MAX WEBER PROGRAMME

EUI Working Papers

MWP 2010/35
MAX WEBER PROGRAMME

CO-EVOLUTIONARY DYNAMICS AND BAYESIAN INTERACTION GAMES

Mathias Staudigl

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE
MAX WEBER PROGRAMME

Co-evolutionary Dynamics and Bayesian Interaction Games

MATHIAS STAUDIGL

EUI Working Paper **MWP** 2010/35

ISSN 1830-7728

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
Hwww.eui.eu
Hcadmus.eui.eu

Abstract
In a Bayesian interaction game players have diverse preferences and are randomly matched according to an inhomogeneous random graph. A co-evolutionary process of networks and play gives a dynamic formalism for the joint evolution of the random graph and the actions the players use per match. Assuming that the players select actions and links according to log-linear functions taking as arguments the reward per match, we provide closed form solutions for the joint invariant distribution of the co-evolutionary process. We give sufficient conditions for the general selection of potential maximizers in the small noise limit, and also discuss concentration of the invariant distribution in the large population limit. Further, we present a general characterization theorem that a co-evolutionary process generates inhomogeneous random graphs, a large and important class of random graphs recently discussed in the economic and mathematical literature on random networks.

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

JEL Classifications: C73, D83, D58

Mathias Staudigl
Max Weber Fellow, 2010-2011

1 Introduction

In an interaction game (Morris, 1997) players are involved in a series of interactions. An interaction, or matching, can be modeled as the outcome of a random graph process. Random graph models provide a flexible modeling environment to study interaction games on general interaction structures, and therefore have been used frequently in the literature on social interactions (Horst and Scheinkman, 2006), communication (Kirman et al., 1986) and diffusion (Jackson and Yariv, 2008).[1] In these papers it is assumed that there exists a certain exogenous and independent probability that two agents are randomly matched in order to play some given game. The purpose of this paper is to propose a model where these interaction probabilities (or *edge-success probabilities*) evolve over time as a function of the actions used by the two players, which are themselves dynamic variables. Such a two-sided process, which we call a *co-evolutionary process of networks and play*, has been studied in Staudigl (2010b) in the context of partnership games (Hofbauer and Sigmund, 1988). We extend our previous work by presenting some more general results in terms of the induced interaction structure and by introducing (random) heterogeneity in the preferences of the players.[2] On the technical side we make in this paper the transition rates dependent on the population size, which gives us an additional adjustable parameter. This added flexibility of the model will lead to new insights, both in terms of the long-run interaction structure as well as the distribution over strategies players use in the long run. In terms of economic interpretations, we prove a general representation theorem that says that a co-evolutionary process of networks and play generates an *inhomogeneous random graph*. This type of random graph has been studied in physics and mathematics intensively (see e.g. Söderberg (2002) and Bollobás et al. (2007)). 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.

As static model we consider *Bayesian interaction games*. For a fixed (commonly known) pattern of interactions we assume that each player's utility function consists of two terms; (i) the expected reward of a player in an interaction game having the partnership structure and (ii) an idiosyn-

[1]The explosive field of social and economic networks is nicely surveyed by the books of Goyal (2007), Jackson (2008b) and Vega-Redondo (2007) where many more possible applications can be found.

[2]See also the models of Ehrhardt et al. (2006a;b; 2008) for early proponents of the co-evolutionary approach in economics.

cratic payoff component interpreted as the *type* of the player. Assumption (i) guarantees that the game is an exact potential game in the sense of Monderer and Shapley (1996). Types are random realizations from an i.i.d. draw from a common finite type space.[3] Utility functions of this kind can be used in various social and economic settings, and we provide two canonical examples, a coordination game on a network and the linear-quadratic game of Ballester et al. (2006), which fall into the class of games we study. For our dynamic study we fix an arbitrary profile of types in the population and study the asymptotic behavior of a continuous-time Markov process similar to Staudigl (2010b) (this could be interpreted as a dynamic process on the ex-post stage of the Bayesian game, i.e. when players get to know their own types). Our results concerning the random dynamics are as follows.

We start in Section 3.1 with the characterization of the structure of the random graph generated by the process. As mentioned above, we show that a co-evolutionary process of networks and play can be used as a dynamic model that explains the formation of inhomogeneous random graphs (IHRG). Our representation theorem is proved as a conditional statement. First, we condition on a type profile, and second we condition on an arbitrary action profile. For every pair of type-action profiles the full structure of the IHRG is proven, so that at each point in the relevant part of the state space a full characterization of the random graph is available. We then proceed in Section 3.2 in presenting a closed form solution of the *joint* invariant distribution of the Markov process under the additional assumption that the attachment mechanism, describing how two players get to form a link with each other, is a log-linear function of the mutual rewards the players receive per interaction.[4] The economic interpretation is that players who expect a higher reward, are more likely to meet in the link creation process.[5] By making this functional assumption we do not only gain the possibility to proceed with a full analytical study

[3]Our terminology of calling the normal form game a Bayesian interaction game comes from the fact that viewing the game from an aggregate perspective the model can be seen as a particular version of the Bayesian population games of Ely and Sandholm (2005). The word interaction should highlight that our games are played on (random) networks, and so the utilities of the agents will depend on the (realized) network.

[4]Log-linear functions are very common in the statistical literature on random graphs. The p^*-models, building on the Markov graphs of Frank and Strauss (1986), are a prominent examples. See Robins et al. (2007) for a recent survey.

[5]As in Staudigl (2010b) we could alternatively interpret the link creation process as a process where players engage in active search for new interaction partners, where the rate at which two players find themselves is an increasing function in the utility they get from being connected.

of the random dynamics, but also gives a modeler a socio-economic interpretation of the IHRG, as we argue in Section 3.1.[6] The closed form of the joint invariant distribution is an extension of Staudigl (2010b) and is key to our final investigation of equilibrium selection, which we start in Section 4. Conditional on a realization of types, the co-evolutionary process depends on two exogenous parameters: The number of players (N), and a noise parameter (β) modeling the degree of (bounded) rationality of the players, how they form links and how they choose actions. When either N goes to infinity or β shrinks to 0 we obtain a sequence of invariant distributions, whose support will concentrate on certain subsets of the state space. Hence, both of these limit exercises can be used for equilibrium selection.[7] In the small noise limit we provide sufficient conditions on the shape of the rate functions of the co-evolutionary process which guarantee the selection of potential maximizers in the long-run. This extends previous results of Staudigl (2010b) to more general network formation dynamics. Further, it gives an alternative way to assess the robustness of potential maximizers against endogenous interaction.[8] In particular, this results implies that in the small noise limit all pairs of networks and action profiles, which are in the support of the limiting invariant distribution, are mutually consistent in the sense that the action profile is a Nash equilibrium on the respective network. This observation is used to give a game-theoretic interpretation of the invariant distribution as a correlation device of the players, which we call a (β, ρ)-correlated equilibrium.[9] Finally, our investigation of the large population limit focuses on the marginal distribution over action profiles. Since the interaction structure is completely characterized in Section 3.1, the study of the marginal distribution on the set of action profiles is the only remaining part in order to achieve a complete characterization of the long-run behavior of the invariant distribution. In the large population limit we are

[6]In his recent book Jackson (2008b) criticizes the use of random graph models exactly because they lack any socio-economic foundation. We think that our model is one possibility to provide such a foundation.

[7]As discussed in more detail in section 3.1 we are not able to take the double limit $(N, \beta) \to (\infty, 0)$ at the same speed, due to the connection with law of the random graph and the noise level β. In essence, as $\beta \to 0$ the random graph tends to be degenerate at the complete network at an exponential rate (this comes from our log-linear specification). Hence, we have to take these two limits separately to get meaningful results.

[8]It has been shown in Ui (2001), and extended in Morris and Ui (2005), that potential maximizers are robust against incomplete information in the sense of Kajii and Morris (1997).

[9]See Mailath et al. (1997) for a similar result. This (approximate) equilibrium concept is inspired by the ex-post Nash equilibrium of Kalai (Kalai, 2004, Definition 4, p. 1641).

interested in the aggregate play of the agents, which we measure in terms of *Bayesian strategies*. A Bayesian strategy records the relative frequency of players of a certain type who play a certain action, and can therefore be interpreted as a map from the type space to the action space. Hence, it formally agrees with the conventional definition of a Bayesian strategy, justifying its name.[10] In this section we show that, for a fixed and positive level of noise, the measure over Bayesian strategies selects maximizers of a so-called "logit-potential function" (Hofbauer and Sandholm, 2002; 2007). Maximizers of a logit-potential function are approximate Nash equilibria, whose distance from Nash equilibrium population profiles depends on the degree of noise. Hence, small noise limits and large population limits will, in general, predict different outcomes in the long-run. Moreover, the large population limit produces a well-defined random graph model, whereas the small noise limit always has the tendency to produce densely connected networks (see section 3.1 for a formal statement). Therefore, we think that in a model of the kind studied in this paper the large population limit should be of more relevance.[11]

The rest of the paper is organized as follows. Section 2 introduces Bayesian interaction games and the co-evolutionary process of networks and play. Section 3 starts our investigation of the long-run properties of the co-evolutionary process, conditional on an arbitrary type profile. The random graph representation theorem is provided in Section 3.1. Section 3.2 gives the closed-form of the joint invariant distribution of the co-evolutionary process. The final section of this paper is concerned with equilibrium selection. Section 4.1 treats the small noise limit behavior of the invariant distribution, while section 4.2 studies the large population limit. Two technical appendices accompany the paper where lengthy and technical proofs, omitted from the main text, are collected.

2 Co-evolutionary dynamics

One of the main concerns of this paper is to propose a dynamic model where players adjust their actions and their interaction network perpetually. We do this in form of a stochastic evolutionary process, which we

[10] This population based interpretation of Bayesian strategies is due to Ely and Sandholm (2005).

[11] The large population limit also captures the fact that most social and economic networks of interest are large and show a complicated topological structure. In evolutionary game theory there is a long lasting debate on which limit should be given precedence. An early account of the different aspects the two limit operations emphasize can be found in Binmore et al. (1995) and Binmore and Samuelson (1997).

call a *co-evolutionary process of networks and play*. Before describing the random dynamics in some detail, we present the static game-theoretic model. The notation we employ is as follows. For any finite set $\mathcal{V}$ we denote by $\mathcal{V}^2$ the set of ordered pairs of elements in this set, while $\mathcal{V}^{(2)}$ is the set of unordered pairs of elements. The set of probability distributions on the finite set $\mathcal{V}$ is denoted by $\Delta(\mathcal{V})$. The indicator function of an arbitrary set A is denoted as $\mathbb{1}_A(\cdot)$, i.e. $\mathbb{1}_A(x) = 1$ if $x \in A$ and 0 otherwise. Bold letters are used to denote tuples and matrices. On the real vector space $\mathbb{R}^n$ we declare the standard inner product as $\langle \boldsymbol{x}, \boldsymbol{y} \rangle := \sum_{i=1}^n x_i y_i$ for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$. We identify networks (or graphs) as a tuple of binary variables $\boldsymbol{g} = (g_{ij})_{(i,j) \in \mathcal{V}^{(2)}}, g_{ij} \in \{0,1\}$. For a given network $\boldsymbol{g}$ we use the notation $\boldsymbol{g} \oplus (i,j)$ to indicate that the (previously not existing) edge (i,j) is added to the network. Similarly, we denote by $\boldsymbol{g} \ominus (i,j)$ the network obtained from $\boldsymbol{g}$ by deleting the edge (i,j).

2.1 Bayesian interaction games

We start by presenting the game theoretic framework on which the co-evolutionary dynamics operate. At the end of the section we provide two concrete examples to which a Bayesian interaction game may be applied. We consider a family of normal form games Γ, whose members are *Bayesian interaction games* $\Gamma_{\boldsymbol{p}}^N = \left\langle [N], A, \Theta, (U_i(\cdot, \boldsymbol{p}, \tilde{\tau}_i))_{i \in [N]} \right\rangle$. $[N] := \{1, 2, \ldots, N\}$ is the set of players, A is a finite set of actions and $\Theta := \{\theta_1, \ldots, \theta_K | \theta_k : A \to \mathbb{R}\}$ is the finite set of *types* of the players. We denote by a_i the action of player i. τ_i is a realization of the random variable $\tilde{\tau}_i \in \Theta$ modeling the type of player $i \in [N]$. Each player can be of K different types, where K is an arbitrary integer. An action profile is a list $\boldsymbol{a} := (a_1, \ldots, a_N)$ and a type profile is a list $\boldsymbol{\tau} := (\tau_1, \ldots, \tau_N)$. The array $\boldsymbol{p} := (p_{ij})_{j>i}$ is a list of *interaction probabilities* which are commonly known to the players. Interaction probabilities have the property that (i) they are *symmetric*, i.e $p_{ij} = p_{ji}$ for all $j \neq i$, and (ii) there is no self-matching, i.e $p_{ii} = 0$ for all $i \in [N]$. The number $p_{ij} \in [0,1]$ is the independent probability that player i will be matched with player j (and vice versa). We therefore have in mind that matchings are realizations of an *inhomogeneous random graph* (IHRG), according to the following definition.

