have modeled the link creation process as only depending on the actions of the players, whereas the link destruction process depends only on the types of players. This modeling strategy is useful in order to disentangle the effect player heterogeneity has on the long-run network structure.

– To be fully clear, the rates of action adjustments (λ_1) and network formation (λ_2) are constant. Hence, although the rate of a meeting of players i and j, Eq. (2), is decreasing in the population size N, the frequency that the process undertakes a network adjustment step is independent of the population size. The interpretation of the dynamics is that with uniform rate λ_2 an edge (i, j) is selected. If the edge does not currently exist, it is created with rate $c^{\beta,N}(a_i, a_j)$. Otherwise it becomes destroyed with rate $w^{\beta,N}_{ij}(\boldsymbol{\tau})$. Note that if we set $\lambda_2 = 0$ then the network is never updated, and the co-evolutionary process reduces to a model of action adjustments on a fixed network (the initial graph). Similarly if $\lambda_1 = 0$ the actions of the individuals are frozen and we obtain a pure network formation dynamics with heterogeneous players. In this sense a co-evolutionary model can be regarded as an extension of an evolutionary model.

3 Asymptotic properties of the process

For positive levels of noise β and finite player population, the co-evolutionary process $\{X_N^{\beta,\boldsymbol{\tau}}(t)\}_{t\geq 0}$ defined in the previous section is ergodic. In Sect. 3.1 we give a complete characterization of its unique invariant distribution, and compare this measure to previous results in evolutionary game theory. Using this result, section 3.2 gives a complete characterization of the long-run interaction structure generated by the co-evolutionary process.

3.1 The invariant distribution

The main result of this section is the following Theorem.

Theorem 1 *The unique invariant distribution of the Markov jump process $\{X_N^{\beta,\boldsymbol{\tau}}(t)\}_{t\geq 0}$ is the Gibbs measure*

$$\mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{x}) = \frac{\exp(\beta^{-1}H_N^{\beta}(\mathbf{x},\boldsymbol{\tau}))}{\sum_{\mathbf{x}'\in\mathcal{X}^N}\exp(\beta^{-1}H_N^{\beta}(\mathbf{x}',\boldsymbol{\tau}))} = \frac{\mu_{0,N}^{\beta,\boldsymbol{\tau}}(\mathbf{x})\exp\left(\beta^{-1}V(\mathbf{x},\boldsymbol{\tau})\right)}{\sum_{\mathbf{x}'\in\mathcal{X}^N}\mu_{0,N}^{\beta,\boldsymbol{\tau}}(\mathbf{x}')\exp\left(\beta^{-1}V(\mathbf{x}',\boldsymbol{\tau})\right)}, \tag{5}$$

where, for all $\mathbf{x} = (\mathbf{a},\mathbf{g}) \in \mathcal{X}^N$,

$$\begin{aligned}
H_N^{\beta}(\mathbf{x},\boldsymbol{\tau}) &:= V(\mathbf{x},\boldsymbol{\tau}) + \beta\log\mu_{0,N}^{\beta,\boldsymbol{\tau}}(\mathbf{x}),\\
\mu_{0,N}^{\beta,\boldsymbol{\tau}}(\mathbf{x}) &:= \prod_{i=1}^{N}\prod_{j>i}\left(\frac{2}{N\xi_{\tau_i,\tau_j}^{\beta,N}}\right)^{g_{ij}},\\
V(\mathbf{x},\boldsymbol{\tau}) &:= \frac{1}{2}\sum_{j\neq i} g_{ij}v(a_i,a_j) + \sum_i \tau_i(a_i).
\end{aligned}$$

Proof It is easy to verify that the detailed balance conditions

$$\mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{x})\eta_N^{\beta,\boldsymbol{\tau}}(\mathbf{x},\mathbf{x}') = \mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{x}')\eta_N^{\beta,\boldsymbol{\tau}}(\mathbf{x}',\mathbf{x})$$

are satisfied for all $\mathbf{x},\mathbf{x}' \in \mathcal{X}^N$. See Theorem 4.1 and Proposition 4.1 in Staudigl (2011). □

The function H_N^{β}, called the graph Hamiltonian in Park and Newman (2004), provides a complete description of the invariant distribution weights. The structure of the invariant distribution is game–theoretically interesting, as it allows us to split the measure into two part: The weight at population state $\mathbf{x}$ is first influenced by graph weight function $\mu_{0,N}^{\beta,\boldsymbol{\tau}}$, which only depends on data concerning the network formation dynamics. Second we find the term $\exp[\beta^{-1}V(\mathbf{x},\boldsymbol{\tau})]$, which both depends on the game and the network structure prevailing at state $\mathbf{x}$. The function $V(\cdot,\boldsymbol{\tau})$ aggregates the individual players' utilities. In particular, it is a potential function for the normal form game $[U_i(\cdot, \mathbf{g}, \tau_i)]_{i\in[N]}$, where the players have learned their own types and the network is fixed at some graph $\mathbf{g} \in \mathcal{G}[N]$. More precisely, one can easily check that

$$V[(a, \mathbf{a}_{-i}), \mathbf{g}, \boldsymbol{\tau}] - V[(b, \mathbf{a}_{-i}), \mathbf{g}, \boldsymbol{\tau}] = U_i[(a, \mathbf{a}_{-i}), \mathbf{g}, \tau_i] - U_i[(b, \mathbf{a}_{-i}), \mathbf{g}, \tau_i]$$

holds for all $i \in [N]$, $a, b \in A$, $\mathbf{a}_{-i} \in A^{N-1}$, $\mathbf{g} \in \mathcal{G}[N]$ and $\boldsymbol{\tau} \in \Theta^N$. Hence, one can view the stationary distribution as a perturbation of the invariant measure found by Blume (1997) in the context of interaction games on fixed networks.

Remark 2 Since the types of the players are determined exogenously, they are independent from the co-evolutionary process of networks and play (but certainly shape the structure of the invariant distribution of the process). Consequently, one should interpret the invariant distribution $\mu_N^{\beta,\boldsymbol{\tau}}$ as a *random measure*, as it depends on the realized type profile. We will elaborate on this point further when we study the large population limit of the co-evolutionary process in Sect. 4.2.

3.2 The random graph

The probability measure $\mu_N^{\beta,\boldsymbol{\tau}}$ collects all information about the long-run behavior of the evolutionary process for a given type profile $\boldsymbol{\tau} \in \Theta^N$. Hence, it measures the probability that a network $\mathbf{g} \in \mathcal{G}[N]$ appears together with an action profile $\mathbf{a} \in A^N$. We now ask the question what graph topologies are most likely to be observed in the long run, when we *fix* the actions of the players (and their types). An answer to this question would give us information on all network topologies that may arise, when we have data on the behavior of the agents and their types. To answer this question we have to condition the process on an $\mathbf{a}$-section of the state space. An $\mathbf{a}$-section is the set $\mathcal{X}_{\mathbf{a}}^N := \{\mathbf{a}\} \times \mathcal{G}[N]$, where $\mathbf{a} = (a_1, \ldots, a_N) \in A^N$ is a given action profile. All population states in the $\mathbf{a}$-section differ only in the interaction network. Thus, if we constrain the process $\{X_N^{\beta,\boldsymbol{\tau}}(t)\}_{t\geq 0}$ to take values in this set only, we obtain a *random graph process* $G_N^{\beta,\boldsymbol{\tau}} = \{G_N^{\beta,\boldsymbol{\tau}}(t)\}_{t\geq 0}$ whose generator describes a birth-death process with "birth rates" of the link (i, j) given by $2\exp(v(a_i, a_j)/\beta)$, and "death-rates" $w_{ij}^{\beta,N}(\boldsymbol{\tau})$. It will be useful to introduce the rate ratio

$$\varphi_{kl}^{\beta,N}(a, b) = \frac{2\beta \exp(v(a, b)/\beta)}{\xi_{kl}^{\beta,N}},$$

for $a, b \in A$ and $1 \leq k, l \leq K$. The main result of this section is the following characterization theorem, categorizing the class of random graphs generated by the co-evolutionary process. The proof of the theorem is actually a corollary of a more general result, which we present in Appendix A.

Theorem 2 *Consider a random graph process $G_N^{\beta,\boldsymbol{\tau}}$ with attachment rates $\mathbf{C}^{\beta,N}$ and volatility rates $\Xi^{\beta,N}$, whose components are defined by Eqs. 2 and 3. This process is ergodic with unique invariant graph measure*

$$\mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{x}|\mathcal{X}_{\mathbf{a}}^N) = \prod_{i=1}^{N}\prod_{j>i} p_{ij}^{\beta,N}(\mathbf{a},\boldsymbol{\tau})^{g_{ij}} \left(1 - p_{ij}^{\beta,N}(\mathbf{a},\boldsymbol{\tau})\right)^{1-g_{ij}} \mathbb{1}_{\mathcal{X}_{\mathbf{a}}^N}(\mathbf{x}), \tag{6}$$

which gives rise to the random graph $\mathcal{G}[N, (p_{ij}^{\beta,N}(\mathbf{a},\boldsymbol{\tau}))_{j>i}]$ with interaction probabilities

$$p_{ij}^{\beta,N}(\mathbf{a},\boldsymbol{\tau}) = \frac{\varphi_{kl}^{\beta,N}(a,b)}{\varphi_{kl}^{\beta,N}(a,b) + N\beta} \quad \text{if } a_i = a, a_j = b, \tau_i = \theta_k, \tau_j = \theta_l \tag{7}$$

for all $i, j \in [N]$.

Theorem 2 allows us to define for each pair of players a conditional probability that they will be matched in the long run for a given action and type profile. In fact, the particular form of the interaction probabilities in Eq. 7, allows us to characterize the interaction model without conditioning on any **a**-section at all. The matrices

$$\mathbf{p}_{kl}^{\beta,N} := \left(p_{kl}^{\beta,N}(a,b)\right)_{(a,b)\in A^2}, \quad \mathbf{p}^{\beta,N} := \left(\mathbf{p}_{kl}^{\beta,N}\right)_{1\le l,k\le K}$$

with

$$p_{kl}^{\beta,N}(a,b) := \frac{\varphi_{kl}^{\beta,N}(a,b)}{\varphi_{kl}^{\beta,N}(a,b) + N\beta}$$

always give us a complete characterization of the random graph. This is an interesting results, as it couples a random graph model with the game theoretic model in a simple and transparent way. Observe that the random graph is essentially independent of the labels of the individual players. Only their types and their used actions identify an agent in the model.

Remark 3 Based on Theorem 2, one can study the large population properties of the random graph ensemble. We shall not pursue such an analysis explicitly in this paper, since it can be done by applying known results from random graph theory. We refer the reader to Bollobás et al. (2007) for a detailed study of this question.