Definition 1. *For a given number of players N let $\mathcal{G}[N]$ denote the set of all undirected and unweighted graphs on the vertex set $[N]$ identified by a tuple $\boldsymbol{g} := (g_{ij})_{(i,j) \in [N]^{(2)}}$. The indicator function $g_{ij} \in \{0,1\}$ is interpreted as a link (edge) between player i and j (and vice versa). An* inhomogeneous random graph *is a probability model $\mathcal{G}[N, \boldsymbol{p}]$ in which each link is an independent* Bernoulli(p_{ij})

random variable.

Remark 1. *Note that the above definition covers "degenerate" situations where the random graph puts probability 1 to a given fixed network* $\boldsymbol{g}$.

Given an action profile $\boldsymbol{a}$ and a profile of types $\boldsymbol{\tau}$, the expected payoff of player i is assumed to be

$$U_i(\boldsymbol{a}, \boldsymbol{p}, \tau_i) = \sum_{j \neq i} \pi_i(a_i, a_j) p_{ij} + \tau_i(a_i). \tag{2.1}$$

The interpretation of this utility function is that each player is involved in a series of 2-player games, each of these is played with the independent probability p_{ij}. In a match with an opponent player i's reward is measured by the function π_i. The weighted sum of all rewards measures the total reward the player makes out of his interactions. On top of this, each player can have an idiosyncratic preference over the set of available actions. For our analytical results we require additional structure on the form of the reward functions of the players. We say that an interaction game is an exact potential game (Monderer and Shapley, 1996) if there exists a function $V(\cdot, \boldsymbol{p}, \boldsymbol{\tau}) : A^N \to \mathbb{R}$ with the property that

$$\begin{aligned} V((a, \boldsymbol{a}_{-i}), \boldsymbol{p}, \boldsymbol{\tau}) - & V((b, \boldsymbol{a}_{-i}), \boldsymbol{p}, \boldsymbol{\tau}) \\ & = U_i((a, \boldsymbol{a}_i), \boldsymbol{p}, \tau_i) - U_i((b, \boldsymbol{a}_{-i}), \boldsymbol{p}, \tau_i) \end{aligned} \tag{2.2}$$

for all $i \in [N], a, b \in A, \boldsymbol{a}_{-i} \in A^{N-1}$. Sufficient (though not necessary) for being a potential game is that $\Gamma^N_{\boldsymbol{p}}$ has the *partnership structure* (Hofbauer and Sigmund, 1988), i.e. if the reward function of every player $i \in [N]$ can be expressed through a single function $v : A^2 \to \mathbb{R}_+$ such that

$$\pi_i(a, b) = v(a, b) = v(b, a)$$

for all $a, b \in A$. Interaction games with the 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.[12] We call $\Gamma^N_{\boldsymbol{p}}(\boldsymbol{\tau})$ the complete information game when the type profile is fixed to be $\boldsymbol{\tau}$.

Lemma 2.1. *The interaction game* $\Gamma^N_{\boldsymbol{p}}(\boldsymbol{\tau})$ *with common reward function* v *is an exact potential game with potential function*

$$V(\boldsymbol{a}, \boldsymbol{p}, \boldsymbol{\tau}) = \frac{1}{2} \sum_{i,j \in [N]} v(a_i, a_j) p_{ij} + \sum_{i \in [N]} \tau_i(a_i) \tag{2.3}$$

[12]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.

Proof. We have to show that the function $V(\cdot, \boldsymbol{p}, \boldsymbol{\tau})$ satisfies condition (2.2). If V is defined as in (2.3), then the left-hand side boils down to

$$\frac{1}{2}\sum_{j\in[N]}\left[v(a,a_j)-v(b,a_j)\right]p_{ij}+\frac{1}{2}\sum_{j\in[N]}\left[v(a_j,a)-v(a_j,b)\right]p_{ji}+\tau_i(a)-\tau_i(b)$$

and the claim follows by symmetry of the reward function v and the interaction probabilities p_{ij}. □

Note that the characterization of the potential function is in particular true in the case of degenerate interactions, i.e. where player interaction takes place on a deterministic network $\boldsymbol{g}$. Further it is "global" in the sense that it holds for every realization of types and every array of interaction probabilities $\boldsymbol{p}$. An interaction game is played as follows. First we fix the number of players $N \geq N^0 \geq 2$. Then each player receives his type τ_i independently of any other player. The distribution of types in the population depends on the population size N and is governed by a reference probability vector $\boldsymbol{q} := (q_1, \dots, q_K) \in \operatorname{int}\Delta(\Theta)$, which can be interpreted as the common prior of the players. A random type profile is an N-sequence of random variables $\tilde{\boldsymbol{\tau}}^{(N)} := (\tilde{\tau}_1^{(N)}, \dots, \tilde{\tau}_N^{(N)})$. A realization defines a type profile $\boldsymbol{\tau} \in \Theta^N$. Every realized type profile generates an (empirical) type distribution, denoted as $\boldsymbol{M}^N = (M_1^N, \dots, M_K^N)$, which is a point in $\Delta(\Theta)$, and defined as

$$M_k^N(\boldsymbol{\tau}) := \frac{1}{N}\sum_{i=1}^{N}\mathbb{1}_{\theta_k}(\tau_i)$$

for all $1 \leq k \leq K$. It is a *random element* of the set

$$\mathcal{L}_N := \{\boldsymbol{m} \in \Delta(\Theta) | (\exists \boldsymbol{\tau} \in \Theta^{\mathbb{N}}) : \boldsymbol{M}^N(\boldsymbol{\tau}) = \boldsymbol{m}\}.$$

Conversely, we can also consider the set of type realizations which result in a targeted distribution $\boldsymbol{m} \in \Delta(\Theta)$. Formally, let us define the *type class set*

$$\mathcal{T}^N(\boldsymbol{m}) := \{\boldsymbol{\tau} \in \Theta^{\mathbb{N}} | \boldsymbol{M}^N(\boldsymbol{\tau}) = \boldsymbol{m}\}.$$

The probability that a type profile $\boldsymbol{\tau} \in \mathcal{T}^N(\boldsymbol{m})$ is realized under $\boldsymbol{q}$ is therefore

$$\mathrm{P}_{\boldsymbol{q}}(\tilde{\boldsymbol{\tau}}^{(N)} = \boldsymbol{\tau}) = \prod_{k=1}^{K} q_k^{N m_k} \tag{2.4}$$

Finally we decide upon the inhomogeneous random graph (which could be a single network $\boldsymbol{g}$, recall remark 1). Players choose actions independently and are paid according to their utility function $U_i(\boldsymbol{a}, \boldsymbol{g}, \tau_i)$, where $\boldsymbol{g}$ is a realized network from the random graph $\mathcal{G}[N, \boldsymbol{p}]$. Let us close this section with two examples illustrating the possible applications of Bayesian interaction games.

Example 1 (Technology-choice game)**.** *Let us consider the following version of a technology choice game. There is a large population of players, where each player must make a choice between two competing technologies, say two competing operating systems for a PC, independently of the choices of the opponent players. For simplicity let us assume that the random graph is degenerate in the sense that all interactions take place on a fixed network* $\boldsymbol{g}$*. For illustration we only need to assume that players know their neighbors but not necessarily the complete topology of the network.*[13] *In a pairwise interaction the reward function is given by the payoff matrix*

$$v = \begin{pmatrix} 1 & 0 \\ 0 & 4 \end{pmatrix}.$$

There are two types of players in the population. The first type, θ_1*, has a strong preference for operating system* 1*, which means that his idiosyncratic preference may look like* $\theta_1(1) = 8$ *and* $\theta_1(2) = 0$*. The second type,* θ_2*, is indifferent to the operating systems so that we may set* $\theta_2(1) = \theta_2(2) = 1$*. The payoffs of an agent of type* θ_1 *using action 1 is* $U_i((1, \boldsymbol{a}_{-i}), \boldsymbol{g}, \theta_1) = d_i(\boldsymbol{g}) - d_{i,2}(\boldsymbol{g}) + 8$*, where* $d_{i,2}(\boldsymbol{g})$ *denotes the number of interactions of player i in which he meets players who choose the second operating system and* $d_i(\boldsymbol{g})$ *measures the total number of interactions of player i (i.e. his degree). Similarly, the payoff to player i from action* 2 *is* $U_i((2, \boldsymbol{a}_{-i}), \boldsymbol{g}, \theta_1) = 4d_{i,2}(\boldsymbol{g})$*. Adopting the second operating system, a choice contradicting the idiosyncratic preference, is a best response for i if and only if* $\frac{d_{i,2}(g)}{d_i(g)} > \frac{1}{5} + \frac{8}{5d_i(g)}$*. Hence, a player who has to coordinate his decision with many other agents, will put less weight on his idiosyncratic preference and more weight on the common reward function.*

Example 2 (Linear-quadratic games)**.** *The linear-quadratic game of Ballester et al. (2006) is a Bayesian interaction game with the partnership structure, if we restrict their utility function to a finite subset of the integers. Let* $A := \{0, 1, \ldots, n\}$*, and define the reward functions of the players as* $v(a, b) := \kappa ab$*.* κ *is a real parameter which defines whether the actions of the players are strategic substitutes* ($\kappa < 0$) *or strategic complements* ($\kappa > 0$)*. The type space is degenerate to a single point* $\Theta = \{\theta\}$*, where* $\theta(a) := a - \frac{1}{2}a^2$*. Hence, the utility function of player i takes the form*

$$U_i(\boldsymbol{a}, \boldsymbol{g}, \tau_i) \equiv U_i(\boldsymbol{a}, \boldsymbol{g}) = \kappa \cdot a_i \sum_{j=1}^{N} g_{ij} a_j + a_i - \frac{1}{2} a_i^2,$$

which is seen to be a particular version of the linear-quadratic game studied by Ballester et al. (2006), restricted to binary interactions and discrete action spaces.

[13]This may not suffice for players to compute a (Bayes) Nash equilibrium, a thing we deliberately don't do in this paper.

This game has an exact potential function given by[14]

$$V(\boldsymbol{a},\boldsymbol{g},\boldsymbol{\tau}) := \sum_{i=1}^{N} a_i - \frac{1}{2}\left(\sum_{i=1}^{N} a_i^2 - \kappa \sum_{i,j} g_{ij} a_i a_j\right),$$

which is easily seen to be of the form (2.3).

2.2 Co-evolution of networks and play

Our dynamic model builds on a class of continuous-time Markov processes, which we have called in Staudigl (2010b) *co-evolutionary processes of networks and play*. A co-evolutionary model is a time-homogeneous Markov jump process $\{X^{\beta,\tau,N}(t)\}_{t\geq 0}$ taking values on the finite state space $\Omega^N = A^N \times \mathcal{G}[N]$. An element of this space is denoted by $\omega = (\boldsymbol{a},\boldsymbol{g})$ and is called a *population state*. The process is parameterized by three variables:

- β the behavioral noise,
- $\boldsymbol{\tau}$ the realized type profile, and
- N the population size.