4 Stochastic stability analysis

Having worked out the ergodic properties of the co-evolutionary process, we now investigate the structure of several limiting measures of the invariant distribution $\mu_N^{\beta,\boldsymbol{\tau}}$. We start with the conceptually simpler case concerning the small noise limit.

4.1 The small noise limit

For the small noise limit, we use the following notion of stochastic stability (cf. Sandholm 2010b)[11]:

Definition 1 A population state $\mathbf{x} \in \mathcal{X}^N$ is called *stochastically stable in the small noise limit* if

$$\lim_{\beta \to 0} \beta \log \mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{x}) = 0.$$

We have the following characterization of small noise stochastically stable states.

Theorem 3 *The family of invariant measures* $\{\mu_N^{\beta,\boldsymbol{\tau}}\}_{\beta>0}$ *satisfies a large deviations principle with rate function* $R(\mathbf{x},\boldsymbol{\tau}) := \max_{\mathbf{x}' \in \mathcal{X}^N} V(\mathbf{x}',\boldsymbol{\tau}) - V(\mathbf{x},\boldsymbol{\tau})$*, i.e.*

$$\lim_{\beta \to 0} \beta \log \mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{x}) = -R(\mathbf{x},\boldsymbol{\tau})$$

for all $\mathbf{x} \in \mathcal{X}^N$ *and* $\boldsymbol{\tau} \in \Theta^N$.

Proof Under our assumptions on the volatility mechanism it is true that

$$\lim_{\beta \to 0} \max_{\mathbf{x} \in \mathcal{X}^N} \beta |\log \mu_{0,N}^{\beta,\boldsymbol{\tau}}(\mathbf{x})| = 0.$$

Hence, the Hamiltonian function H_N^β converges uniformly to the game potential function as $\beta \to 0$. Thus,

$$\begin{aligned} -\lim_{\beta \to 0} \beta \log \mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{x}) &= \max_{\mathbf{x}' \in \mathcal{X}^N} \lim_{\beta \to 0} H_N^\beta(\mathbf{x}',\boldsymbol{\tau}) - \lim_{\beta \to 0} H_N^\beta(\mathbf{x},\boldsymbol{\tau}) \\ &= \max_{\mathbf{x}' \in \mathcal{X}^N} V(\mathbf{x}',\boldsymbol{\tau}) - V(\mathbf{x},\boldsymbol{\tau}) = R(\mathbf{x},\boldsymbol{\tau}) \end{aligned}$$

for all $\mathbf{x} \in \mathcal{X}^N$, completing the proof. □

An immediate corollary of this theorem is the following:

Corollary 1 *The set of stochastically stable states in the small noise limit is the set of population states which maximize the game potential function, i.e.*

$$\{\mathbf{x} \in \mathcal{X}^N | R(\mathbf{x},\boldsymbol{\tau}) = 0\} = \arg \max_{\mathbf{x} \in \mathcal{X}^N} V(\mathbf{x},\boldsymbol{\tau}).$$

[11] The small noise limit has been emphasized in the early literature on evolutionary equilibrium selection in the seminal papers by Kandori et al. (1993) and Young (1993). The differences between these papers and our definition of small noise stochastic stability is explained in detail in Chapter 12 of Sandholm (2010b).

This results has the pleasant property that stochastically stable states in the small noise limit are socially efficient states, taking the sum of utilities of the players as an indicator for welfare, for *every* realization of types. Further, small noise stochastically stable states give us a *joint* prediction of action profiles and networks, which are mutually consistent in the sense that the action profile played on the corresponding network must be a Nash equilibrium, given the types of the players.

4.2 The large population limit

In a large population, it is natural to consider situations where each individual player has a minor impact on the evolution of the population state. Hence, it is natural to focus on *aggregative properties* of the model. Thus, in this section of the paper we shift our interest from the micro-level of the process to macroeconomic phenomena. Therefore we have to introduce several aggregation operators, which allow us to transport the mass the the invariant distribution puts on subsets of the set of type-action *profiles* to the subsets of type-action *distributions*. Moreover, we are interested in the weak-convergence properties of this law as the number of players becomes large. As this is rather technical, no proofs will be given in this section. Instead we try to convey the main motivation and intuition for our results, referring the interested reader to Appendix B for all the mathematical details.

4.2.1 Aggregation operators and notation

Recall that in our co-evolutionary process, before the dynamics starts to unfold, each player randomly receives his private type τ_i, which is the realization of an i.i.d. random variable with law $\mathbf{q} \in \text{int}\,\Delta(\Theta)$. A realized type profile $\boldsymbol{\tau} \in \Theta^N$ defines an (empirical) distribution over types, defined as

$$M_k^N(\boldsymbol{\tau}) := \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\theta_k}(\tau_i) \quad \forall k \in \{1, 2, \ldots, K\}.$$

The components of the (random) distribution $\mathbf{M}^N(\boldsymbol{\tau}) := (M_1^N(\boldsymbol{\tau}), \ldots, M_K^N(\boldsymbol{\tau})) \in \Delta(\Theta)$ measure the frequency of players of type $k = 1, 2, \ldots, K$, when the type profile is $\boldsymbol{\tau} \in \Theta^N$. This empirical type distribution is a (random) element of the set $\mathcal{L}_N := \{\mathbf{M}^N(\boldsymbol{\tau}) | \boldsymbol{\tau} \in \Theta^N\}$. Conversely, we can also consider the set of type realizations which result in a targeted distribution $\mathbf{m} \in \Delta(\Theta)$. This is done by defining the *type class set* $\mathcal{T}^N(\mathbf{m}) := \{\boldsymbol{\tau} \in \Theta^N | \mathbf{M}^N(\boldsymbol{\tau}) = \mathbf{m}\}$. Hence, $\mathcal{T}^N(\mathbf{m}) \neq \emptyset$ if and only if $\mathbf{m} \in \mathcal{L}_N$. Given that each player receives his type independently and with the same law $\mathbf{q}$, the joint law of a profile $\boldsymbol{\tau} \in \mathcal{T}^N(\mathbf{m})$ is given by

$$\mathsf{P}_{\mathbf{q}}(\tilde{\boldsymbol{\tau}}_N = \boldsymbol{\tau}) = \prod_{k=1}^{K} q_k^{Nm_k}.$$

This probability is seen to be constant on a type class set $\mathcal{T}^N(\mathbf{m})$, which implies that

$$\mathsf{P}_{\mathbf{q}}(\tilde{\boldsymbol{\tau}}_N \in \mathcal{T}^N(\mathbf{m})) \equiv \mathsf{P}^N_{\mathbf{q}}(\mathbf{m}) := |\mathcal{T}^N(\mathbf{m})| \prod_{k=1}^{K} q_k^{Nm_k},$$

where $|\mathcal{T}^N(\mathbf{m})| = \frac{N!}{\prod_{k=1}^K (Nm_k)!}$ is the number of type profiles $\boldsymbol{\tau} \in \Theta^N$ which generate the type distribution $\mathbf{m} \in \mathcal{L}_N$.

Denote by $\Sigma := \Delta(A)^K$ the K-fold copy of the mixed strategy simplex, and $\mathcal{K} := \Sigma \times \Delta(\Theta)$ the space of *action-type distributions*. Finally, denote by $\Omega^N := A^N \times \Theta^N$ the set of action-type profiles. We now come to our definition of a Bayesian strategy.

Definition 2 A finite population *Bayesian strategy* is a distribution $\mathbf{S}^N = (S_1^N, \ldots, S_K^N) \in \Sigma$, where $S_k^N : \Omega^N \to \Delta(A)$ is the distribution over pure actions chosen in aggregate by the players of type θ_k, i.e.

$$S^N_{b,k}(\mathbf{a}, \boldsymbol{\tau}) := \frac{1}{NM_k^N(\boldsymbol{\tau})} \sum_{i=1}^{N} \mathbb{1}_{(b,\theta_k)}(a_i, \tau_i)$$

for all $(\mathbf{a}, \boldsymbol{\tau}) \in \Omega^N,\ b \in A, 1 \leq k \leq K$.

Every component $S_k^N = (S^N_{a,k})_{a \in A}$ is formally equivalent to a mixed strategy, the mixed strategy chosen by a "representative" player of type θ_k.[12]

Remark 4 Our definition of a Bayesian strategy is the appropriate finite-population version of the continuum population model of Ely and Sandholm (2005). Let us stress again that, although players belonging to the same population share the same preferences, their payoff in the game will in general not be the same, since their personal interaction network may be different. This is an important difference to the Bayesian population game model of Ely and Sandholm (2005), where players of the same type have the same (expected) payoff, and henceforth the same incentives in the game.

4.2.2 The law of action-type distributions

Once we have specified the common prior of the types, we can define an invariant distribution on the set $\mathcal{X}^N \times \Theta^N = A^N \times \mathcal{G}[N] \times \Theta^N$ as

$$\mu_N^{\beta}(\{(\mathbf{a}, \mathbf{g}, \boldsymbol{\tau})\}) := \mathsf{P}_{\mathbf{q}}(\tilde{\boldsymbol{\tau}}_N = \boldsymbol{\tau})\mu_N^{\beta,\boldsymbol{\tau}}(\{(\mathbf{a}, \mathbf{g})\}),$$

where $\mu_N^{\beta,\boldsymbol{\tau}}$ is the invariant measure over population states characterized in Theorem 1. In our large population investigations, we would like to extract from this measure information on the joint law of the random pair $(\mathbf{S}^N, \mathbf{M}^N)$ under the co-evolutionary process. In order to get this information, we need to transport the mass which the invariant distribution assigns to subsets of $\mathcal{X}^N \times \Theta^N$ to the space Ω^N, and then to the space $\mathcal{K}$. We proceed step-by-step by transporting first the mass from $\mathcal{X}^N \times \Theta^N$ to Ω^N. The following proposition is the first result along these lines.

[12] With this interpretation in mind, calling the empirical measure $\mathbf{S}^N$ a Bayesian strategy is justified by the same logic as the population-based interpretation of mixed strategies in evolutionary games.

Proposition 1 *For all* $(\mathbf{a}, \boldsymbol{\tau}) \in \Omega^N$, *define the conditional measure*

$$\nu_N^\beta(\mathbf{a}|\boldsymbol{\tau}) := \mu_N^{\beta,\boldsymbol{\tau}}(\mathcal{X}_{\mathbf{a}}^N).$$

Then we can compute

$$\nu_N^\beta(\mathbf{a}|\boldsymbol{\tau}) = \frac{1}{Z_N^\beta(\boldsymbol{\tau})} \prod_{k=1}^K \prod_{a=1}^n \Phi_k^a(\mathbf{S}^N(\mathbf{a}, \boldsymbol{\tau}), \beta, N)^{Nm_k S_{a,k}^N(\mathbf{a},\boldsymbol{\tau})} \tag{8}$$

where, for all types $1 \leq k < l \leq K$, *and actions* $1 \leq a \leq n$, *the function* $\Phi_k^a(\cdot, \beta, N) : \Sigma \to \mathbb{R}_+$ *is defined as*

$$\Phi_k^a(\boldsymbol{\sigma}, \beta, N) := \prod_{l \geq k} \Phi_{kl}^a(\boldsymbol{\sigma}, \beta, N),$$

$$\Phi_{kk}^a(\boldsymbol{\sigma}, \beta, N) := \exp\left(\frac{\theta_k(a)}{\beta}\right) \prod_{b \geq a} \left(1 + \frac{1}{N\beta}\varphi_{kk}^{\beta,N}(a, b)\right)^{\frac{Nm_k\sigma_k(b) - \delta_{a,b}}{1+\delta_{a,b}}},$$

$$\Phi_{kl}^a(\boldsymbol{\sigma}, \beta, N) := \prod_{b=1}^n \left(1 + \frac{1}{N\beta}\varphi_{kl}^{\beta,N}(a, b)\right)^{Nm_l\sigma_l(b)}.$$

The factor $Z_N^\beta(\boldsymbol{\tau})$ *is the normalization constant.*

Proof See, Appendix B.1. □

The measure $\nu_N^\beta(\cdot|\boldsymbol{\tau})$ is defined conditional on the type realization $\boldsymbol{\tau}$. Given this tuple, the mass $\nu_N^\beta(\mathbf{a}|\boldsymbol{\tau})$ is the mass that the stationary distribution $\mu_N^{\beta,\boldsymbol{\tau}}(\cdot)$ puts on all action-network configurations, where the action profile of the players is $\mathbf{a}$, but the network topologies vary. In other words it is the marginal distribution $\sum_{\mathbf{g} \in \mathcal{G}[N]} \mu_N^{\beta,\boldsymbol{\tau}}(\mathbf{a}, \mathbf{g})$.[13] A formal application of the rules of conditional probability gives us the *joint law* of the pair $(\mathbf{a}, \boldsymbol{\tau}) \in \Omega^N$ as

$$\nu_N^\beta(\mathbf{a}, \boldsymbol{\tau}) := \mathsf{P}_{\mathbf{q}}(\tilde{\boldsymbol{\tau}}_N = \boldsymbol{\tau})\nu_N^\beta(\mathbf{a}|\boldsymbol{\tau}) = \mu_N^\beta(\mathcal{X}_{\mathbf{a}}^N \times \{\boldsymbol{\tau}\}). \tag{9}$$

Instead of explicitly writing down this rather cumbersome expression, we introduce some new concepts, which will turn out to be very useful in the sequel. We define the *interaction potential functions*

[13] As an interpretation for this measure, think of the case with mean-field interactions (Horst and Scheinkman 2006), i.e. a situation where the utility functions of the agents only depend on the average behavior of all other players. Then a strategy revision process for a fixed type profile $\boldsymbol{\tau} \in \Theta^N$ with the logit dynamics (1) would produce an invariant distribution on A^N of the form of $\nu_N^\beta(\cdot|\boldsymbol{\tau})$ (see, Sandholm 2010b, Exercise 11.5.17). Proposition 1 goes beyond the case of mean-field interactions, since it results from an aggregation procedure where we aggregate over all possible network structures for a given profile of actions and types.

$$
\begin{aligned}
(1 \leq k \leq K): f^{\beta}_{N,k}(\boldsymbol{\sigma}, \mathbf{m}) &:= \sum_{a \in A} \sigma_k(a) \sum_{l \geq k} \log \Phi^a_{kl}(\boldsymbol{\sigma}, \beta, N), \\
f^{\beta}_N(\boldsymbol{\sigma}, \mathbf{m}) &:= \sum_{k=1}^{K} m_k f^{\beta}_{N,k}(\boldsymbol{\sigma}, \mathbf{m}).
\end{aligned} \tag{10}
$$

Using these mappings, it is a simple matter of rewriting terms in Eq. 9 to see that

$$\nu^{\beta}_N(\mathbf{a}, \boldsymbol{\tau}) \propto \mathsf{P}_{\mathbf{q}}(\tilde{\boldsymbol{\tau}}_N = \boldsymbol{\tau}) e^{N f^{\beta}_N(\mathbf{S}^N(\mathbf{a},\boldsymbol{\tau}), \mathbf{M}^N(\boldsymbol{\tau}))} \quad \forall (\mathbf{a}, \boldsymbol{\tau}) \in \Omega^N.$$

Now we have obtained an expression for the weight of an action-type profile $(\mathbf{a}, \boldsymbol{\tau})$, which actually only depends on the induced action-type *distribution*. Thus, all pairs of action-type profiles which give rise to the same action-type distribution, receive the same mass under the distribution ν^{β}_N. This simple observation allows us to derive a probability measure on the set of action-type distributions $\mathcal{K} = \Sigma \times \Delta(\Theta)$, as follows. Since the mapping $(\mathbf{S}^N, \mathbf{M}^N)$ maps action-type profiles in Ω^N to action-type distributions in $\mathcal{K}$, we can define the pre-image of a couple $(\boldsymbol{\sigma}, \mathbf{m}) \in \mathcal{K}$ as $(\mathbf{S}^N, \mathbf{M}^N)^{-1}(\boldsymbol{\sigma}, \mathbf{m}) := \{(\mathbf{a}, \boldsymbol{\tau}) \in \Omega^N | \mathbf{S}^N(\mathbf{a}, \boldsymbol{\tau}) = \boldsymbol{\sigma}, \mathbf{M}^N(\boldsymbol{\tau}) = \mathbf{m}\}$. An element of this set is called a *distributionally equivalent* action-type profile. Now, since all distributionally equivalent action-type profiles must receive the same mass under the measure ν^{β}_N, we need to count the number of elements in the set $(\mathbf{S}^N, \mathbf{M}^N)^{-1}(\boldsymbol{\sigma}, \mathbf{m})$. This is a straightforward exercise in combinatorics. By its very definition, the set $(\mathbf{S}^N, \mathbf{M}^N)^{-1}(\boldsymbol{\sigma}, \mathbf{m})$ is nonempty if and only if the type distribution is contained in the set $\mathcal{L}_N$, and $\boldsymbol{\sigma}$ is an element of the set $\Sigma^N(\mathbf{m}) = \times_{k=1}^{K} \Sigma^N_k(\mathbf{m})$, where each set $\Sigma^N_k(\mathbf{m})$ is defined as the discrete grid $\Delta(A) \cap \frac{1}{N m_k} \mathbb{Z}^n$ if $m_k > 0$, and $\Delta(A)$ otherwise. Consider a pair $(\boldsymbol{\sigma}, \mathbf{m})$ which satisfies these conditions. Then the array of numbers $\{\{N m_k \sigma_k(a)\}_{a \in A}\}_{1 \leq k \leq K}$ sums up to N. The number of distributionally equivalent action-type profiles is given by the number of all possible permutations of action profiles and type profiles which give rise to the distribution $(\boldsymbol{\sigma}, \mathbf{m})$. Elementary combinatorics tells us that the number of distributionally equivalent action-type profiles is given by

$$\frac{N!}{\prod_{k=1}^{K} \prod_{a=1}^{n} (N m_k \sigma_k(a))!} = \prod_{k=1}^{K} \frac{(N m_k)!}{\prod_{a \in A} (N m_k \sigma_k(a))!} \frac{N!}{\prod_{k=1}^{K} (N m_k)!}.$$

Hence, we can define a measure

$$\gamma^{\beta}_N(\{(\boldsymbol{\sigma}, \mathbf{m})\}) := \frac{N!}{\prod_{k=1}^{K} \prod_{a=1}^{n} (N m_k \sigma_k(a))!} \exp\left\{ N[f^{\beta}_N(\boldsymbol{\sigma}, \mathbf{m}) + \sum_{k=1}^{K} m_k \log(q_k)] \right\}$$

for all couples $(\boldsymbol{\sigma}, \mathbf{m})$ in range of the random pair $(\mathbf{S}^N, \mathbf{M}^N)$.[14] The interpretation of this measure is given by the following lemma, whose straightforward proof we omit.

[14] Thus, this measure is a counting measure on $\mathcal{K}$, which puts mass $\gamma^{\beta}_N(\{(\boldsymbol{\sigma}, \mathbf{m})\})$ on the individual points in the range of $(\mathbf{S}^N, \mathbf{M}^N)$. See, Appendix B.1 for the details.

Lemma 1 *For all couples* $(\boldsymbol{\sigma}, \mathbf{m}) \in \mathcal{K}$ *we have*

$$\gamma_N^\beta(\{(\boldsymbol{\sigma}, \mathbf{m})\}) = \mathsf{P}_{\mathbf{q}}^N(\mathbf{m}) e^{N f_N^\beta(\boldsymbol{\sigma}, \mathbf{m})}. \tag{11}$$

Note that $\gamma_N^\beta(\{(\boldsymbol{\sigma}, \mathbf{m})\}) > 0$ if and only if $\mathbf{m} \in \mathcal{L}_N$ and $\boldsymbol{\sigma}$ is an action distribution such that $N m_k \sigma_k(a)$ is an integer for all $1 \leq k \leq K, 1 \leq a \leq n$. Thus, only those pairs $(\boldsymbol{\sigma}, \mathbf{m})$ have positive mass, which are in the range of the pair $(\mathbf{S}^N, \mathbf{M}^N)$. Normalizing the finite measure (11) leaves us with a probability measure on the set $\mathcal{K}$, defined as

$$P_N^\beta(\Gamma) := \frac{\gamma_N^\beta(\Gamma)}{\gamma_N^\beta(\mathcal{K})}, \quad \forall \Gamma \subset \mathcal{K}. \tag{12}$$

This probability measure is a joint law on the set of action-type distributions $\mathcal{K}$. A law for Bayesian strategies, for a given type distribution $\mathbf{m} \in \mathcal{L}_N$ is then obtained from the conditional distribution induced by P_N^β on the set $\Sigma^N(\mathbf{m})$, which is characterized by the probability mass function

$$\psi_N^\beta(\boldsymbol{\sigma}|\mathbf{m}) = \frac{e^{N f_N^\beta(\boldsymbol{\sigma}, \mathbf{m})}}{\sum_{\boldsymbol{\sigma}' \in \Sigma^N(\mathbf{m})} e^{N f_N^\beta(\boldsymbol{\sigma}', \mathbf{m})}}.$$

4.2.3 The large deviations principle

In this section we prove that the sequence of laws $\{P_N^\beta\}_{N \in \mathbb{N}}$, under some mild assumptions on the data of the co-evolutionary process, satisfies a large deviations principle (LDP) in the limit of large player populations. From the representation of the probability measure P_N^β via the measures γ_N^β, it is clear that convergence properties of the interaction potential functions will play a decisive role here. It turns out that, in order to establish such a convergence, we need to make some assumptions on the population dependence of the volatility rates of the network formation dynamics. Specifically, we require the following.