For every population state ω we can define projection mappings $\alpha : \Omega^N \to A^N$ and $\gamma : \Omega^N \to \mathcal{G}[N]$. The evolutionary process must be specified for the events of action revision, link creation and link destruction. These events are modeled as conditionally independent random processes, each defined by a collection of rate functions so that the dynamics is described by the infinitesimal generator $(\eta^{\beta,\tau,N}_{\omega,\omega'})_{\omega,\omega'\in\Omega^N}$. The *rate* of a transition from a state ω to some other state ω' is defined as

$$\eta^{\beta,\tau,N}_{\omega,\omega'} = \begin{cases} \ell^{i,\beta}_a(\omega|\tau_i) & \text{if } \omega' = ((a,\alpha_{-i}(\omega)),\gamma(\omega)), \\ c^{\beta,N}_{ij}(\alpha(\omega),\boldsymbol{\tau}) & \text{if } \omega' = (\alpha(\omega),\gamma(\omega)\oplus(i,j)), \\ \xi^{\beta,N}_{ij}(\boldsymbol{\tau}) & \text{if } \omega' = (\alpha(\omega),\gamma(\omega)\ominus(i,j)). \end{cases}$$

The meaning of these rate functions will be clarified below. Sample paths of the process are characterized by random jump times $\{J_n\}_{n\geq 0}$. The chain sampled at the jump times has values $X^{\beta,\tau,N}(J_n) = X^{\beta,\tau,N}_n$, and appears formally as a discrete-time Markov chain.[15] Each J_{n+1} is measurable with respect to the filtration $\mathcal{F}_n = \sigma\left(\{J_0,\ldots,J_n\},\{X^{\beta,\tau,N}_0,\ldots,X^{\beta,\tau,N}_n\}\right), n\geq 0$. We set $J_0 = 0$ and $X^{\beta,\tau,N}(0) = X^{\beta,\tau,N}_0$, which is an arbitrarily chosen

[14]See also König and Staudigl (2010) for the case with continuous action spaces, and further investigations of this model in the light of network formation algorithms.

[15]This is called the embedded jump chain of the process. See e.g. Stroock (2005).

network g and an action profile $\boldsymbol{a} \in A^N$. The rates will be constructed such that $J_n < \infty$ for all $n \geq 1$. For all $\omega \in \Omega^N$ let us define

$$\eta_{\omega}^{\beta,\tau,N} := \sum_{\omega' \in \Omega^N \setminus \{\omega\}} \eta_{\omega,\omega'}^{\beta,\tau,N} \in (0, \infty).$$

This measures the rate of leaving state ω. For all $t \geq 0$ transitions of $X^{\beta,\tau,N}(t)$ are described by

$$\begin{aligned}\mathbb{P}\left(X_{J_{n+1}}^{\beta,\tau,N} = \omega', J_{n+1} - J_n > t | \mathcal{F}_n\right) &= \mathbb{P}\left(X_{J_{n+1}}^{\beta,\tau,N} = \omega', J_{n+1} - J_n > t | X_{J_n}^{\beta,\tau,N}\right) \\ &= \exp(-t\eta_{\omega}^{\beta,\tau,N}) \frac{\eta_{\omega,\omega'}^{\beta,\tau,N}}{\eta_{\omega}^{\beta,\tau,N}}\end{aligned}$$

We turn now to the interpretation of the rate functions. If not stated differently, these are all the properties we assume on the nature of the transition rates.

Action adjustment: Each agent is endowed with an independent Poisson alarm clock of intensity 1, so that the conditional probability that agent i receives an action revision opportunity is $1/N$. In case of such an event we assume that the agent switches to action $a \in A$ with probability determined by the log-linear response function

$$(\forall a \in A) : \ell_a^{i,\beta}(\omega|\tau_i) = \frac{\exp\left(\beta^{-1} U_i((a, \alpha_{-i}(\omega)), \gamma(\omega), \tau_i)\right)}{\sum_{b \in A} \exp\left(\beta^{-1} U_i((b, \alpha_{-i}(\omega)), \gamma(\omega), \tau_i)\right)}$$

The *rate* of the transition $\omega \to \omega' = ((a, \alpha_{-i}(\omega)), \gamma(\omega))$ is therefore

$$\eta_{\omega,\omega'}^{\beta,\tau,N} = \ell_a^{i,\beta}(\omega|\tau_i). \tag{2.5}$$

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*.

Definition 2. *An* attachment mechanism $\boldsymbol{C}^{\beta,\tau,N}$ *is a collection of functions* $c_{ij}^{\beta,N}(\cdot, \boldsymbol{\tau}) : A^N \to \mathbb{R}_+$, *such that for all* $\beta > 0$ *the conditions*

(B) $\sup_{\boldsymbol{a} \in A^N} \sum_{j>i} c_{ij}^{\beta,N}(\boldsymbol{a}, \boldsymbol{\tau}) < \infty$, *and*

(Sym) $(\forall i, j \in [N])(\forall \boldsymbol{a} \in A^N) : c_{ij}^{\beta,N}(\boldsymbol{a}, \boldsymbol{\tau}) = c_{ji}^{\beta,N}(\boldsymbol{a}, \boldsymbol{\tau})$

are satisfied.

The rate of the transition $\omega \to \omega' = (\alpha(\omega), \gamma(\omega) \oplus (i, j))$ is therefore

$$\eta_{\omega,\omega'}^{\beta,\tau,N} = (1 - \gamma_{ij}(\omega)) c_{ij}^{\beta,N}(\alpha(\omega), \boldsymbol{\tau}). \tag{2.6}$$

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 are described via a *volatility mechanism*.

Definition 3. *A* volatility mechanism $\Xi^{\beta,\tau,N}$ *is a collection of positive functions* $(\xi_{ij}^{\beta,N}(\boldsymbol{\tau}))_{i,j\in[N]}$*, where each* $\xi_{ij}^{\beta,N}(\boldsymbol{\tau})$ *depends on the type realizations of players i and j alone. It is called an* admissible volatility mechanism *if every function satisfies conditions* (B) *and* (Sym) *of definition 2, and additionally*

(SNB) *For all* $\boldsymbol{\tau} \in \Theta^N$ *and all* $i, j \in [N]$*, it satisfies* small noise boundedness*, i.e. the bound*

$$0 < \xi_-^N \leq \xi_{ij}^{\beta,N}(\boldsymbol{\tau}) \leq c \exp(f^\beta_{\tau_i,\tau_j})$$

holds uniformly in $\beta \in (0, \infty)$ *for some function* $f^\beta : \Theta^2 \to \mathbb{R}_+$ *that satisfies* $\lim_{\beta\to 0} \beta f^\beta_{\theta_k,\theta_l} = 0$ *for all* $1 \leq k, l \leq K$.

(LPB) *For all* $\boldsymbol{\tau} \in \Theta^{\mathbb{N}}$ *and all* $i, j \in [N]$*, it* satisfies large population boundedness

$$0 < \xi_-^\beta \leq \lim_{N\to\infty} \left\{ \inf_{j>i} \xi_{ij}^{\beta,N}(\boldsymbol{\tau}) \right\} \leq \lim_{N\to\infty} \left\{ \sup_{j>i} \xi_{ij}^{\beta,N}(\boldsymbol{\tau}) \right\} \leq \xi_+^\beta < \infty$$

The numbers $\xi_{ij}^{\beta,N}(\boldsymbol{\tau})$ are the volatility rates and are fixed scalars once the types of the players are fixed. The rate of the transition $\omega \to \omega' = (\alpha(\omega), \gamma(\omega) \ominus (i, j))$ is

$$\eta^{\beta,\tau,N}_{\omega,\omega'} = \gamma_{ij}(\omega)\xi_{ij}^{\beta,N}(\boldsymbol{\tau}). \tag{2.7}$$

3 Asymptotic properties of the process

A co-evolutionary process of networks and play describes a random dynamics on the joint space of action profiles and networks. In the following sections we give a complete description of the long-run properties of this process.

3.1 *Inhomogeneous random graphs*

A first characterization of the invariant regime of the dynamics is obtained when we condition the process to the $\boldsymbol{a}$-section of the state space, by which we mean the set $\Omega^N_{\boldsymbol{a}} = \{\omega \in \Omega^N | \alpha(\omega) = \boldsymbol{a}\}$. All population states in this set differ only in the interaction network. Thus, if we constrain the process

$\{X^{\beta,\tau,N}(t)\}_{t\geq 0}$ to take values in this set only, we obtain an ergodic random graph process $G^{\beta,\tau,N} = \{G^{\beta,\tau,N}(t)\}_{t\geq 0}$ whose generator describes a multi-type birth-death process with "birth rates" of the link (i,j) given by the deterministic scalar $c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})$, and "death-rates" $\xi_{ij}^{\beta,N}(\boldsymbol{\tau})$. In this section we make no specific assumptions on these quantities, beside that they are admissible. The main result of this section is the following characterization theorem of the class of random graphs generated by the co-evolutionary process. Note that this characterization theorem is rather general since we do not impose any specific functional form on the attachment and volatility mechanism, despite that they are admissible.[16]

Theorem 3.1. *The random graph process $G^{\beta,\tau,N}$ with admissible attachment mechanism $\boldsymbol{C}^{\beta,\tau,N}$ and volatility mechanism $\Xi^{\beta,\tau,N}$ has a unique invariant graph measure*

$$\mu^{\beta,\tau,N}(\omega|\Omega_{\boldsymbol{a}}^N) = \prod_{i=1}^{N}\prod_{j>i} p_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})^{\gamma_{ij}(\omega)}\left(1 - p_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})\right)^{1-\gamma_{ij}(\omega)}, \tag{3.1}$$

which generates the interaction model $\mathcal{G}[N,(p_{ij}^{\beta,\tau,N}(\boldsymbol{a}))_{j>i}]$ with interaction probabilities

$$p_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) = \frac{c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})}{c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) + \xi_{ij}^{\beta,N}(\boldsymbol{\tau})}. \tag{3.2}$$

Proof. See Appendix A. □

Hence, for each pair of players we can derive the conditional probability that they will be matched in the long run, given that we know the action and the type profile.

The IHRG in theorem 3.1 respects the labels of the individual players and might be useful in models where the labels of the players are important, or there is a subgroup of agents that influence the behavior of all other agents in a non-negligible way, as illustrated by the following example.

Example 3 (Almost a star)**.** *Consider the co-evolutionary process of networks and play with admissible attachment mechanism*

$$c_{1j}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) = N^{1+\epsilon}, \quad \forall j > 1,$$
$$c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) = N^{-\epsilon}, \quad \forall i,j > 1,$$

[16] However, in many applications one would like to formulate the volatility rates as functions of the rewards of the players (see e.g. Jackson and Watts, 2002, for such a model). In Staudigl (2010a) we allow for such a scenario and prove that a co-evolutionary process still generates inhomogeneous random graphs such as in Theorem 3.1.

where $\epsilon > 0$ is a given number, and admissible volatility mechanism

$$\xi_{ij}^{\beta,N}(\boldsymbol{\tau}) \equiv \xi > 0.$$

This generates a random graph process which is independent of the types and the actions of the players. However, there is asymmetry in the process since there is one distinguished player, player 1, to whom all other players want to be connected with very high probability. Indeed, plugging into formula (3.2), we see that

$$p_{1j}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) = \frac{N^{1+\epsilon}}{N^{1+\epsilon}+\xi} \quad ,\forall j > 1$$
$$p_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) = \frac{N^{-\epsilon}}{N^{-\epsilon}+\xi} \quad ,\forall i,j > 1.$$

For N sufficiently large player 1 will be incident to almost all edges and can therefore be called a star-player in the network. Edges not incident to player 1 appear with vanishing probability as N gets large.

If individual players are small it might be of more economic significance to model the co-evolutionary process by dropping the names of the players while focusing on the observable characteristics of the players: their actions and their types. We would like to do so in a manner that allows us to give a comprehensive characterization of the long-run properties of the co-evolutionary process. Therefore we need to impose more structure on the random graph process. A specifically interesting functional form for an attachment mechanism is the type-independent exponential flip rate

$$c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) = \frac{2}{N}\exp(v(a_i,a_j)/\beta), \tag{3.3}$$

which arises naturally in a random utility formulation of active search for new interaction partners.[17]

Assumption 1. *Unless indicated differently we henceforth assume that the attachment mechanism $\boldsymbol{C}^{\beta,\tau,N}$ is given by the exponential flip rates (3.3).*

To capture the types of the players in the linking probabilities we introduce the concept of a semi-anonymous volatility mechanism.

Definition 4. *We call an admissible volatility mechanism $\Xi^{\beta,\cdot,N}$* semi-anonymous *if for all $N \geq 2, \boldsymbol{\tau} \in \Theta^N$, and all agents $i,j \in [N]$, it is true that*

$$\xi_{ij}^{\beta,N}(\boldsymbol{\tau}) = \xi_{kl}^{\beta,N} \text{ whenever } \tau_i = \theta_k \text{ and } \tau_j = \theta_l.$$

[17] See Staudigl (2010b) for the details of the construction.

Remark 2. *Under the assumption that our feasible attachment mechanism is given by (3.3), the reward function is already incorporated into the random graph. Effects coming from the types of the players are incorporated by the volatility mechanism. A good reason why this separation of effects may be useful can be given in light of potential statistical applications. Once we know the reward functions, the actions and the interaction probabilities of the players, one can identify the volatility rates by estimating the likelihood ratios of interactions among players. To see this, consider the odds-ratio*

$$\frac{p_{ij}^{\beta,N}(\boldsymbol{a})}{1-p_{ij}^{\beta,N}(\boldsymbol{a})} = \frac{c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})}{\xi_{ij}^{\beta,N}}. \tag{3.4}$$

Thus, the model is theoretically and empirically flexible with respect to the volatility mechanism, once we agree on a functional form on the attachment mechanism.

Given that the attachment mechanism is of the form (3.3) and the volatility mechanism is semi-anonymous we can define for all $1 \le k,l \le K$ and $a,b \in A$ the scalars

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

and the symmetric matrices

$$\boldsymbol{\varphi}_{kl}^{\beta,N} := \left(\varphi_{kl}^{\beta,N}(a,b)\right)_{(a,b)\in A^2}, \ \boldsymbol{\varphi}^{\beta,N} = \left(\boldsymbol{\varphi}_{kl}^{\beta,N}\right)_{1\le k,l\le K}.$$

In terms of these quantities the edge-success probabilities can be written as

$$p_{kl}^{\beta,N}(a,b) = \frac{\varphi_{k,l}^{\beta,N}(a,b)}{N\beta+\varphi_{kl}^{\beta,N}(a,b)}. \tag{3.5}$$

We see that the interaction model can in fact be characterized without conditioning on any $\boldsymbol{a}$-section at all. The matrices

$$\boldsymbol{p}_{kl}^{\beta,N} := \left(p_{kl}^{\beta,N}(a,b)\right)_{(a,b)\in A^2} \in \mathbb{R}_+^{n\times n}, \ \boldsymbol{p}^{\beta,N} := \left(\boldsymbol{p}_{kl}^{\beta,N}\right)_{1\le l,k\le K} \tag{3.6}$$

always give a complete characterization of the long-run interaction model of the society. Moreover, once we know the actions of the players and the reward function, it should be possible to identify the volatility rates from the likelihood ratio (3.4). Given these functional assumptions we can collect the following facts on the interaction probabilities.

Observation 1. *(i) Each matrix $\boldsymbol{p}_{kl}^{\beta,N}$ is symmetric and for $\beta > 0$ positive.*

(ii) Under assumption (LPB) for all $1 \leq k, l \leq K$ and $a, b \in A$ we have $\lim_{N\to\infty} p_{kl}^{\beta,N}(a,b) = 0$.

(iii) Under assumption (SNB) for all $1 \leq k, l \leq K$ and $a, b \in A$ we have

$$\lim_{\beta\to 0} p_{kl}^{\beta,N}(a,b) \begin{cases} = 1 & \text{if } v(a,b) > 0 \\ \leq \frac{2}{2+N\xi_-^N} & \text{if } v(a,b) = 0 \\ = 0 & \text{if } v(a,b) < 0 \end{cases}$$

Item (i) just says that the generated random graph is undirected and the interaction pattern is irreducible, in the sense that the network will almost surely have a single connected component. Item (ii) in turn says that all interaction probabilities are $o(N)$. In particular this rules out the presence of a star player such as in example 3. Item (iii) investigates the small noise behavior of the interaction probabilities. To verify that the stated condition holds consider the odds-ratios of interaction probabilities

$$\frac{p_{kl}^{\beta,N}(a,b)}{1 - p_{kl}^{\beta,N}(a,b)} = \frac{2\exp(v(a,b)/\beta)}{N\xi_{kl}^{\beta,N}}.$$

By (SNB) we can upper and lower bound this ratio as

$$\frac{2}{Nc}\exp(v(a,b)/\beta - f_{kl}^{\beta}) \leq \frac{p_{kl}^{\beta,N}(a,b)}{1 - p_{kl}^{\beta,N}(a,b)} \leq \frac{2\exp(v(a,b)/\beta)}{N\xi_-^N}.$$

We see that if $v(a,b) < 0$ the lower and the upper bounds go to 0 and consequently $p_{kl}^{\beta,N}(a,b) \overset{\beta\to 0}{\to} 0$. If $v(a,b) > 0$ the lower bound goes to ∞ implying that $p_{kl}^{\beta,N}(a,b) \overset{\beta\to 0}{\to} 1$. Finally, if $v(a,b) = 0$ then it must be true that $p_{kl}^{\beta,N}(a,b) \in [0, \frac{2}{2+N\xi_-^N}]$. This shows that in the small noise limit interaction networks are almost deterministic and players interact with positive probability if and only if they obtain a non-negative reward. This holds for arbitrary large (but finite) population size N.

Remark 3. *In terms of play this has the consequence that the network formation process clusters players together that play actions which yield high mutual reward. The observation is actually troublesome, since it makes it very likely that in the limit of infinitely many players the utility functions of the players are not defined anymore.*[18] *This follows because each agents' utility function is the sum of all rewards. As $\beta \to 0$ there is the possibility that interaction probabilities between*

[18] This problem is also addressed in the general study of Horst and Scheinkman (2006) concerning existence and uniqueness of equilibria in systems of general (random) social interactions.

two player groups are not scaled by population size, and so the expected number of interactions will go to infinity as N grows large. This calls for our attention to consider the small noise limit for finite populations and the large population limit for positive noise as two separate parameters, which should not be taken to their limit jointly. One could perform a double limit analysis by picking a sequence $(\beta_N, N) \to (0, \infty)$*, where N goes to infinity at a much faster rate than* β*. The necessary speed regulation will be exponential. We have not pursued such an analysis yet.*

Finally, we would like to emphasize that inhomogeneous random graphs can be used in a quite satisfactory way to capture the phenomenon of homophily.[19] Motivated by Jackson (2008a), suppose we have a bounded and monotonically increasing function $w : \mathbb{R} \to \mathbb{R}_+$ such that $\xi_{kl}^{\beta,N} = w(||\theta_k - \theta_l||)$, where we interpret the types of the players as points in $\mathbb{R}^n$ and $||\cdot||$ is some norm on this vector space. Assume that this specification of a volatility mechanism satisfies (SNB) and (LPB). This volatility mechanism has the interpretation that larger differences in the idiosyncratic preferences of the players lead to a higher rate of link destruction. In equilibrium this manifests itself into a lower probability of interaction. This induces a natural ordering on odds-ratios. To illustrate this let us consider three players h, i, j with $a_i = a$ and $a_j = a_h = b$. Suppose that player i is of type θ_k while player j is of type θ_l and player h is of type θ_m. If $||\theta_k - \theta_l|| < ||\theta_k - \theta_m||$ then it follows that

$$\frac{p_{kl}^{\beta,N}(a,b)}{p_{km}^{\beta,N}(a,b)} = \frac{2\exp(v(a,b)/\beta) + \xi_{km}^{\beta,N}}{2\exp(v(a,b)/\beta) + \xi_{kl}^{\beta,N}} > 1.$$

Hence, although the interaction between player i and j would yield the same reward as the interaction between player i and h, the likelihood that i interacts with player j is relatively higher than the interaction with player h, and this can be explained only through the difference in idiosyncratic preferences.

3.2 *Joint invariant distribution*

In this section we provide a closed form expression for the unique invariant distribution of the co-evolutionary process under the assumption that the attachment mechanism is of the exponential form (3.3). The volatility mechanism is admissible, but need not be semi-anonymous for the results to hold. Let $\mathsf{M}(\Omega^N)$ denote the set of probability measures on the

[19] On the role of homophily in social networks we refer to Currarini et al. (2009) and the references therein.

finite space Ω^N endowed with its σ-algebra 2^{Ω^N}. The main result of this section is the following theorem, which specifies for every type sequence $\boldsymbol{\tau}$ a unique invariant measure $\mu^{\beta,\boldsymbol{\tau},N} \in \mathsf{M}(\Omega^N)$ of the Markov process $\{X^{\beta,\boldsymbol{\tau},N}(t)\}_{t\geq 0}$. Recall that the potential function of the interaction game is denoted by $V(\omega,\boldsymbol{\tau})$.

Theorem 3.2. *The unique invariant distribution of the Markov jump process* $\{X^{\beta,\boldsymbol{\tau},N}(t)\}_{t\geq 0}$ *is the Gibbs measure*

$$\mu^{\beta,\boldsymbol{\tau},N}(\omega) = \frac{\exp(\beta^{-1}H^{\beta,N}(\omega,\boldsymbol{\tau}))}{\sum_{\omega'\in\Omega^N}\exp(\beta^{-1}H^{\beta,N}(\omega',\boldsymbol{\tau}))} \tag{3.7}$$

$$= \frac{\mu_0^{\beta,\boldsymbol{\tau},N}(\omega)\exp\left(\beta^{-1}V(\omega,\boldsymbol{\tau})\right)}{\sum_{\omega'\in\Omega^N}\mu_0^{\beta,\boldsymbol{\tau},N}(\omega')\exp\left(\beta^{-1}V(\omega',\boldsymbol{\tau})\right)} \tag{3.8}$$

where for all $\omega \in \Omega^N$

$$H^{\beta,N}(\omega,\boldsymbol{\tau}) := V(\omega,\boldsymbol{\tau}) + \beta\log\mu_0^{\beta,\boldsymbol{\tau},N}(\omega), \tag{3.9}$$

$$\mu_0^{\beta,\boldsymbol{\tau},N}(\omega) := \prod_{i=1}^{N}\prod_{j>i}\left(\frac{2}{N\xi_{ij}^{\beta,N}(\boldsymbol{\tau})}\right)^{\gamma_{ij}(\omega)}. \tag{3.10}$$

Proof. See Appendix A. □

The function $H^{\beta,N}$ provides a complete description of the invariant distribution weights. The Gibbs measure is game-theoretically interesting, since it shows that we can formulate the invariant distribution as a product of a "graph weight function" $\mu_0^{\beta,\boldsymbol{\tau},N}$, and the term $\exp(\beta^{-1}V(\omega,\boldsymbol{\tau}))$.[20] The second term is only driven by the value of the potential function at the population state ω, which is, recall eq. (2.3), given by the welfare measure

$$V(\omega,\boldsymbol{\tau}) = \sum_{j>i} v(\alpha_i(\omega),\alpha_j(\omega))\gamma_{ij}(\omega) + \sum_{i=1}^{N}\tau_i(\alpha_i(\omega)).$$

The graph weight function $\mu_0^{\beta,\boldsymbol{\tau},N}$ collects the effects arising from the volatility of the interaction structure. It is clear that the exact value of these weights depends on the realization of the types of the agents in the population. As types are random realizations the resulting invariant measure is a *random element* of the set of measures $\mathsf{M}(\Omega^N)$. Each measure $\mu^{\beta,\boldsymbol{\tau},N}$ is realized with the probability that the population is characterized by the

[20] In this sense one can view the stationary distribution (3.7) as a perturbation of the invariant measure found by Blume (1993; 1997) in the context of interaction games on fixed networks.

type profile $\boldsymbol{\tau}$. Hence, we can define the joint probability measure on the extended state space $\Omega^N \times \Theta^N$ as

$$\mu^{\beta,N}(\omega, \boldsymbol{\tau}) := \mathrm{P}_{\boldsymbol{q}}\left(\tilde{\boldsymbol{\tau}}^{(N)} = \boldsymbol{\tau}\right)\mu^{\beta,\boldsymbol{\tau},N}(\omega). \tag{3.11}$$

4 Limit behavior of the invariant distribution

The transition rates of our co-evolutionary process depend on two exogenous parameters, β and N. In both cases it can be shown that the invariant distribution has a tendency to single out certain subsets of the state space when these two parameters are taken to their respective limits, i.e. $\beta \to 0$ and/or $N \to \infty$. In view of Remark 3 we treat these two limits as complementary ways to study the concentration of the invariant distribution, and *a fortiori* two complementarity ways to obtain exact results concerning equilibrium selection.