Theorem 1 The array of volatility rates satisfies *large population boundedness*:

$$\lim_{N \to \infty} \xi_{kl}^{\beta, N} = \xi_{kl}^\beta \in (0, \infty) \quad 1 \leq k, l \leq K, \tag{LPB}$$

Under this assumption, we can prove that the sequence of interaction potential functions $\{\{f_k^{\beta, N}\}_{k=1}^K\}_{N \geq 1}$ converges to a well-defined sequence of (continuous) limit functions $\{f_k^\beta\}_{k=1}^K$.

Lemma 2 *Assume that the volatility rates satisfy (LPB). Then, for all* $1 \leq k \leq K$ *and along any sequence* $\left\{(\boldsymbol{\sigma}^N, \mathbf{m}^N)\right\}_{N \geq 1}$, $\boldsymbol{\sigma}^N \in \Sigma^N(\mathbf{m}^N), \mathbf{m}^N \in \mathcal{L}_N$, *with limit* $(\boldsymbol{\sigma}, \mathbf{m}) \in \Sigma \times \operatorname{int} \Delta(\Theta)$, *we have*

$$\lim_{N \to \infty} f_k^{\beta, N}(\boldsymbol{\sigma}^N, \mathbf{m}^N) = \frac{1}{\beta} f_k^\beta(\boldsymbol{\sigma}, \mathbf{m}),$$

where $f_k^\beta : \Sigma \times \Delta(\Theta) \to \mathbb{R}$ *is the continuous function*

$$f_k^\beta(\boldsymbol{\sigma}, \mathbf{m}) := \langle \boldsymbol{\sigma}_k, \boldsymbol{\theta}_k \rangle + \sum_{l \geq k} \frac{m_l}{1+\delta_{kl}} \left\langle \boldsymbol{\sigma}_k, \boldsymbol{\varphi}_{kl}^\beta \boldsymbol{\sigma}_l \right\rangle.$$

Proof See, Appendix B.2. □

Given this convergence result, we define

$$f^\beta(\boldsymbol{\sigma}, \mathbf{m}) := \sum_{k=1}^K m_k f_k^\beta(\boldsymbol{\sigma}, \mathbf{m}) \quad \forall (\boldsymbol{\sigma}, \mathbf{m}) \in \Sigma \times \Delta(\Theta).$$

For a full statement of the the large deviations principle, we still need one more piece of notation. Given that the measure γ_N^β involves a combinatorial factor which depends on the number of players, we will need to take care of the large N properties of this combinatorial term. To do this, we will need to introduce the (relative) entropy of a discrete probability measure. Given a type distribution $\mathbf{m}$ we define the *relative entropy* of this distribution with respect to the prior $\mathbf{q}$ as

$$h(\mathbf{m}||\mathbf{q}) := \sum_{k=1}^K m_k \log \frac{m_k}{q_k}.$$

Similarly, and with an abuse of notation, we define the entropy of an action distribution $\boldsymbol{\sigma}_k \in \Delta(A)$ as

$$h(\boldsymbol{\sigma}_k) := -\sum_{a \in A} \sigma_k(a) \log(\sigma_k(a)).$$

These functions will show up in the large deviations estimate combined with the limiting interaction potential function f^β in form of the *perturbed interaction potential function*

$$\tilde{f}^\beta(\boldsymbol{\sigma}, \mathbf{m}) := f^\beta(\boldsymbol{\sigma}, \mathbf{m}) + \beta \sum_{k=1}^K m_k h(\boldsymbol{\sigma}_k). \tag{13}$$

Remark 5 Perturbed potential functions of the form (13) play an important role in models of noisy evolution in games. Hofbauer and Sandholm (2007) identify functions of the form (13) as strict Lyapunov functions for an appropriately defined "mean-field" dynamics, arising in their model of stochastic evolution in potential games (see in particular their Theorem 3.2). Our analysis shows that perturbed potential functions play an important role in understanding the long-run properties of the co-evolutionary process in the limit of large populations.

The next Theorem uses the perturbed interaction potential function (13) to establish the large deviations principle for the measures P_N^β.

Theorem 4 *The family of measures* $\{P_N^\beta\}_{N\in\mathbb{N}}$ *satisfies a large deviations principle with speed* N *and rate function*

$$R_{\mathbf{q}}^\beta(\boldsymbol{\sigma}, \mathbf{m}) := \max_{(\boldsymbol{\sigma}', \mathbf{m}')\in\mathcal{K}} \left[\frac{1}{\beta}\tilde{f}^\beta(\boldsymbol{\sigma}', \mathbf{m}') - h(\mathbf{m}'||\mathbf{q})\right] - \left[\frac{1}{\beta}\tilde{f}^\beta(\boldsymbol{\sigma}, \mathbf{m}) - h(\mathbf{m}||\mathbf{q})\right]. \tag{14}$$

Hence, we have the large-deviations lower bound

$$\liminf_{N\to\infty} \frac{1}{N} \log P_N^\beta(G) \geq -\inf_{(\boldsymbol{\sigma},\mathbf{m})\in G} R_{\mathbf{q}}^\beta(\boldsymbol{\sigma}, \mathbf{m})$$

for all open $G \subset \mathcal{K}$*, and the large deviations upper bound*

$$\limsup_{N\to\infty} \frac{1}{N} \log P_N^\beta(F) \leq -\inf_{(\boldsymbol{\sigma},\mathbf{m})\in F} R_{\mathbf{q}}^\beta(\boldsymbol{\sigma}, \mathbf{m})$$

for all closed $F \subset \mathcal{K}$.

The proof of this Theorem is given in Appendix B.3. This result shows how one can estimate the probability that the generated action-type distribution falls into a certain subset $\Gamma \subset \mathcal{K}$, by using the rate function (14). Writing $\asymp$ for asymptotic equivalence on an exponential scale, the content of Theorem 4 is that

$$P_N^\beta(\Gamma) \asymp \exp\left(-N \inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma} R_{\mathbf{q}}^\beta(\boldsymbol{\sigma}, \mathbf{m})\right)$$

for all $\Gamma \subseteq \mathcal{K}$. Hence, action-type distributions which are close to the global optimum of the function $(\boldsymbol{\sigma}, \mathbf{m}) \mapsto \frac{1}{\beta}\tilde{f}^\beta(\boldsymbol{\sigma}, \mathbf{m}) - h(\mathbf{m}||\mathbf{q})$ are going to be observed with non-vanishing probability on an exponential scale. This relation between this concentration tendency of the sequence of laws P_N^β and the rate function $R_{\mathbf{q}}^\beta$ will be exploited in the next section, to relate the large-deviations principle to a well-known game theoretic concept.

4.2.4 Logit equilibria

The strong law of large numbers implies that $\mathbf{M}^N \to \mathbf{q}$ almost surely. Thus, let us specialize the setting in this section by assuming right away that $\mathbf{m} = \mathbf{q}$ is fixed deterministically. This allows us to get rid of the relative entropy function appearing in the rate function (14), since $h(\mathbf{q}||\mathbf{q}) = 0$. Under this assumption, the Bayesian strategies which describe the most likely distribution of play in the population (aggregated over all networks) are the maximizers of the perturbed interaction potential function $\tilde{f}^\beta$. We are therefore naturally led to investigate the structure of the set of solutions of the optimization problem

$$\max_{\boldsymbol{\sigma}'\in\Sigma} \tilde{f}^\beta(\boldsymbol{\sigma}', \mathbf{q}).$$

It is well known (see for instance Fudenberg and Levine 1998) that solutions of this program are *logit equilibria*, i.e. Bayesian strategies which are defined by the fixed-point condition

$$\sigma_k^*(a) = \frac{\exp\left(\beta^{-1}(\pi_a^k(\boldsymbol{\sigma}^*, \mathbf{q}) + \theta_k(a))\right)}{\sum_{b\in A}\exp\left(\beta^{-1}(\pi_b^k(\boldsymbol{\sigma}^*, \mathbf{q}) + \theta_k(b))\right)},$$

where

$$\pi_a^k(\boldsymbol{\sigma}, \mathbf{q}) := \sum_{l=1}^{K} q_l \sum_{b\in A} \varphi_{kl}^{\beta}(a, b)\sigma_l(b) \equiv \sum_{l=1}^{K} q_l \left(\boldsymbol{\varphi}_{kl}^{\beta}\boldsymbol{\sigma}_l\right)_a$$

for all $a \in A$ and $1 \leq k \leq K$. To show this, introduce the Lagrangian

$$L(\boldsymbol{\sigma}, \boldsymbol{\lambda}) := \tilde{f}^{\beta}(\boldsymbol{\sigma}, \mathbf{q}) - \sum_{k=1}^{K} \lambda_k \left(\sum_{a\in A} \sigma_k(a) - 1\right),$$

where $\boldsymbol{\lambda} := (\lambda_1, \ldots, \lambda_K)$ are the Lagrangian multipliers corresponding to the state-space constraints. First order conditions will give necessary and sufficient conditions for an optimum, which will be interior and unique for β sufficiently large and positive. Straightforward algebra shows that the the first-order conditions are

$$\frac{\partial \tilde{f}^{\beta}(\boldsymbol{\sigma}, \mathbf{q})}{\partial \sigma_k(a)} - \frac{\partial \tilde{f}^{\beta}(\boldsymbol{\sigma}, \mathbf{q})}{\partial \sigma_k(b)} = 0$$

for all $1 \leq k \leq K$ and $a, b \in A$. By symmetry of the matrices $\boldsymbol{\varphi}_{kl}^{\beta}$, we have

$$\frac{\partial f_l^{\beta}(\boldsymbol{\sigma}, \mathbf{q})}{\partial \sigma_k(a)} = \begin{cases} 0 & \text{if } l > k, \\ \theta_k(a) + \sum_{l' \geq k} q_{l'} \left(\boldsymbol{\varphi}_{kl'}^{\beta}\boldsymbol{\sigma}_{l'}\right)_a & \text{if } l = k, \\ q_k \left(\boldsymbol{\varphi}_{kl}^{\beta}\boldsymbol{\sigma}_l\right)_a & \text{if } l < k. \end{cases}$$

From this we obtain

$$\frac{\partial \tilde{f}^{\beta}(\boldsymbol{\sigma}, \mathbf{q})}{\partial \sigma_k(a)} = q_k \left[\theta_k(a) + \sum_{l=1}^{K} q_l \left(\boldsymbol{\varphi}_{kl}^{\beta}\boldsymbol{\sigma}_l\right)_a - \beta(\log \sigma_k(a) + 1)\right].$$

Using this expression for the first-order conditions shows that in an optimum we need that

$$\begin{aligned} \log\frac{\sigma_k(a)}{\sigma_k(b)} &= \frac{1}{\beta}\left[\sum_{l=1}^{K} q_l \left(\boldsymbol{\varphi}_{kl}^{\beta}\boldsymbol{\sigma}_l + \theta_k\right)_a - \sum_{l=1}^{K} q_l \left(\boldsymbol{\varphi}_{kl}^{\beta}\boldsymbol{\sigma}_l + \theta_k\right)_b\right] \\ &= \beta^{-1}\left[(\pi_a^k(\boldsymbol{\sigma}, \mathbf{q}) + \theta_k(a)) - (\pi_b^k(\boldsymbol{\sigma}, \mathbf{q}) + \theta_k(b))\right]. \end{aligned}$$

From this odds-ratio condition the rest follows immediately by using the constraint $\sum_{a\in A}\sigma_k(a)=1$. To interpret this result, note that $\pi_a^k(\boldsymbol{\sigma},\mathbf{q})$ has actually a quite interesting interpretation. In a dynamic choice model without network formation dynamics, a logit equilibrium would correspond to a perturbed Nash equilibrium of the underlying normal form game. Hence, instead of π_a^k the logit equilibrium of a normal form game would have to take the pure-strategy payoff function of the game into account. Viewed from this angle, the network formation dynamic acts on the normal form game *as if* the payoffs are (non-linearly) transformed, so that the (ex-post) payoff matrix of a player of type θ_k meeting a player of type θ_l is $\boldsymbol{\varphi}_{kl}^{\beta}$.