4.1 Small Noise Limit

A standard result in the study of potential games with log-linear response functions is that the family of stationary distributions $\{\mu^{\beta,\boldsymbol{\tau},N}\}_{\beta>0}$ concentrates on the set of potential maximizers.[21] In order to examine the limiting invariant distribution we employ the following notion of stochastic stability.[22]

Definition 5. *A population state $\omega \in \Omega^N$ is called* stochastically stable in the small noise limit *if*

$$\lim_{\beta\to 0} \beta \log \mu^{\beta,\boldsymbol{\tau},N}(\omega) = 0.$$

As discussed in the Introduction the argmax of a potential function serves as a natural device for equilibrium selection in games. Moreover points in this set possess in our model the important property that they are social welfare maximizing, taking the sum of utilities of the agents as a welfare measure (recall the definition of the potential function in eq. (2.3)). It is therefore important to have sufficient conditions under which the potential maximizers retain their robustness property. A characterization in this direction is given by the following lemma.

[21] See Blume (1997) and Sandholm (2010, ch. 12) for the most general results in this direction.

[22] The notion of stochastic stability we employ in this paper is weaker than the one of Young (1993) or Ellison (2000). We refer to Sandholm (2010) for a more thorough discussion of the differences between the two definitions.

Lemma 4.1. *Fix $N \geq 2$ and an arbitrary type profile $\boldsymbol{\tau} \in \Theta^N$. Then*

$$\lim_{\beta \to 0} \max_{\omega \in \Omega^N} |H^{\beta,N}(\omega, \boldsymbol{\tau}) - V(\omega, \boldsymbol{\tau})| = 0 \tag{4.1}$$

if and only if

$$\lim_{\beta \to 0} \max_{\omega \in \Omega^N} \beta |\log \mu_0^{\beta,\tau,N}(\omega)| = 0. \tag{4.2}$$

Proof. This follows immediately from eq. (3.9). □

We see that if the co-evolutionary dynamics has an admissible volatility mechanism, or only satisfies $(B), (Sym)$ and (SNB), then the potential function dominates the perturbation coming from the graph weight function $\mu_0^{\beta,\tau,N}$ in the small noise limit. The main result of this section is the following theorem, which proves that the family of invariant measures $\{\mu^{\beta,\tau,N}\}_{\beta>0}$ satisfies a *large deviations principle*. This result gives us not only the information that the invariant distribution concentrates (on a logarithmic scale) on certain subsets of the state space (to be precise minimizers of a *rate function* to be specified in the statement of the theorem), but moreover it gives us an estimate on the *rate* of this convergence (again at a logarithmic scale).

Theorem 4.1. *If $\Xi^{\beta,\tau,N}$ is an admissible volatility mechanism then the family of invariant measures $\{\mu^{\beta,\tau,N}\}_{\beta>0}$ satisfies a large deviations principle with rate function $R(\omega, \boldsymbol{\tau}) := \max_{\omega' \in \Omega^N} V(\omega', \boldsymbol{\tau}) - V(\omega, \boldsymbol{\tau})$, i.e.*

$$\lim_{\beta \to 0} \beta \log \mu^{\beta,\tau,N}(\omega) = -R(\omega, \boldsymbol{\tau}) \tag{4.3}$$

for all $\omega \in \Omega^N$.

Proof. Recall from equation (3.7) that

$$\mu^{\beta,\tau,N}(\omega) = \frac{\exp(\beta^{-1} H^{\beta,N}(\omega, \boldsymbol{\tau}))}{\sum_{\omega' \in \Omega^N} \exp(\beta^{-1} H^{\beta,N}(\omega', \boldsymbol{\tau}))}$$

for all $\omega \in \Omega^N$. Further, if the volatility mechanism is admissible it satisfies in particular (SNB). Thus, we know from Lemma 4.1 that the Hamiltonian function converges uniformly to the game potential function as $\beta \to 0$. Thus,

$$\begin{aligned}
-\lim_{\beta \to 0} \beta \log \mu^{\beta,\tau,N}(\omega) &= \max_{\omega' \in \Omega^N} \lim_{\beta \to 0} H^{\beta,N}(\omega', \boldsymbol{\tau}) - \lim_{\beta \to 0} H^{\beta,N}(\omega, \tau) \\
&= \max_{\omega' \in \Omega^N} V(\omega', \boldsymbol{\tau}) - V(\omega, \boldsymbol{\tau}) \\
&= R(\omega, \boldsymbol{\tau}).
\end{aligned}$$

for all $\omega \in \Omega^N$. □

An immediate consequence of this theorem are the following two observations:

Corollary 4.1. *Let $\Xi^{\beta,\boldsymbol{\tau},N}$ be an admissible volatility mechanism.*

(i) Let

$$\Omega^{*,N}(\boldsymbol{\tau}) := \{\omega \in \Omega^N | \lim_{\beta\to 0} \beta \log \mu^{\beta,\boldsymbol{\tau},N}(\omega) = 0\}$$

denote the set of stochastically stable states in the small noise limit for the realized type profile $\boldsymbol{\tau}$. We have the equivalence

$$\Omega^{*,N}(\boldsymbol{\tau}) = \arg\max_{\omega\in\Omega} V(\omega, \boldsymbol{\tau}).$$

(ii) The invariant distribution is exponentially tight in the sense that for every $\epsilon > 0$ there exists a subset $X_\epsilon \subseteq \Omega^N$ such that

$$\lim_{\beta\to 0} \beta \log \mu^{\beta,\boldsymbol{\tau},N}(X_\epsilon) < -\epsilon$$

Proof. Only the second point requires a proof. Denote the level sets of the rate function R as $L_R(\epsilon) := \{\omega \in \Omega^N | R(\omega, \boldsymbol{\tau}) \leq \epsilon\}$. For all $\epsilon > 0$ these sets are nonempty, since always $\Omega^{*,N}(\boldsymbol{\tau}) \subseteq L_R(\epsilon)$, with equality as $\epsilon \downarrow 0$. Fix an $\epsilon > 0$ and consider the set $X_\epsilon := \Omega^N \setminus L_R(\epsilon)$. Then $R_{X_\epsilon}(\boldsymbol{\tau}) := \min_{\omega\in X_\epsilon} R(\omega, \boldsymbol{\tau}) > \epsilon$. We claim that

$$\begin{aligned}
\lim_{\beta\to 0} \beta \log \mu^{\beta,\boldsymbol{\tau},N}(X_\epsilon) &= \lim_{\beta\to 0} \beta \log \left(\sum_{\omega\in X_\epsilon} \mu^{\beta,\boldsymbol{\tau},N}(\omega) \right) \\
&= \max_{\omega\in X_\epsilon} \lim_{\beta\to 0} \beta \log \mu^{\beta,\boldsymbol{\tau},N}(\omega) \\
&= -\min_{\omega\in X_\epsilon} R(\omega) < -\epsilon.
\end{aligned}$$

The argument for this is as in the proof of Theorem 4.1. We have

$$\sum_{\omega\in X_\epsilon} \mu^{\beta,\boldsymbol{\tau},N}(\omega) = \exp(-\beta^{-1} R_{X_\epsilon}(\boldsymbol{\tau})) B^{\beta,N}_{X_\epsilon}(\boldsymbol{\tau}) r_{X_\epsilon}(\beta)$$

for some functions $B^{\beta,N}_{X_\epsilon}, r_{X_\epsilon}(\beta)$. Taking logarithms on both sides and multiplying by β shows that

$$\beta \log \left(\sum_{\omega\in X_\epsilon} \mu^{\beta,\boldsymbol{\tau},N}(\omega) \right) = -R_{X_\epsilon}(\boldsymbol{\tau}) + \beta \log B^{\beta,N}_{X_\epsilon}(\boldsymbol{\tau}) + \beta \log r_{X_\epsilon}(\beta)$$

and it is easy to see that the left hand side is $-R_{X_\epsilon}(\boldsymbol{\tau})(1 + o(1))$, for $\beta \to 0$. $\square$

Theorem 4.1 is an extension of a result of Staudigl (2010b). It shows that in the small noise limit and for every type profile $\boldsymbol{\tau} \in \Theta^N$ the invariant distribution concentrates on the set of potential maximizers. The only crucial assumption on the nature of the volatility rates needed for this result to hold is (SNB). The message of Corollary 4.1 is that stochastically stable states are global maximizers of the potential function. Part (ii) of the Corollary says that the family of measure $\{\mu^{\beta,\tau,N}\}_{\beta>0}$ puts arbitrary small weight on certain subsets of the state space.

4.1.1 Game theoretic interpretation of the invariant distribution

Before studying the large population case, we provide an interesting game-theoretic interpretation of the invariant distribution $\mu^{\beta,\tau,N}$. In the dynamic model we assumed that players always play pure actions and we would like to know whether the actions chosen by the players are individually rational in some sense. This requires a notion of equilibrium in actions. It turns out that the right notion of equilibrium in actions is Aumann's *correlated equilibrium* (Aumann, 1987). We take the state space Ω^N as the set of possible states of the world. The common prior of the players is $\mu^{\beta,\tau,N}$, which for simplicity is denoted by μ. The information partition $\mathcal{P}_i$ of player i, consists of the sets $P_i(a) := \{\omega \in \Omega^N | \alpha_i(\omega) = a\}, a \in A$. A *strategy* of player i is a map $s_i : \Omega^N \to A$ that is measurable with respect to his information partition $\mathcal{P}_i$, i.e. $s_i(\omega) = s_i(\omega') = a$, whenever $\omega, \omega' \in P_i(a)$.[23] A profile of strategies is denoted as $\boldsymbol{s} = (s_i)_{i\in[N]}$. The interpretation is the following; Suppose the evolutionary process has settled to a dynamic equilibrium. Now we replicate the population by introducing N players who have the types $(\tau_1, \dots, \tau_N)$ and the information structure $\langle \Omega^N, \mu, \mathcal{P}_i \rangle$. Suppose we recommend each player to follow the strategy $s_i(\omega) = \alpha_i(\omega), \forall \omega \in \Omega^N$. Would player i follow that recommendation?
Since for positive noise the measure μ has a full support, all states of the world are realized with a positive probability. However, in general not all strategies (i.e. functions that are measurable w.r.t. $\mathcal{P}_i$) will be "good" recommendations for a player. But we know from Theorem 4.1 that in the small noise limit only states of the world appear with non-vanishing probability which are potential maximizers. It is not surprising that the players' actions are compatible on this set in the sense of Nash equilibrium.

Lemma 4.2. *Let $\boldsymbol{\tau} \in \Theta^N$ be an arbitrary type-profile. For all $i \in [N]$ consider the strategy $s_i(\omega) = \alpha_i(\omega), \omega \in \Omega^N$. The profile $\boldsymbol{s}$ is a Nash equilibrium in*

[23] Technically, we would need to index this function with the type of the player. We omit this dependence for notational simplicity.

actions at all states $\omega \in \Omega^{*,N}(\boldsymbol{\tau})$.

Proof. The proposed mappings s_i are clearly $\mathcal{P}_i$ measurable, i.e. they are strategies. Fix an arbitrary player $i \in [N]$ and let $\hat{s}_i$ be some other strategy. At a fixed $\omega \in \Omega^N$ the deviation payoffs of player i are

$$\begin{aligned}&U_i[\boldsymbol{s}(\omega), \gamma(\omega), \tau_i] - U_i[(\hat{s}_i(\omega), \boldsymbol{s}_{-i}(\omega)), \gamma(\omega)), \tau_i]\\&= V((\boldsymbol{s}(\omega), \gamma(\omega)), \boldsymbol{\tau}) - V((\hat{s}_i(\omega), \boldsymbol{s}_{-i}(\omega), \gamma(\omega)), \boldsymbol{\tau})\end{aligned}$$

by definition of the potential function (2.3). This expression is non-negative on $\Omega^{*,N}(\boldsymbol{\tau})$. □

To get a global characterization of individually rational actions, we consider the following notion of equilibrium.[24]

Definition 6. *The quadruple $\left\langle \Omega^N, \mu^{\beta,\boldsymbol{\tau},N}, (\mathcal{P}_i)_{i\in[N]}, (s_i)_{i\in[N]} \right\rangle$ is a (β, ρ)-correlated equilibrium if for every $\rho > 0$ there exists a $\beta > 0$, such that for all $i \in [N]$ and all strategies $\hat{s}_i$ we have*

$$\sum_{\omega\in\Omega^N} \mu^{\beta',\boldsymbol{\tau},N}(\omega) U_i(\boldsymbol{s}(\omega), \gamma(\omega), \tau_i) \geq \sum_{\omega\in\Omega^N} \mu^{\beta',\boldsymbol{\tau},N}(\omega) U_i[(\hat{s}_i(\omega), s_{-i}(\omega)), \gamma(\omega), \tau_i] - \rho$$

for all $\beta' < \beta$.