5 Conclusion

This paper presents an analytically solvable model on the co-evolution of networks and play in settings where players have diverse preferences. We restrict the class of games by assuming that players' preferences can be additively decomposed into a common utility term and a random idiosyncratic payoff term. If the common payoff function is a game of common interest and the players use log–linear functions in the action choice and linking choice, we give a closed-form solution of the (unique) invariant distribution of the process. This in turn allows us to perform a stochastic stability analysis in the limits of small noise and large populations, respectively. Many results presented in the paper hinge on the specific assumptions made in order to proceed with analytical methods. However, some qualitative features of the model are quite robust to changes in the model setting. Among these is the creation of inhomogeneous random graphs, as shown in Staudigl (2010b). It remains an open problem to work out the fine details of a co-evolutionary model, outside the world of exact potential games. This is, however, a general problem of stochastic evolutionary dynamics, where little is known about the exact long-run behavior of the dynamics once no closed-form solution of the invariant distribution is available. Extending our knowledge about co-evolutionary dynamics is a challenging task for future research.

Acknowledgements This paper is a largely revised version of the working paper Staudigl (2010a), and uses results from chapters of my dissertation Staudigl (2010b). I would like to thank Immanuel Bomze and Manfred Nermuth for their constant encouragement and support. I have also benefited a lot from insightful conversations with Fernando Vega-Redondo, and I am deeply indebted to William H. Sandholm for giving me very detailed feedback on earlier drafts of this paper. Finally, I would like to express my gratefulness to an anonymous referee and an associate editor for very useful comments and suggestions, which improved this paper a lot. Financial support form the Vienna Science and Technology Fund (WWTF) under project fund MA 09-017 is gratefully acknowledged.

Appendix A: The random graph model

Theorem 2 is a corollary of the following, more general, characterization theorem.

Theorem 5 *Consider a random graph process* $\{G(t)\}_{t\geq 0}$ *on* $\mathcal{G}[N]$ *with symmetric link creation rates* $\mathbf{C}=(c_{ij})_{(i,j)\in[N]^2}$ *and symmetric link destruction rates* $\mathbf{W}=$

$(w_{ij})_{(i,j)\in[N]^2}$, *where* **C** *and* **W** *are positive matrices. The generator of the process is defined as*

$$\eta_{g,g\oplus(i,j)} = (1-g_{ij})c_{ij},\ \eta_{g,g\ominus(i,j)} = g_{ij}w_{ij}.$$

Define $p_{ij} := \frac{c_{ij}}{c_{ij}+w_{ij}}$ *for all* $i, j \in [N]$. *Then the unique invariant distribution of this process is given by*

$$\mu(g) = \prod_{i,j>i} (p_{ij})^{g_{ij}} (1-p_{ij})^{1-g_{ij}}$$

Proof By positivity of **C** and **W** the graph-valued Markov process $\{G(t)\}_{t\geq 0}$ is ergodic. Since these matrices are also symmetric the Markov process is reversible in equilibrium. Consider the detailed balance conditions

$$\mu(g)\eta_{g,g\oplus(i,j)} = \mu(g\oplus(i,j))\eta_{g\oplus(i,j),g}. \tag{15}$$

for all $g \in \mathcal{G}[N]$ and $i, j \in [N]$, $j \neq i$. By force of normalization, given by a constant Z, this system of equations has a unique solution

$$\mu(g) = Z^{-1}\prod_{i=1}^{N}\prod_{j>i}\left(\frac{c_{ij}}{w_{ij}}\right)^{g_{ij}}. \tag{16}$$

Define for all $i = 1, 2, \dots, N$ and $j > i$ the numbers $x_{ij} := \log\left(\frac{c_{ij}}{w_{ij}}\right)$, and the function

$$H_0(g) := \sum_{i=1}^{N}\sum_{j>i} x_{ij}^{g_{ij}}.$$

Direct substitution into Eq. 16 gives the alternative representation of the invariant distribution as

$$\mu(g) = \frac{\exp(H_0(g))}{\sum_{g'\in\mathcal{G}[N]}\exp(H_0(g'))}. \tag{17}$$

We can compute the numerator of Eq. 17 as

$$\begin{aligned}\sum_{g\in\mathcal{G}[N]}\exp(H_0(g)) &= \sum_{g\in\mathcal{G}[N]}\prod_{i=1}^{N}\prod_{j>i}\exp(x_{ij}^{g_{ij}}) = \prod_{i=1}^{N}\prod_{j>i}(1+\exp(x_{ij}))\\ &= \prod_{i=1}^{N}\prod_{j>i}\left(1+\frac{c_{ij}}{w_{ij}}\right)\\ &= \prod_{i=1}^{N}\prod_{j>i}\left(1-p_{ij}\right)^{-1}.\end{aligned}$$

Further, for all $g \in \mathcal{G}[N]$ we have

$$\exp(H_0(g)) = \prod_{i=1}^{N} \prod_{j>i} \left(\frac{p_{ij}}{1 - p_{ij}} \right)^{g_{ij}}.$$

Combining these last two observations, we obtain the desired product measure

$$\mu(g) = \prod_{i=1}^{N} \prod_{j>i} \left(p_{ij}\right)^{g_{ij}} \left(1 - p_{ij}\right)^{1-g_{ij}}.$$

□

Proof of theorem 2 On $\mathbf{x} \in \mathcal{X}_{\mathbf{a}}^N$ the actions of the players are frozen. If $a_i = a$ and $a_j = b$ on $\mathbf{x}$, we set $c_{ij} = \frac{2}{N} \exp(v(a, b)/\beta)$ and $w_{ij} = \xi^{\beta,N}_{\tau_i,\tau_j}$. All the conditions for Theorem 5 are clearly satisfied, and the proof is complete. □

Appendix B: Aggregation and the large population limit

In this appendix we give a rigorous analysis of the large population behavior of the probability measure over Bayesian strategies. The analysis proceeds in three major steps.

B.1 Proof of Proposition 1

For aggregation purposes it is useful to have a partition of the set of players at hand, that categorizes them according to their action and their type. Let us define the sets

$$I_k^{\boldsymbol{\tau}}(a)(\mathbf{a}, \mathbf{g}) := \{i \in [N] | a_i = a \ \& \ \tau_i = \theta_k\}$$

for all $1 \leq k \leq K$ and $a \in A$. This is the set of players, who play action a at state $\mathbf{x}$ and are of type θ_k. Clearly, for a given type profile $\boldsymbol{\tau} \in \Theta^N$ the family of sets $\left\{\{I_k^{\boldsymbol{\tau}}(a)\}_{a \in A}\right\}_{k=1}^K$ is a partition on $[N]$. We call this collection of sets the *action-type-*partition of the population. The random graph measure (6) treats all edges between players $i \in I_k^{\boldsymbol{\tau}}(a)$ and $j \in I_l^{\boldsymbol{\tau}}(b)$ as i.i.d. random variables. Therefore, we can define a Binomially distributed random variable (with parameter $p_{kl}^{\beta,N}(a, b)$)

$$\mathcal{E}_{kl}^{N,\boldsymbol{\tau}}(a, b)(\mathbf{x}) := \frac{1}{1 + \delta_{ab}\delta_{lk}} \sum_{i \in I_k^{\boldsymbol{\tau}}(a)} \sum_{j \in I_l^{\boldsymbol{\tau}}(b)} g_{ij}.$$

Given a type profile $\boldsymbol{\tau}$ and an $\mathbf{a}$-section $\mathcal{X}_{\mathbf{a}}^N$ we denote by $E_{kl}^{N,\boldsymbol{\tau}}(a, b)$ the maximal number of edges that can be formed between agents of type k who play action a and agents of type l who play action b. $e_{kl}(a, b)$ denotes a particular realization of the random variable $\mathcal{E}_{kl}^{N,\boldsymbol{\tau}}(a, b)(\cdot)$.