Proposition 4.1. *For all $i \in [N]$ consider the strategy $s_i(\omega) = \alpha_i(\omega), \omega \in \Omega^N$. $\left\langle \Omega^N, \mu^{\beta,\boldsymbol{\tau},N}, (\mathcal{P}_i)_{i\in[N]}, (s_i)_{i\in[N]} \right\rangle$ is a (β, ρ)-correlated equilibrium.*

Proof. Consider a player $i \in [N]$ and an arbitrary alternative strategy $\hat{s}_i$. The deviation payoffs of player i are bounded by

$$\begin{aligned}&\sum_{\omega\in\Omega^N} \mu^{\beta,\boldsymbol{\tau},N}(\omega) \left\{ U_i[(\hat{s}_i(\omega), \boldsymbol{s}_{-i}(\omega)), \gamma(\omega), \tau_i] - U_i(\boldsymbol{s}(\omega), \gamma(\omega), \tau_i) \right\}\\&\quad= \sum_{\omega\in\Omega^{*,N}(\boldsymbol{\tau})} \mu^{\beta,\boldsymbol{\tau},N}(\omega) \left\{ V((\hat{s}_i(\omega), \boldsymbol{s}_{-i}(\omega), \gamma(\omega)), \boldsymbol{\tau}) - V(\omega, \boldsymbol{\tau}) \right\}\\&\quad+ \sum_{\omega\notin\Omega^{*,N}(\boldsymbol{\tau})} \mu^{\beta,\boldsymbol{\tau},N}(\omega) \left\{ V((\hat{s}_i(\omega), \boldsymbol{s}_{-i}(\omega), \gamma(\omega)), \boldsymbol{\tau}) - V(\omega, \boldsymbol{\tau}) \right\}\\&\quad\leq \mu^{\beta,\boldsymbol{\tau},N}(\Omega^N \setminus \Omega^{*,N}(\boldsymbol{\tau}))C\end{aligned}$$

where $C := \max_{\omega\notin\Omega^{*,N}(\boldsymbol{\tau})} \left\{ V((\hat{s}_i(\omega), \boldsymbol{s}_{-i}(\omega), \gamma(\omega)), \boldsymbol{\tau}) - V(\omega, \boldsymbol{\tau}) \right\}$. This upper bound follows from the fact that the first summand in the second line is non-positive by definition of the set $\Omega^{*,N}(\boldsymbol{\tau})$. If $C < 0$ we are done. If $C \geq 0$ we apply exponential tightness of the invariant distribution.

[24] This definition is inspired by the concept of ex-post Nash equilibrium of Kalai (Kalai, 2004, Definition 4, p. 1641).

Corollary 4.1 tells us that we can find an $\epsilon > 0$ such that $\mu^{\beta,\boldsymbol{\tau},N}(\Omega^N \setminus \Omega^{*,N}(\boldsymbol{\tau})) \leq \exp\left(-\frac{\epsilon}{\beta}(1+o(1))\right) =: \delta(\beta)$ for $\beta \downarrow 0$, and $\delta(\beta) \to 0$ in the respective limit. From this it follows that for every given $\rho > 0$ we can push down the established upper bound to be below ρ by choosing β sufficiently small. □

Thus, at least for sufficiently small noise we can be sure that the players are willing to follow the recommendation if we allow for small deviations from pure best responding. In a sense this means that if we sample states ω from the invariant distribution, then we can replicate the process by giving players the respective information partition.

4.2 *Large population limit*

In this final section we fix a positive noise level $\beta > 0$ and take the population size as a selection parameter. As in our study of the small noise limit we assume henceforth that the attachment mechanism is given by the log-linear function (3.3) and the volatility mechanism is semi-anonymous. Once we make this assumption we achieve that the labels of the players are unimportant for the weight of the invariant distribution $\mu^{\beta,\boldsymbol{\tau},N}$ at the various population states $\omega \in \Omega^N$. This will allow us to define population aggregates in a meaningful way and their distribution induced by the invariant measure. More specifically, in this section we will be concerned with the empirical distribution over actions, denoted as $\hat{\boldsymbol{\sigma}}^N = (\hat{\boldsymbol{\sigma}}_1^N, \ldots, \hat{\boldsymbol{\sigma}}_K^N)$, which we interpret as *Bayesian strategies*.[25] Each component of a Bayesian strategy $\hat{\boldsymbol{\sigma}}_k$ is a probability distribution on the action set A, whose coordinates are denoted by $\hat{\sigma}_k(a), a \in A$. Formally it is the empirical measure

$$\hat{\sigma}_k^N(a)(\omega, \boldsymbol{\tau}) := \frac{1}{NM_k^N(\boldsymbol{\tau})} \sum_{i=1}^{N} \mathbb{1}_a(\alpha_i(\omega))\mathbb{1}_{\theta_k}(\tau_i) \tag{4.4}$$

for all $a \in A$ and $1 \leq k \leq K$. Since we can view $\hat{\boldsymbol{\sigma}}$ as a map from the type space Θ to the set of mixed strategies $\Delta(A)$ we have formally indeed a Bayesian strategy in the classical game theoretic sense. Denote by $\Sigma := \Delta(A)^K$ the set of Bayesian strategies. Measurable sets generated by the mapping $\hat{\boldsymbol{\sigma}}^N$ are defined as

$$[\boldsymbol{\sigma}, \boldsymbol{m}] := \{(\omega, \boldsymbol{\tau}) \in \Omega^N \times \Theta^N | \hat{\boldsymbol{\sigma}}^N(\omega, \boldsymbol{\tau}) = \boldsymbol{\sigma} \ \& \ \boldsymbol{M}^N(\boldsymbol{\tau}) = \boldsymbol{m}\}. \tag{4.5}$$

[25] The interaction structure has been completely characterized in Section 3.1. Hence, the study of the marginal distribution on the set of action profiles is the only remaining part in order to achieve a complete characterization of the long-run behavior of the invariant distribution.

for $(\sigma, m) \in \Sigma \times \Delta(\Theta)$. Our goal in this section will be to derive a closed-form expression for the mass the invariant measure $\mu^{\beta,N} \in \mathsf{M}\left(\Omega^N \times \Theta^N\right)$, defined in eq. (3.11), puts on the set $[\sigma, m]$. By definition we can compute this mass as

$$\begin{aligned}\mu^{\beta,N}([\sigma, m]) &= \sum_{(\omega,\tau)\in[\sigma,m]} \mu^{\beta,N}(\omega, \tau) \\ &= \sum_{\tau\in\mathcal{T}^N(m)} \mathrm{P}_q(\tilde{\tau}^{(N)} = \tau) \sum_{\omega\in(\hat{\sigma}^{N,\tau})^{-1}(\sigma)} \mu^{\beta,\tau,N}(\omega),\end{aligned}$$

where $\hat{\sigma}^{N,\tau}(\cdot) = \hat{\sigma}^N(\cdot, \tau)$ and

$$\left(\hat{\sigma}^{N,\tau}\right)^{-1}(\sigma) = \{\omega \in \Omega^N | \hat{\sigma}^N(\omega, \tau) = \sigma\}.$$

From the definition of a Bayesian strategy it follows that all states in $\Omega^N_a \times \{\tau\}$, for $\tau \in \mathcal{T}^N(m)$, are in the set $[\sigma, m]$, and with it so are all sections resulting from permuting the players labels. Therefore we can define an equivalence relation $\sim_{[\sigma,m]}$ such that $(a, \tau) \sim_{[\sigma,m]} (a', \tau')$ if and only if $\tau, \tau' \in \mathcal{T}^N(m)$ and $\hat{\sigma}^N(\Omega^N_a, \tau) = \hat{\sigma}^N(\Omega^N_{a'}, \tau') = \sigma$. Putting it differently, with the equivalence relation $\sim_{[\sigma,m]}$ we declare a pair of action and type profiles as $[\sigma, m]$-equivalent if these profiles generate the same aggregate statistic (σ, m). From this it follows that $\left(\hat{\sigma}^{N,\tau}\right)^{-1}(\sigma)$ is a union of a-sections, In Appendix B we derive a closed form expression for the masses $\mu^{\beta,\tau,N}(\Omega^N_a)$ for all $a \in A^N$. Under semi-anonymity this measure only depends on the frequency of players of a certain type who play a certain action. By taking care of the combinatorial terms coming from all possible permutations of the players, we arrive at the following measure over (finite population) Bayesian strategies (see Theorem B.1 in Appendix B)

$$\psi^{\beta,N}(\sigma|m) = \mathcal{K}^{\beta,N}(m)^{-1} \prod_{k=1}^{K} \frac{(Nm_k)!}{\prod_{a\in A}(Nm_k\sigma_k(a))!} \exp\left(Nm_k f^{\beta,N}_k(\sigma, m)\right). \tag{4.6}$$

In this expression the factor $\mathcal{K}^{\beta,n}(m)$ is a function that normalizes the measure to be a probability measure. The support of the probability distribution $\psi^{\beta,N}(\cdot|m)$ is the subset

$$\mathrm{supp}(\psi^{\beta,N}(\cdot|m)) = \Sigma^N(m) := \{\sigma \in \Sigma | Nm_k\sigma_k(a) \in \mathbb{N}, \forall a \in A, 1 \leq k \leq K\},$$

a finite inner approximation of the polyhedron of Bayesian strategies Σ. The functions $f^{\beta,N}_k : \Sigma \times \Delta(\Theta) \to \mathbb{R}$ capture interaction payoffs players of type k earn from matches with players of type $l \geq k$, after we have aggregated over all possible networks and have taken care of their own idiosyncratic preferences.[26] While their finite N form is not very pretty[27], as N goes to infinity we can show that the sequence $\{f_k^{\beta,N}\}_{N\geq N^0}$ converges almost everywhere to the limit function

$$f_k^{\beta}(\sigma,\boldsymbol{m}) := \langle \boldsymbol{\sigma}_k, \boldsymbol{\theta}_k \rangle + \sum_{l\geq k} \frac{m_l}{1+\delta_{kl}} \left\langle \boldsymbol{\sigma}_k, \boldsymbol{\varphi}_{kl}^{\beta} \boldsymbol{\sigma}_l \right\rangle. \tag{4.7}$$

In this formulation we have identified the type $\theta_k : A \to \mathbb{R}$ with the n-dimensional vector $\boldsymbol{\theta}_k = (\theta_k(a))_{a\in A}$. The probability measure (4.6) is our candidate for proving a large deviations principle for the family $\{\psi^{\beta,N}(\cdot|\boldsymbol{M}^N)\}_{N\geq N^0}$. As in Theorem 4.1 this requires the identification of a rate function $r : \Sigma \times \Delta(\Theta) \to [0,\infty]$ such that for all converging sequences $\{(\boldsymbol{\sigma}^N, \boldsymbol{m}^N)\}_{N\geq N^0}$ with limit $(\boldsymbol{\sigma}, \boldsymbol{m}) \in \Sigma \times \Delta(\Theta)$ we have

$$-\lim_{N\to\infty} \frac{\beta}{N} \log \psi^{\beta,N}(\boldsymbol{\sigma}^N|\boldsymbol{m}^N) = r^{\beta}(\sigma, \boldsymbol{m}).$$

Unpacking this expression shows that, in the limit of large populations, the probability of observing a Bayesian strategy $\boldsymbol{\sigma} \in \Sigma$, given the limiting type distribution $\boldsymbol{m}$, is on the (logarithmic) order of $\exp\left(-\frac{N}{\beta} r(\sigma,\boldsymbol{m})\right)$. Hence, Bayesian strategies appearing with highest probability (on a logarithmic scale) are those for which $r(\sigma,\boldsymbol{m}) = 0$. We turn now to the identification of this rate function. While in Theorem 4.1 the rate function corresponds to a rescaled potential function of the game this will not be the right rate function for the present purposes. It turns out that the "logit potential functions" (Hofbauer and Sandholm, 2002; 2007)

$$\begin{aligned} (1 \leq k \leq K): \quad \tilde{f}_k^{\beta,N}(\boldsymbol{\sigma},\boldsymbol{m}) &:= f_k^{\beta,N}(\boldsymbol{\sigma},\boldsymbol{m}) + \beta h(\boldsymbol{\sigma}_k), \\ \tilde{f}^{\beta,N}(\boldsymbol{\sigma},\boldsymbol{m}) &:= \sum_{k=1}^{K} m_k \tilde{f}_k^{\beta,N}(\boldsymbol{\sigma},\boldsymbol{m}) \end{aligned} \tag{4.8}$$

play a key role, where $h(\boldsymbol{x}) = -\sum_i x_i \log x_i$ is the entropy of a distribution $\boldsymbol{x}$.[28] These functions depend on the distribution of types in the population. Since we consider a growing population of players, the distribution of types changes over time, and in fact, a version of the strong law of large numbers, proved in Appendix B (Lemma B.5), shows that

[26]The reason why only interactions with types $l \geq k$ are included in this function is to avoid double counting. Please see Appendix B for the details.

[27]For a definition of these functions we refer the reader to our results presented in Appendix B.

[28]Similar functions have been identified by Hofbauer and Sandholm (2002; 2007) to serve as Lyapunov functions for the mean-field dynamic generated by a stochastic evolutionary process where players revise their actions according to a general perturbed best-response rule. We do not look at mean-field equations here but instead are interested in large-deviations of the Bayesian strategies produced by the co-evolutionary process.

Lemma 4.3. $\boldsymbol{M}^N \overset{a.s.}{\to} \boldsymbol{q}$ *for* $N \to \infty$.

Hence, almost all realizations of nature's type assignment process will lead to a type distribution that is close to the prior $\boldsymbol{q}$, provided N is sufficiently large. We may thus focus on the set of type realizations on which $\boldsymbol{M}^N \to \boldsymbol{q}$, and ignore the limiting behavior of the measure $\psi^{\beta,N}(\cdot|\boldsymbol{M}^N)$ on the $\mathrm{P}_{\boldsymbol{q}}$-measure 0 set where this convergence may fail. Our main result of this section is the following theorem which establishes the desired large deviations principle for the family of measures $\left\{\psi^{\beta,N}(\cdot|\boldsymbol{m}^N)\right\}_{N\geq N^0}$, where $(\boldsymbol{m}^N)_{N\geq N^0}$ is a type distribution that converges to $\boldsymbol{q}$.

Theorem 4.2. *Let* $(\boldsymbol{m}^N)_{N\geq N^0}$ *be a sequence of type distributions which converges to the prior distribution* $\boldsymbol{q}$. *The family* $\{\psi^{\beta,N}(\cdot|\boldsymbol{m}^N)\}_{N\geq N_0}$, *generated under an admissible semi-anonymous volatility mechanism, satisfies a large deviations principle with rate function* $r^\beta(\boldsymbol{\sigma},\boldsymbol{q}) := \max_{\boldsymbol{\sigma}'\in\Sigma} \tilde{f}^\beta(\boldsymbol{\sigma}',\boldsymbol{q}) - \tilde{f}^\beta(\boldsymbol{\sigma},\boldsymbol{q})$, *for all* $\boldsymbol{\sigma} \in \Sigma$, *in the sense that*

$$\lim_{N\to\infty} \frac{\beta}{N} \log \psi^{\beta,N}(\boldsymbol{\sigma}^N|\boldsymbol{m}^N) = -r^\beta(\boldsymbol{\sigma},\boldsymbol{q}) \tag{4.9}$$

for every sequence $\left\{\boldsymbol{\sigma}^N\right\}_{N\geq N^0}$ *such that* $\boldsymbol{\sigma}^N \in \Sigma^N(\boldsymbol{m}^N), \forall N \geq N^0$ *and* $\boldsymbol{\sigma}^N \to \boldsymbol{\sigma}$.

Proof. See Appendix B. ☐

As the large deviations principle derived in Theorem 4.1 the message of Theorem 4.2 is that the family of measures $\{\psi^{\beta,N}(\cdot|\boldsymbol{m}^N)\}_{N\geq N^0}$ concentrates on a logarithmic scale at Bayesian strategies which solve the program

$$\max_{\boldsymbol{\sigma}'\in\Sigma} \tilde{f}^\beta(\boldsymbol{\sigma}',\boldsymbol{q}). \tag{4.10}$$

It is well known (see for instance Fudenberg and Levine, 1998, Hofbauer and Sandholm, 2002) that solutions of this program are *logit equilibria* (McKelvey and Palfrey, 1995), i.e. Bayesian strategies which are defined by the fixed-point condition

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

where

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

for all $a \in A$ and $1 \leq k \leq K$. To show this, simply note that solutions of the program (4.10) can be determined by solving a standard constrained optimization problem with Lagrangian

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

First order conditions give necessary and sufficient conditions for an optimum, which will be interior and unique for β sufficiently large and positive. Formally, the first-order conditions are

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

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

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

we obtain

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

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

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

from which the rest follows immediately by using the constrained $\sum_{a \in A} \sigma_k(a) = 1$.

To summarize, in the small noise limit we have seen in section 4.1 that the invariant distribution of the co-evolutionary process concentrates on the set of potential maximizers. In the large population limit a completely different prediction obeys. For finite N and in the small noise limit we have seen in section 3.1 that networks tend to cluster strongly. In the larger population limit with positive noise networks do not tend to cluster strongly (in general) but are well defined inhomogeneous random graphs, whose connectivity depends on the underlying volatility mechanism. Action profiles do not maximize a potential function but rather a "logit potential function".[29]

[29]See also Hofbauer and Sandholm (2007) for qualitatively the same result but in a purely game theoretic setting.

5 Conclusion

This paper presents an analytically solvable model on the co-evolution of networks of play in settings where players have diverse preferences. Under the assumptions that players have reward functions of the partnership type and use log-linear functions in the action choice and linking choice, we can give a closed-form solution of the (unique) invariant distribution of the process. This in turn allows us to investigate the robustness of certain equilibria in the case of small noise in the behavioral rules, and large populations. Many results presented in the paper hinge on the specific assumptions made in order to proceed with analytical methods. However, as mentioned in the main text, there are also some results which extend beyond the present framework. Among these is the creation of inhomogeneous random graphs. However, it remains an open problem how a co-evolutionary model behaves 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.

Appendix A The joint invariant distribution and the small noise limit

A.1 Proof of Theorem 3.1

As an ansatz for the invariant distribution we consider the detailed balance conditions

$$\mu^{\beta,\tau,N}(\omega|\Omega_{\boldsymbol{a}}^N)\eta_{\omega,\omega'}^{\beta,\tau,N} = \mu^{\beta,\tau,N}(\omega'|\Omega_{\boldsymbol{a}}^N)\eta_{\omega',\omega}^{\beta,\tau,N} \tag{A.1}$$

for all $\omega, \omega' \in \Omega_{\boldsymbol{a}}^N$. By force of normalization, given by the constant $Z^{\beta,\tau,N}(\boldsymbol{a})$, this system of equations has a unique solution

$$\mu^{\beta,\tau,N}(\omega|\Omega_{\boldsymbol{a}}^n) = Z^{\beta,\tau,N}(\boldsymbol{a})^{-1}\prod_{i=1}^N\prod_{j>i}\left(\frac{c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})}{\zeta_{ij}^{\beta,N}(\boldsymbol{\tau})}\right)^{\gamma_{ij}(\omega)}. \tag{A.2}$$

Define for all $i = 1, 2, \dots, N$ and $j > i$ the numbers

$$x_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau}) := \log\left(\frac{c_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})}{\zeta_{ij}^{\beta,N}(\boldsymbol{\tau})}\right),$$

and on $\Omega_{\boldsymbol{a}}^N$ the function

$$H_0(\omega|\Omega_{\boldsymbol{a}}^N) := \sum_{i=1}^N\sum_{j>i} x_{ij}^{\beta,N}(\boldsymbol{a},\boldsymbol{\tau})\gamma_{ij}(\omega).$$

Direct substitution into eq. (A.2) gives the alternative representation of the invariant distribution as

$$\mu^{\beta,\boldsymbol{\tau},N}(\omega|\Omega^N_{\boldsymbol{a}}) = \frac{\exp(H_0(\omega|\Omega^N_{\boldsymbol{a}}))}{\sum_{\omega\in\Omega^N_{\boldsymbol{a}}}\exp(H_0(\omega|\Omega^N_{\boldsymbol{a}}))}. \tag{A.3}$$

Define the probabilities $p^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})$ as in eq. (3.2) we can compute the numerator of eq. (A.3) as

$$\begin{aligned}
\sum_{\omega'\in\Omega^N_{\boldsymbol{a}}}\exp(H_0(\omega'|\Omega^N_{\boldsymbol{a}})) &= \sum_{\omega'\in\Omega^N_{\boldsymbol{a}}}\prod_{i=1}^N\prod_{j>i}\exp(x^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})\gamma_{ij}(\omega')) \\
&= \prod_{i=1}^N\prod_{j>i}(1+\exp(x^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})) \\
&= \prod_{i=1}^N\prod_{j>i}\left(1+\frac{c^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})}{\xi^{\beta,N}_{ij}(\boldsymbol{\tau})}\right) \\
&= \prod_{i=1}^N\prod_{j>i}\left(1-p^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})\right)^{-1}
\end{aligned}$$

Further, for all $\omega\in\Omega^N_{\boldsymbol{a}}$ we have

$$\exp(H_0(\omega|\Omega^N_{\boldsymbol{a}})) = \prod_{i=1}^N\prod_{j>i}\left(\frac{p^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})}{1-p^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})}\right)^{\gamma_{ij}(\omega)}.$$

Combining these last two observations, we obtain the desired product measure on $\Omega^N_{\boldsymbol{a}}$

$$\mu^{\beta,\boldsymbol{\tau},N}(\omega|\Omega^N_{\boldsymbol{a}}) = \prod_{i=1}^N\prod_{j>i}\left(p^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})\right)^{\gamma_{ij}(\omega)}\left(1-p^{\beta,N}_{ij}(\boldsymbol{a},\boldsymbol{\tau})\right)^{1-\gamma_{ij}(\omega)}.$$

□

A.2 Proof of Theorem 3.2

In order to proof Theorem 3.2 we need the following intermediate result.

Lemma A.1. *The Markov jump process $\{X^{\beta,\boldsymbol{\tau},N}_t\}_{t\geq 0}$ with infinitesimal generator $\{\eta^{\beta,\boldsymbol{\tau},N}_{\omega,\omega'}\}_{\omega,\omega'\in\Omega^N}$ whose rate functions are admissible has the unique invariant distribution*

$$\mu^{\beta,\boldsymbol{\tau},N}(\omega) = (Z^{\beta,\boldsymbol{\tau},N})^{-1}\prod_{i=1}^N\prod_{j>i}\left(\frac{2}{N}\frac{\exp(v(\alpha_i(\omega),\alpha_j(\omega))/\beta)}{\xi^{\beta,N}_{ij}(\boldsymbol{\tau})}\right)^{\gamma_{ij}(\omega)}\exp(\tau_i(\alpha_i(\omega))/\beta) \tag{A.4}$$

Proof. It is sufficient to show that the measure (A.4) satisfies the detailed balance condition (A.1) for all $\omega,\omega'\in\Omega^N$. We can split the proof of the claim in two parts:

Action revision and link creation. The case of link destruction is symmetric to the case of link creation, thus must not be treated separately. Hence, start with the process of link creation. Let $\omega = (\boldsymbol{a}, \boldsymbol{g}), \hat{\omega} = (\boldsymbol{a}, \boldsymbol{g} \oplus (i,j))$. We see that

$$\frac{\eta^{\beta,\boldsymbol{\tau},N}_{\omega,\hat{\omega}}}{\eta^{\beta,\boldsymbol{\tau},N}_{\hat{\omega},\omega}} = \frac{2\exp(v(a_i,a_j)/\beta)}{N\xi^{\beta,N}_{\tau_i,\tau_j}}. \tag{A.5}$$

Now, consider the ratio $\mu^{\beta,\boldsymbol{\tau},N}(\hat{\omega})/\mu^{\beta,\boldsymbol{\tau},N}(\omega)$. Due to the product structure, we see easily that all factors appearing in these two measures cancel out, and the only remaining factor is (A.5).