Proof of Proposition 1 For notational simplicity let us drop the dependence of $\beta, \boldsymbol{\tau}$ and N from the involved functions and distributions, whenever no confusion can arise. Let us denote the absolute number of a-players of type θ_k as $z_k(a) := Nm_k\sigma_k(a)$. Item (i) of the Proposition is obvious. We therefore turn immediately to part (ii). Define for all $1 \leq k, l \leq K$ and $a, b \in A$

$$x_{kl}(a, b) := \frac{1}{\beta}v(a, b) + \log\left(\frac{2}{N\xi_{kl}}\right),$$

and

$$\rho(\mathbf{x}, \boldsymbol{\tau}) = \mu_0(\mathbf{x})\exp(V(\mathbf{x}, \boldsymbol{\tau})/\beta).$$

For all $\mathbf{x} \in \mathcal{X}_{\mathbf{a}}^N$ the action-type partition is fixed by definition, and therefore $I_k^{\boldsymbol{\tau}}(a)(\mathbf{x}) = I_k^{\boldsymbol{\tau}}(a)$ for all $1 \leq k \leq K, a \in A$ and $\mathbf{x} \in \mathcal{X}_{\mathbf{a}}^N$. Thus, we can write the function $\rho(\mathbf{x}, \boldsymbol{\tau})$ as[15]

$$\begin{aligned}\rho(\mathbf{x}, \boldsymbol{\tau}) = &\prod_{k=1}^{K}\prod_{a=1}^{n}\exp\left(\frac{\theta_k(a)z_k(a)}{\beta}\right)\prod_{b\geq a}\exp\left[x_{kk}(a, b)\right]^{\mathcal{E}_{kk}(a,b)(\mathbf{x})}\\ &\times\prod_{k,l>k}\prod_{a,b\in A}\exp\left[x_{kl}(a, b)\right]^{\mathcal{E}_{kl}(a,b)(\mathbf{x})}\end{aligned} \tag{18}$$

which is seen only to depend on the population state via the number of edges the network at $\mathbf{x}$ has. Now we aggregate this expression over all states $\mathbf{x} \in \mathcal{X}_{\mathbf{a}}^N$. This requires integrating over all possible edges that connect players playing a specific action and being of a specific type. The integration procedure can be performed iteratively by running the following algorithm:

Initialization: Set $k = 1$ and $a = 1$.
Loop 1: Consider $l = k$. Integrate over all possible edges $e_{kl}(a, b)$ for $b \geq a$. If $b = n$ set $l \to l + 1$ and go to Loop 2.
Loop 2: Integrate over possible edges $e_{kl}(a, b)$ for $b \in A$. If $l \leq K - 1$ set $l \to l + 1$ and repeat this procedure; otherwise go to Loop 3.
Loop 3: If $a \leq n - 1$ and $k \leq K - 1$ go to Loop 1 with the same k and $a \to a + 1$. If $a = n$ and $k \leq K - 1$ go to Loop 1 with $k \to k + 1$ and $a \to 1$. If $a = n$ and $k = K$ STOP.

To illustrate what this algorithm does we present the result after the initialization step and Loop 1 has been executed. Loop 1 starts with integrating over all possible edges connecting agents belonging to action class $I_1^{\boldsymbol{\tau}}(1)$ with itself. To perform this summation exercise, note that the only factor affected by the aggregation is $\exp[x_{11}(1, 1)]^{\mathcal{E}_{11}(1,1)(\mathbf{x})}$, $\mathbf{x} \in \mathcal{X}_{\mathbf{a}}^N$. Hence, if we collect terms unaffected by the aggregation under the placeholder B_1, we see that $\rho(\mathbf{x}, \boldsymbol{\tau}) = B_1\exp[x_{11}(1, 1]^{\mathcal{E}_{11}(1,1)(\mathbf{x})}$.

[15] The notation $\prod_{k,l>k}$ should be read as $\prod_{k=1}^{K}\prod_{l>k}$.

Next, we have to take care of combinatorial identities since there are many possibilities to connect agents in the respective action classes in order to produce the event $\{\mathcal{E}_{11}(1,1) = e_{11}(1,1)\}$. Adjusting for this we see that the output of the algorithm after the first round is

$$B_1 \sum_{e_{11}(1,1)=0}^{E_{11}(1,1)} \binom{E_{11}(1,1)}{e_{11}(1,1)} \exp(x_{11}(1,1))^{e_{11}(1,1)} = B_1\,(1+\exp(x_{11}(1,1)))^{E_{11}(1,1)}$$
$$= B_1 \left(1 + \tfrac{1}{N\beta}\varphi_{11}(1,1)\right)^{\frac{z_1(1)(z_1(1)-1)}{2}}$$

where $z_k(a) = Nm_k\sigma_k(a)$ for $1 \le k \le K$ and $a \in A$. The next step performed by the algorithm inside Loop 1 will be to sum over all possible connections between players in the action cells $I_1^{\tau}(1)$ and $I_1^{\tau}(2)$. Therefore we have to take the relevant factor out of the placeholder B_1 and perform the integral as above. This gives the intermediate result

$$B_2 \left(1 + \frac{1}{N\beta}\varphi_{11}(1,1)\right)^{\frac{z_1(1)(z_1(1)-1)}{2}} \left(1 + \frac{1}{N\beta}\varphi_{11}(1,2)\right)^{z_1(1)z_1(2)}.$$

Repeating this, as prescribed by the algorithm, we obtain after n steps the function

$$\Phi_{11}^1(\boldsymbol{\sigma},\beta,N)^{z_1(1)} = \exp\left(\frac{\theta_1(1)z_1(1)}{\beta}\right) \prod_{b\ge 1} \left(1 + \frac{1}{N\beta}\varphi_{11}(1,b)\right)^{\frac{z_1(1)(z_1(b)-\delta_{1,b})}{1+\delta_{1,b}}}.$$

Recalling that $z_k(a) = Nm_k\sigma_k(a)$, we see that this agrees with the definition of the function $\Phi_{11}^1(\boldsymbol{\sigma},\beta,N)$ in the text of the Proposition. Executing the remaining steps of the algorithm gives the desired result. □

B.2 Some large N approximation results

Bayesian strategies are defined as empirical distributions over actions used by the players of the several types and given the realized type distribution $\mathbf{m} \in \mathcal{L}_N$. Recall that, for any realized type distribution $\mathbf{m} \in \mathcal{L}_N$, the space of Bayesian strategies is the finite set $\Sigma^N(\mathbf{m}) = \times_{k=1}^K \Sigma_k^N(\mathbf{m})$, where $\Sigma_k^N(\mathbf{m}) = \Delta(A) \cap \frac{1}{Nm_k}\mathbb{Z}^n$ if $m_k > 0$, and $\Delta(A)$ otherwise. Consequently, calling $\mathcal{K} := \Sigma \times \Delta(\Theta)$, the finite set $\mathcal{K}^N := \{(\boldsymbol{\sigma},\mathbf{m}) \in \mathcal{K} | \boldsymbol{\sigma} \in \Sigma^N(\mathbf{m}), \mathbf{m} \in \mathcal{L}_N\}$ is the set of action-type statistics which are in the range of the random pair $(\mathbf{S}^N, \mathbf{M}^N)$. It is clear that any sequence $(\boldsymbol{\sigma}^N, \mathbf{m}^N) \in \mathcal{K}^N$ possesses a convergent subsequence with limit in $\mathcal{K}$. The following Lemma provides the converse to this result, by showing that we can arbitrarily well approximate any given pair $(\boldsymbol{\sigma},\mathbf{m}) \in \mathcal{K}$ with a sequence $\{(\boldsymbol{\sigma}^N, \mathbf{m}^N)\}_{N\in\mathbb{N}}$, chosen such that $\boldsymbol{\sigma}^N \in \Sigma^N(\mathbf{m}^N)$ for all $N \in \mathbb{N}$.

Lemma 3 *For every Bayesian strategy $\boldsymbol{\sigma} \in \Sigma$ and type distribution $\mathbf{m} \in \operatorname{int}\Delta(\Theta)$ there exists a sequence $\{(\boldsymbol{\sigma}^N, \mathbf{m}^N)\}_{N\in\mathbb{N}}$, with $\boldsymbol{\sigma}^N \in \Sigma^N(\mathbf{m}^N)$ and $\mathbf{m}^N \in \mathcal{L}_N$ for all N, such that $(\boldsymbol{\sigma}^N, \mathbf{m}^N) \to (\boldsymbol{\sigma},\mathbf{m})$ as $N \to \infty$.*

Proof The proof proceeds in two steps. First we show that we can find a sequence $\mathbf{m}^N \in \mathcal{L}_N$ that converges to $\mathbf{m}$ in total variation distance as $N \to \infty$. Then we use this sequence to construct the sequence of Bayesian strategies.

(i) On $\Delta(\Theta)$ define the total variation distance between two distributions $\mathbf{x}, \mathbf{y} \in \Delta(\Theta)$ as[16]

$$||\mathbf{x} - \mathbf{y}||_{TV,\Theta} := \frac{1}{2}\sum_{k=1}^{K} |x_k - y_k|,$$

If $\mathbf{m}^N \in \mathcal{L}_N$ then each coordinate $m_k^N \in \{0, \frac{1}{N}, \ldots, \frac{N}{N}\}$. Thus, if $\mathbf{m} \in \Delta(\Theta)$ then for every $1 \leq k \leq K$ there is a $m_k^N \in \{0, \frac{1}{N}, \ldots, \frac{N}{N}\}$ such that $|m_k - m_k^N| \leq \frac{1}{N}$. Thus, for every N we find a vector $\mathbf{m}^N$ such that $||\mathbf{m}^N - \mathbf{m}||_{TV,\Theta} \leq \frac{K}{2N}$. Hence $\mathbf{m}^N \to \mathbf{m}$ in total variation distance as $N \to \infty$.

(ii) Given the sequence of empirical type distribution $(m^N)_{N \geq N_0}$ identified in item (i), let $\boldsymbol{\sigma}^N \in \Sigma^N(\mathbf{m}^N)$ for all $N \geq N_0$. On the product space Σ we measure distance via the maximum-norm, that is

$$||\boldsymbol{\sigma} - \boldsymbol{\sigma}'||_{TV,\Sigma} := \max_{1 \leq k \leq K} ||\boldsymbol{\sigma}_k - \boldsymbol{\sigma}_k'||_{TV}.$$

for all $\boldsymbol{\sigma}, \boldsymbol{\sigma}' \in \Sigma$. As in (i) we see that for all $1 \leq k \leq K$ one can bound the distance between $\boldsymbol{\sigma}_k^N$ and $\boldsymbol{\sigma}_k$ by

$$||\boldsymbol{\sigma}_k^N - \boldsymbol{\sigma}_k||_{TV} \leq \frac{n}{2Nm_k^N}$$

Consequently for all N sufficiently large we have

$$||\boldsymbol{\sigma} - \boldsymbol{\sigma}^N||_{TV,\Sigma} \leq \frac{n}{2N} \max_{1 \leq k \leq K} \frac{1}{m_k^N}.$$

Since $\mathbf{m}^N \to \mathbf{m} \in \operatorname{int}\Delta(\Theta)$ it follows that for all $\epsilon > 0$ there exists a N sufficiently large so that $||\boldsymbol{\sigma} - \boldsymbol{\sigma}^N||_{TV,\Sigma} \leq \epsilon$. This completes the proof. □

B.2.1 Proof of Lemma 2

Proof As a first step we have to determine the asymptotic behavior of the factors determining the functions $\Phi_k^a(\cdot)$, i.e. the large population behavior of the numbers $\varphi_{k,l}^{\beta,N}(a,b) = \frac{2\beta \exp(v(a,b)/\beta)}{\xi_{kl}^{\beta,N}}$. Assuming that the volatility rates satisfy (LPB), then, for all $1 \leq k, l \leq K$ and $a, b \in A$,

$$\lim_{N\to\infty} \frac{1}{N}\varphi_{k,l}^{\beta,N}(a,b) = 0, \quad \lim_{N\to\infty} \varphi_{k,l}^{\beta,N}(a,b) = \frac{2\exp(v(a,b)/\beta)}{\xi_{kl}^{\beta}}.$$

[16] The choice of norm is, of course, not essential here.

This implies that the first-order approximation

$$\log\left(1+\frac{\varphi_{kl}^{\beta,N}(a,b)}{N\beta}\right)=\frac{\varphi_{kl}^{\beta,N}(a,b)}{N\beta}+O(N^{-2}\beta^{-1})$$

gives the right asymptotic behavior for sufficiently large N. For all $a\in A$ and $1\leq k<l\leq K$ observe that

$$\log\left[\Phi_{kk}^{a}(\boldsymbol{\sigma}^N,\beta,N)\right]=\frac{1}{\beta}\theta_k(a)+\sum_{b\geq a}\left(\frac{Nm_k^N\sigma_k^N(b)-\delta_{a,b}}{1+\delta_{a,b}}\right)\log\left(1+\frac{\varphi_{kk}^{\beta,N}(a,b)}{N\beta}\right)$$
$$=\frac{1}{\beta}\left[\theta_k(a)+\frac{1}{2}m_k^N\sigma_k^N(a)\varphi_{kk}^{\beta,N}(a,a)+\sum_{b>a}m_k^N\sigma_k^N(b)\varphi_{kk}^{\beta,N}(a,b)+O(1/N)\right],$$

and

$$\log\left[\Phi_{kl}^{a}(\boldsymbol{\sigma}^N,\beta,N)\right]=\frac{1}{\beta}\left[m_l^N\sum_{b\in A}\sigma_l^N(b)\varphi_{kl}^{\beta,N}(a,b)+O(1/N)\right].$$

Thus, for all $1\leq k\leq K$ we see that

$$f_k^{\beta,N}(\boldsymbol{\sigma}^N,\mathbf{m}^N)=\sum_{a\in A}\sigma_k^N(a)\sum_{l\geq k}\log\Phi_{kl}^{a}(\boldsymbol{\sigma}^N,\beta,N)$$
$$=\frac{1}{\beta}\left[\left\langle\boldsymbol{\sigma}_k^N,\boldsymbol{\theta}_k\right\rangle+\sum_{l\geq k}\frac{m_l^N}{1+\delta_{kl}}\left\langle\boldsymbol{\sigma}_k^N,\boldsymbol{\varphi}_{kl}^{\beta,N}\boldsymbol{\sigma}_l^N\right\rangle+O(1/N)\right]$$
$$=\frac{1}{\beta}\left(f_k^{\beta}(\boldsymbol{\sigma}^N,\mathbf{m}^N)+O(1/N)\right).$$

All the functions appearing in the definition of $f_k^{\beta}(\boldsymbol{\sigma}^N,\mathbf{m}^N)$ have a well defined limit as $N\to\infty$, and therefore the proof is completed. □

Corollary 2 *The sequence of functions* $\{f^{\beta,N}\}_{N\geq N^0}$ *converges uniformly to the limit function* f^{β}.

Proof This follows from Lemma 3 together with Lemma 2. □

B.3 The large deviations principle

The proof of the large deviations principle for the family of measures $\{P_N^{\beta}\}_{N\in\mathbb{N}}$ proceeds in two steps. We first introduce a simple reference probability space for which a large deviations principle can be obtained by an application of Sanov's Theorem in a quite straightforward way (Lemma 4). We will use this result in order to reformulate the law P_N^{β} in terms of this reference measure, so that we can derive the LDP for P_N^{β}

from Sanov's theorem proved for the reference measure, by applying a theorem from den Hollander (2000).[17]

B.3.1 An auxiliary probability space

We start with the construction of an auxiliary sequence of probability spaces $\{(\Omega^N, \mathcal{F}^N, \mathbb{P}^N)\}_{N\in\mathbb{N}}$, where the family of probability measures $\{\mathbb{P}^N\}_{N\in\mathbb{N}}$ will be used as reference measures. The state space is the set of action-type profiles $\Omega^N := A^N \times \Theta^N$. The sigma-algebra is simply the set of subsets 2^{Ω^N}. On this measurable pair $(\Omega^N, \mathcal{F}^N)$, we define the probability measure

$$\mathbb{P}^N(\{\omega\}) := e^{-NI_q(\mathbf{M}^N(\omega))} \quad \forall\omega = (\mathbf{a}, \boldsymbol{\tau}) \in \Omega^N,$$

where the function $I_q : \Delta(\Theta) \to \mathbb{R}$ is defined as

$$\mathbf{m} \mapsto I_q(\mathbf{m}) := \log(n) - \sum_{k=1}^{K} m_k \log(q_k).$$

Under this probability measure the types of the players are i.i.d distributed with law $\mathbf{q}$, and the players choose actions with uniform probability $1/n$. Thus, the marginal distribution on A^N implied by this measure is $n^{-N} = e^{-N\log(n)}$. The law of the empirical processes $(\mathbf{S}^N, \mathbf{M}^N)$, under the measure $\mathbb{P}^N$, is denoted by $\hat{P}^N := \mathbb{P}^N \circ (\mathbf{S}^N, \mathbf{M}^N)^{-1}$, so that

$$\hat{P}^N(\boldsymbol{\sigma}, \mathbf{m}) := \begin{cases} \frac{N!}{\prod_{k=1}^{K}\prod_{a=1}^{n}(Nm_k\sigma_k(a))!} e^{-NI_q(\mathbf{m})} & \text{if } (\boldsymbol{\sigma}, \mathbf{m}) \in \mathcal{K}^N, \\ 0 & \text{otherwise.} \end{cases}$$

Our first result, Lemma 4, is a version of Sanov's Theorem (see e.g. den Hollander 2000). This Lemma is needed as it characterizes the rate function implied by the large deviations principle satisfied by the law $\hat{P}^N$ as $N \to \infty$. This rate function will then be used to state and proof a large deviations principle for the family $\{P_N^\beta\}_{N\in\mathbb{N}}$. To state this lemma, we need slightly more notation. Let P and Q be two probability measures defined on a common discrete set X. The *relative entropy* of P with respect to Q is defined as

$$h(P||Q) := \sum_{x\in X} P(x) \log \frac{P(x)}{Q(x)}.$$

The *entropy* of P is defined as

$$h(P) := -\sum_{x\in X} P(x) \log P(x).$$

[17] I am deeply indebted to an anonymous referee for suggesting this proof strategy.

Remark 6 In the following we will always use the letter h to denote the (relative) entropy of a probability measure, independent of the underlying space. This abuse of notation should cause no confusion.

Lemma 4 *The sequence $\{\hat{P}^N\}_{N\in\mathbb{N}}$ satisfies a large deviations principle with rate function*

$$F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m}) := \log(n) + h(\mathbf{m}||\mathbf{q}) - \sum_{k=1}^{K} m_k h(\boldsymbol{\sigma}_k).$$

That is, for an arbitrary set $A \times B =: \Gamma \subset \mathcal{K}$, we have

$$-\inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma^o} F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m}) \leq \liminf_{N\to\infty} \frac{1}{N}\log \hat{P}^N(\Gamma)$$

$$\limsup_{N\to\infty} \frac{1}{N}\log \hat{P}^N(\Gamma) \leq -\inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma} F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m}),$$

where Γ^o is the relative interior of Γ in $\mathcal{K}$.

Proof For any set $A \times B \subseteq \mathcal{K}$ its $\hat{P}^N$-probability is given by

$$\hat{P}^N(A\times B) = \mathbb{P}^N\left(\mathbf{S}^N \in A, \mathbf{M}^N \in B\right).$$

We have the elementary lower bound

$$\max_{(\boldsymbol{\sigma},\mathbf{m})\in A\times B\cap\mathcal{K}^N} \hat{P}^N(\boldsymbol{\sigma},\mathbf{m}) \leq \hat{P}^N(A\times B),$$

From the formal computation

$$\begin{aligned}
\hat{P}^N(A\times B) &= \sum_{\mathbf{m}\in B\cap\mathcal{L}_N}\ \sum_{\boldsymbol{\sigma}\in A\cap\Sigma^N(\mathbf{m}^N)} \hat{P}^N(\boldsymbol{\sigma},\mathbf{m}) \\
&= \sum_{\mathbf{m}\in B\cap\mathcal{L}_N} \mathsf{P}^N_{\mathbf{q}}(\mathbf{m})\left(\sum_{\boldsymbol{\sigma}\in A\cap\Sigma^N(\mathbf{m})} \hat{P}^N(\boldsymbol{\sigma}|\mathbf{m})\right) \\
&\leq |\mathcal{L}_N\cap B| \max_{\mathbf{m}\in\mathcal{L}_N\cap B}\left\{\mathsf{P}^N_{\mathbf{q}}(\mathbf{m})\left(\sum_{\boldsymbol{\sigma}\in A\cap\Sigma^N(\mathbf{m})} \hat{P}^N(\boldsymbol{\sigma}|\mathbf{m})\right)\right\} \\
&\leq |\mathcal{L}_N| \max_{\mathbf{m}\in\mathcal{L}_N\cap B}\left\{\mathsf{P}^N_{\mathbf{q}}(\mathbf{m})|\Sigma^N(\mathbf{m})| \max_{\boldsymbol{\sigma}\in\Sigma^N(\mathbf{m})} \hat{P}^N(\boldsymbol{\sigma}|\mathbf{m})\right\}
\end{aligned}$$

we obtain the upper bound

$$\hat{P}^N(A\times B) \leq |\mathcal{L}_N| \max_{\mathbf{m}\in\mathcal{L}_N\cap B}\left\{\mathsf{P}^N_{\mathbf{q}}(\mathbf{m})|\Sigma^N(\mathbf{m})| \max_{\boldsymbol{\sigma}\in\Sigma^N(\mathbf{m})} \hat{P}^N(\boldsymbol{\sigma}|\mathbf{m})\right\}$$

Observe that $|\mathcal{L}_N| = \binom{N+K-1}{K-1} = O(N^{K-1})$ and $|\Sigma_k^N(\mathbf{m})| = \binom{Nm_k+n-1}{n-1} = O(N^{n-1})$ for every $k \in \{1, 2, \ldots, K\}$. Hence $|\Sigma^N(\mathbf{m})| = O(KN^{n-1})$. Therefore, for the upper bound, we get

$$\frac{1}{N}\log \hat{P}^N(\Gamma) \leq O(\log(N)/N) + \frac{1}{N}\log\left(\max_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma\cap\mathcal{K}^N} \hat{P}^N(\boldsymbol{\sigma},\mathbf{m})\right). \tag{19}$$

The lower bound gives us in turn

$$\frac{1}{N}\log \hat{P}^N(\Gamma) \geq \frac{1}{N}\log\left(\max_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma\cap\mathcal{K}^N} \hat{P}^N(\boldsymbol{\sigma},\mathbf{m})\right). \tag{20}$$

To estimate the probability in the brackets, pick a sequence $(\boldsymbol{\sigma}^N, \mathbf{m}^N) \to (\boldsymbol{\sigma}, \mathbf{m}) \in \mathcal{K}$ with $(\boldsymbol{\sigma}^N, \mathbf{m}^N) \in \mathcal{K}^N$ for each N, and assume that $\mathbf{m} \in \operatorname{int}\Delta(\Theta)$. Then, an iterative application of Stirling's formula gives us

$$\begin{aligned}
\frac{1}{N}\log \hat{P}^N(\boldsymbol{\sigma}^N,\mathbf{m}^N) &= -I_{\mathbf{q}}(\mathbf{m}^N) + h(\mathbf{m}^N) + \sum_{k=1}^K m_k^N h(\boldsymbol{\sigma}_k^N) + O(\log(N)/N)\\
&= -\log(n) - h(\mathbf{m}^N||\mathbf{q}) + \sum_{k=1}^K m_k^N h(\boldsymbol{\sigma}_k^N) + O(\log(N)/N)\\
&= -F_{\mathbf{q}}(\boldsymbol{\sigma}^N,\mathbf{m}^N) + O(\log(N)/N)
\end{aligned}$$

as $N \to \infty$. Hereby we have used the continuity of the involved functions and the fact that the sequence of type distributions is assumed to be interior. Hence, for the upper bound (19), we get that

$$\limsup_{N\to\infty}\frac{1}{N}\log \hat{P}^N(\Gamma) \leq -\liminf_{N\to\infty}\left\{\inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma\cap\mathcal{K}^N} F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m})\right\}$$

Since $\Gamma \cap \mathcal{K}^N \subset \Gamma$ for all N, it follows that $\inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma\cap\mathcal{K}^N} F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m}) \geq \inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma} F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m})$ for all N. Hence,

$$\limsup_{N\to\infty}\frac{1}{N}\log \hat{P}^N(\Gamma) \leq -\inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma} F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m}).$$

To finish with the lower bound, fix a point $(\boldsymbol{\sigma}, \mathbf{m}) \in \Gamma^o$ and pick an interior sequence $(\boldsymbol{\sigma}^N, \mathbf{m}^N) \in \Gamma \cap \mathcal{K}^N$ converging to this point. This is possible by Lemma 3. Then, using the estimate (20), we see that

$$\begin{aligned}\liminf_{N\to\infty}\frac{1}{N}\log\hat{P}^N(\Gamma) &\geq -\limsup_{N\to\infty}\left\{\inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma\cap\mathcal{K}^N}F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m})\right\}\\ &\geq -\lim_{N\to\infty}F_{\mathbf{q}}(\boldsymbol{\sigma}^N,\mathbf{m}^N)\\ &= -F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m}) \geq -\inf_{(\boldsymbol{\sigma},\mathbf{m})\in\Gamma^o}F_{\mathbf{q}}(\boldsymbol{\sigma},\mathbf{m}).\end{aligned}$$

□

B.3.2 An expression for the law of action-type distributions

We will now use the constructions from the previous section to reformulate the law of action-type distributions P_N^β defined in eq. (12). From this reformulation, a large deviations principle will follow immediately. First, we define a measure on the set $\mathcal{K}$, by

$$\gamma_N^\beta(\{(\boldsymbol{\sigma},\mathbf{m})\}) := \int_{(\mathbf{S}^N,\mathbf{M}^N)^{-1}(\boldsymbol{\sigma},\mathbf{m})} e^{N[f_N^\beta(\mathbf{S}^N(\omega),\mathbf{M}^N(\omega))+\log(n)]}\mathrm{d}\mathbb{P}^N(\omega),$$

where $(\mathbf{S}^N,\mathbf{M}^N)^{-1}(\boldsymbol{\sigma},\mathbf{m}) := \{\omega\in\Omega^N|(\mathbf{S}^N(\omega),\mathbf{M}^N(\omega)) = (\boldsymbol{\sigma},\mathbf{m})\}$ for $(\boldsymbol{\sigma},\mathbf{m})\in\mathcal{K}$. Given the definition of the reference measure $\mathbb{P}^N$, it is straightforward to verify that this measure agrees with the one introduced in Eq. 11. Based on this, we recover (12) by

$$P_N^\beta(\{(\boldsymbol{\sigma},\mathbf{m})\}) = \frac{\gamma_N^\beta(\{(\boldsymbol{\sigma},\mathbf{m})\})}{\gamma_N^\beta(\mathcal{K})} \quad \forall(\boldsymbol{\sigma},\mathbf{m})\in\mathcal{K}.$$

Using this formula we prove Theorem 4, which we restate here for the readers' convenience, using Lemma (4).

Theorem 6 *The family of measures $\{P_N^\beta\}_{N\in\mathbb{N}}$ satisfies a large deviations principle with speed N and rate function*

$$R_{\mathbf{q}}^\beta(\boldsymbol{\sigma},\mathbf{m}) := \max_{(\boldsymbol{\sigma}',\mathbf{m}')\in\mathcal{K}}\left[\frac{1}{\beta}\tilde{f}^\beta(\boldsymbol{\sigma}',\mathbf{m}') - h(\mathbf{m}'||\mathbf{q})\right] - \left[\frac{1}{\beta}\tilde{f}^\beta(\boldsymbol{\sigma},\mathbf{m}) - h(\mathbf{m}||\mathbf{q})\right].$$

Proof Define first a measure $\hat{\gamma}_N^\beta$ on the set $\mathcal{K}$, absolutely continuous with respect to the probability measure $\mathbb{P}^N$, as

$$\hat{\gamma}_N^\beta(\{(\boldsymbol{\sigma},\mathbf{m})\}) := \int_{(\mathbf{S}^N,\mathbf{M}^N)^{-1}(\boldsymbol{\sigma},\mathbf{m})} e^{N[\frac{1}{\beta}f^\beta(\mathbf{S}^N(\omega),\mathbf{M}^N(\omega))+\log(n)]}\mathrm{d}\mathbb{P}^N(\omega),$$

for all $(\boldsymbol{\sigma},\mathbf{m})\in\mathcal{K}$. Then, we can define a probability measure on $\mathcal{K}$ as

$$Q_N^\beta(\Gamma) := \frac{\hat{\gamma}_N^\beta(\Gamma)}{\hat{\gamma}_N^\beta(\mathcal{K})},\quad \Gamma\subseteq\mathcal{K}.$$

From Lemma 4 we know that the reference measure $\mathbb{P}^N$ satisfies a large-deviations principle with rate function $F_{\mathbf{q}}$. Applying Theorem III.17 in den Hollander (2000), we conclude that the sequence of probability measures $\{Q_N^\beta\}_{N\in\mathbb{N}}$ satisfies a "tilted" large deviations principle with rate function $R_{\mathbf{q}}^\beta$. The only remaining step is to show that the family of measures $\{P_N^\beta\}_{N\in\mathbb{N}}$ satisfies the same LDP. This follows, however, from the uniform convergence of the interaction potential function. Specifically, by Lemma 2, for every $\varepsilon > 0$ there exists a population size N_0, such that for all $N \geq N_0$ we have $|f_N^\beta(\boldsymbol{\sigma}, \mathbf{m}) - \frac{1}{\beta} f^\beta(\boldsymbol{\sigma}, \mathbf{m})| < \varepsilon$ for all $(\boldsymbol{\sigma}, \mathbf{m}) \in \mathcal{K}^N$. Hence, $P_N^\beta(\{(\boldsymbol{\sigma}, \mathbf{m})\}) = Q_N^\beta(\{(\boldsymbol{\sigma}, \mathbf{m})\})e^{o(N)}$, where $o(N)$ is a remainder term satisfying $\frac{o(N)}{N} \to 0$ uniformly (everything takes place on the compact set $\mathcal{K}$) as $N \to \infty$. $\square$

References

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

Blume L (1993) Statistical mechanics of strategic interaction. Games Econ Behav 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, Menlo-Park, pp 425–462

Bollobás B, Janson S, Riordan O (2007) The phase transition in inhomogeneous random graphs. Random Struct Algorithm 31(1):3–122

den Hollander F (2000) Large deviations. Fields institute monographs. American Mathematical Society, Providence

Durrett R (2007) Random graph dynamics, Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge

Ely J, Sandholm WH (2005) Evolution in Bayesian games I: theory. Games Econ Behav 53:83–109

Ethier SN, Kurtz TG (1986) Markov processes: characterization and convergence. Wiley, New York

Fudenberg D, Levine DK (1998) The theory of learning in games. The MIT Press, Cambridge

Golub B, Jackson MO (2010) How homophily affects diffusion and learning in networks. Mimeo. http://arxivorg/abs/08114013. Accessed 26 September 2010

Goyal S, Vega-Redondo F (2005) Network formation and social coordination. Games Econ Behav 50:178–207

Hofbauer J, Sandholm WH (2007) Evolution in games with randomly disturbed payoffs. J Econ Theory 132:47–69

Hofbauer J, Sigmund K (1988) The theory of evolution and dynamical systems. Cambridge University Press, Cambridge

Horst U, Scheinkman JA (2006) Equilibria in systems of social interactions. J Econ Theory 130(1):44–77

Jackson MO, Watts A (2002) On the formation of interaction networks in social coordination games. Games Econ Behav 41:265–291

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

Monderer D, Shapley L (1996) Potential games. Games Econ Behav 14:124–143

Park J, Newman M (2004) The statistical mechanics of networks. Phys Rev E 70:066117

Sandholm WH (2007) Pigouvian pricing and stochastic evolutionary implementation. J Econ Theory 132:367–382

Sandholm WH (2010a) Orders of limits for stationary distributions, stochastic dominance, and stochastic stability. Theor Econ 5:1–26

Sandholm WH (2010b) Population games and evolutionary dynamics. The MIT Press, Cambridge

Söderberg B (2002) General formalism for inhomogeneous random graphs. Phys Rev E 66:066121–066126

Staudigl M (2010a) Co-evolutionary dynamics and bayesian interaction games. EUI Working papers. European University Institute, Florence

Staudigl M (2010b) Co-evolutionary dynamics of networks and play. PhD Thesis. Department of Economics, University of Vienna, Vienna

Staudigl M (2010c) On a general class of stochastic co-evolutionary dynamics. Working Paper No: 1001. Department of Economics, University of Vienna, Vienna

Staudigl M (2011) Potential games in volatile environments. Games Econ Behav 72(1):271–287

Vega-Redondo F (2007) Complex social networks, econometric society monograph series. Cambridge University Press, Cambridge

Young H (1993) The evolution of conventions. Econometrica 61(1):57–84

Young HP (2003) The diffusion of innovations in social networks. In: Blume LE, Durlauf SN (eds) The economy as an evolving complex system III. Oxford University Press, New York