Next consider two states $\omega = (\boldsymbol{a}, \boldsymbol{g}), \hat{\omega} = ((a', \boldsymbol{a}_{-i}), \boldsymbol{g}), a' \in A$. We see that

$$\frac{\eta^{\beta,\boldsymbol{\tau},N}_{\omega,\hat{\omega}}}{\eta^{\beta,\boldsymbol{\tau},N}_{\hat{\omega},\omega}} = \frac{\ell^{i,\beta}_{a'}(\omega|\tau_i)}{\ell^{i,\beta}_{a}(\hat{\omega}|\tau_i)} = \exp\left[\frac{1}{\beta}\left(U_i((a', a_{-i}), \boldsymbol{g}) - U_i(\boldsymbol{a}, \boldsymbol{g})\right)\right].$$

To compute the odds ratio observe that

$$\begin{aligned}\frac{\mu^{\beta,\boldsymbol{\tau},N}(\omega)}{\mu^{\beta,\boldsymbol{\tau},N}(\hat{\omega})} &= \prod_{j=1}^{i}\prod_{k>j} \frac{\left[\frac{2\exp(v(\alpha_j(\omega),\alpha_k(\omega))/\beta)}{\xi^{\beta,N}_{\tau_j,\tau_k}}\right]^{\gamma_{jk}(\omega)} \exp(\tau_j(\alpha_j(\omega))/\beta)}{\left[\frac{2\exp(v(\alpha_j(\hat{\omega}),\alpha_k(\hat{\omega}))/\beta)}{\xi^{\beta,N}_{\tau_j,\tau_k}}\right]^{\gamma_{jk}(\hat{\omega})} \exp(\tau_j(\alpha_j(\hat{\omega}))/\beta)} \\ &\times \prod_{j=i+1}^{N}\prod_{k>j} \frac{\left[\frac{2\exp(v(\alpha_j(\omega),\alpha_k(\omega))/\beta)}{\xi^{\beta,N}_{\tau_j,\tau_k}}\right]^{\gamma_{jk}(\omega)} \exp(\tau_j(\alpha_j(\omega))/\beta)}{\left[\frac{2\exp(v(\alpha_j(\hat{\omega}),\alpha_k(\hat{\omega}))/\beta)}{\xi^{\beta,N}_{\tau_j,\tau_k}}\right]^{\gamma_{jk}(\hat{\omega})} \exp(\tau_j(\alpha_j(\hat{\omega})/\beta)}\end{aligned}$$

The second multiplicative term is independent of player i and so the ratio is equal to 1. Multiplying out the first term, and taking care of the symmetry of the reward function v, shows that

$$\begin{aligned}\frac{\mu^{\beta,N}(\omega)}{\mu^{\beta,N}(\hat{\omega})} &= \exp\left[\frac{1}{\beta}\left(\sum_{j=1}^{N}(v(\alpha_i(\omega),\alpha_j(\omega)) - v(\alpha_i(\hat{\omega}),\alpha_j(\hat{\omega}))) + \tau_i(\alpha_i(\omega) - \tau_i(\alpha_i(\hat{\omega}))\right)\right] \\ &= \exp\left[\frac{1}{\beta}\left(U_i(\boldsymbol{a}, \boldsymbol{g}, \tau_i) - U_i((a', \boldsymbol{a}_{-i}), \boldsymbol{g})\right)\right] = \frac{\eta^{\beta,\boldsymbol{\tau},N}_{\hat{\omega},\omega}}{\eta^{\beta,\boldsymbol{\tau},N}_{\omega,\hat{\omega}}}.\end{aligned}$$

$\square$

Proof of Theorem 3.2. As in Lemma A.1 we define, with a slight abuse of notation, for all $i, j \in [N]$ and $\omega \in \Omega^N$ the functions

$$x^{\beta,N}_{ij}(\omega, \boldsymbol{\tau}) := \log\left(\frac{2}{N}\frac{\exp(v(\alpha_i(\omega),\alpha_j(\omega))/\beta)}{\xi^{\beta,N}_{ij}(\boldsymbol{\tau})}\right) \tag{A.6}$$

$$= \frac{1}{\beta}v(\alpha_i(\omega),\alpha_j(\omega)) + \log\left(\frac{2}{N\xi^{\beta,N}_{ij}(\boldsymbol{\tau})}\right). \tag{A.7}$$

Then (A.4) can be written as

$$\begin{aligned}
\mu^{\beta,\boldsymbol{\tau},N}(\omega) &= (Z^{\beta,\boldsymbol{\tau},N})^{-1}\prod_{i=1}^{N}\prod_{j>i}\exp\left(x_{ij}^{\beta,N}(\omega,\boldsymbol{\tau})\gamma_{ij}(\omega)+\beta^{-1}\tau_i(\alpha_i(\omega))\right)\\
&= (Z^{\beta,\boldsymbol{\tau},N})^{-1}\exp\left[\beta^{-1}\sum_{i=1}^{N}\sum_{j>i}\left(x_{ij}^{\beta,N}(\omega,\boldsymbol{\tau})\gamma_{ij}(\omega)+\tau_i(\alpha_i(\omega))\right)\right]\\
&= (Z^{\beta,\boldsymbol{\tau},N})^{-1}\exp\left[\beta^{-1}\left(V(\omega,\boldsymbol{\tau})+\beta\log\mu_0^{\beta,\boldsymbol{\tau},N}(\omega)\right)\right]\\
&= (Z^{\beta,\boldsymbol{\tau},N})^{-1}\exp\left[\beta^{-1}H^{\beta,N}(\omega,\boldsymbol{\tau})\right].
\end{aligned}$$

□

Appendix B Aggregation and the large population limit

We are interested in the distribution of actions in the large population limit. Together with our results on the characterization of the inhomogeneous random graph in Section 3.1 this gives a complete description of the asymptotics of the co-evolutionary process. Our first step is to determine the marginal distribution on A^N for a given (but arbitrary) type profile $\boldsymbol{\tau}\in\Theta^N$. Under semi-anonymity this gives already (up to a combinatorial term) a distribution on the set of Bayesian strategies Σ with a support which depends on the type class $\mathcal{T}^N(\boldsymbol{m}),\boldsymbol{m}\in\mathcal{L}_N$ in which the type profile $\boldsymbol{\tau}$ happens to fall in. Before coming to the statement of some technical preparatory lemmas we need some additional notation. For aggregation purposes it is useful to have a partition of the set of players at hand that categorizes the players according to their action and their type. Formally, let us define the set

$$I_k^{\boldsymbol{\tau}}(a)(\omega):=\{i\in[N]|\alpha_i(\omega)=a\ \&\ \tau_i=\theta_k\}$$

for all $1\leq k\leq K$ and $a\in A$. Obviously, 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 $[N]$. Under semi-anonymous volatility mechanisms the random graph measure (3.1) treats all edges between players $\in I_k^{\boldsymbol{\tau}}(a)$ 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 $p_{kl}(a,b)$)

$$\mathcal{E}_{kl}^{N,\boldsymbol{\tau}}(a,b)(\omega):=\sum_{(i,j)\in[I_k^{\boldsymbol{\tau}}(a)\cup I_l^{\boldsymbol{\tau}}(b)]^{(2)}}\gamma_{ij}(\omega).$$

Given a type profile $\boldsymbol{\tau}$ and an $\boldsymbol{a}$-section $\Omega_{\boldsymbol{a}}^N$ 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 and $e_{kl}(a,b)$ a realization of the random variable $\mathcal{E}_{kl}^{N,\boldsymbol{\tau}}(a,b)(\cdot)$. We start with some important Lemmas.

Lemma B.1. *Consider a given type profile $\boldsymbol{\tau}\in\mathcal{T}^N(\boldsymbol{m})$ and a semi-anonymous volatility mechanism $\Xi^{\beta,\boldsymbol{\tau},N}$.*

(i) On the $\boldsymbol{a}$-section $\Omega^N_{\boldsymbol{a}}$, the fraction of a-players is fixed at

$$\sigma_k(a) = \hat{\sigma}_k(a)(\omega, \boldsymbol{\tau}) \quad \forall \omega \in \Omega^N_{\boldsymbol{a}}.$$

Call $\boldsymbol{\sigma} := (\sigma_k(a); 1 \leq k \leq K, a \in A) \in \Sigma^N(\boldsymbol{m})$.

(ii) We have

$$\mu^{\beta,\boldsymbol{\tau},N}(\Omega^N_{\boldsymbol{a}}) \propto \prod_{k=1}^{K}\prod_{a=1}^{n} \Phi^a_k(\boldsymbol{\sigma}, \beta, N)^{Nm_k\sigma_k(a)} \tag{B.1}$$

where, for all types $1 \leq k < l \leq K$, and actions $1 \leq a \leq n$, $\Phi^a_k(\cdot)$ is defined as

$$\Phi^a_k(\boldsymbol{\sigma}, \beta, N) := \prod_{l \geq k} \Phi^a_{kl}(\boldsymbol{\sigma}, \beta, N),$$

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

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

Proof. For notational simplicity let us drop the dependence of $\beta, \boldsymbol{\tau}$ and N from the involved functions and distributions. 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 Lemma is trivial. To prove item (ii) we proceed as follows;
For all $\omega \in \Omega^N$ define $\rho(\omega, \boldsymbol{\tau}) := \mu_0(\omega)\exp(V(\omega, \boldsymbol{\tau})/\beta)$. Using the functions $x_{ij}(\cdot, \cdot)$ of eq. (A.6), we can formulate this map as

$$\rho(\omega, \boldsymbol{\tau}) = \prod_{i=1}^{N}\prod_{j>i} \exp(x_{ij}(\omega, \boldsymbol{\tau})\gamma_{ij}(\omega)) \exp(\tau_i(\alpha_i(\omega))/\beta). \tag{B.2}$$

For all $\omega \in \Omega^N_{\boldsymbol{a}}$ the action classes are fixed by definition, and therefore $I^{\boldsymbol{\tau}}_k(a)(\omega) = I^{\boldsymbol{\tau}}_k(a)$ for all $1 \leq k \leq K, a \in A$ and $\omega \in \Omega^N_{\boldsymbol{a}}$. For all agents $i \in I^{\boldsymbol{\tau}}_k(a), j \in I^{\boldsymbol{\tau}}_l(b)$ we observe that

$$x_{ij}(\omega, \boldsymbol{\tau}) \equiv x_{kl}(a,b) := \frac{1}{\beta}v(a,b) + \log(\frac{2}{N\xi_{kl}}).$$

Thus, we can aggregate the product of (B.2) as

$$\begin{aligned} \rho(\omega, \boldsymbol{\tau}) &= \tilde{\rho}_{[\sigma,\boldsymbol{m}]}(\omega) := \prod_{k=1}^{K}\prod_{a=1}^{n} \exp\left(\frac{\theta_k(a)z_k(a)}{\beta}\right) \prod_{b \geq a} \exp\left[x_{kk}(a,b)\right]^{\mathcal{E}^{N,\boldsymbol{\tau}}_{kk}(a,b)(\omega)} \\ &\quad \times \prod_{k,l>k}\prod_{a,b \in A} \exp\left[x_{kl}(a,b)\right]^{\mathcal{E}^{N,\boldsymbol{\tau}}_{kl}(a,b)(\omega)} \end{aligned} \tag{B.3}$$

which is seen only to depend on the population state via the number of edges the network at ω has.[30] Now we aggregate this expression over all states $\omega \in \Omega^N_{\boldsymbol{a}}$.

[30]It does also not depend on the specific type *profile* but rather the realized type *distribution*. For notational unity we keep the index $\boldsymbol{\tau}$ however.

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 (an elementary version of Fubini's Theorem) by 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 integral, note that the only factor affected by the aggregation is $\exp(x_{11}(1,1))^{\mathcal{E}_{11}^{N,\boldsymbol{\tau}}(1,1)(\omega)}, \omega \in \Omega_{\boldsymbol{a}}^N$. Hence, if we collect terms unaffected by the aggregation under the placeholder B_1, we see that $\rho(\omega, \boldsymbol{\tau}) = B_1 \exp(x_{11}(1,1))^{\mathcal{E}_{11}^{N,\boldsymbol{\tau}}(1,1)(\omega)}$. 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 $\left\{\mathcal{E}_{11}^{N,\boldsymbol{\tau}}(1,1) = e_{11}(1,1)\right\}$. Adjusting for this we see that the output of the algorithm after the first round is

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

The next step performed by the algorithm inside Loop 1 will be to sum over all possible connections between players in the action cells $I_1^{\boldsymbol{\tau}}(1)$ and $I_1^{\boldsymbol{\tau}}(2)$. 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

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

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

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

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

This proposition shows how the invariant distribution weights on the $\boldsymbol{a}$-sections of the state space. From the proof it is clear that not the specific action profile is important for the invariant distribution weight, but only the Bayesian strategy it generates.

Lemma B.2. *If $\boldsymbol{\tau}, \boldsymbol{\tau}' \in \mathcal{T}^N(\boldsymbol{m})$ and $\hat{\sigma}^N(\Omega^N_{\boldsymbol{a}}, \boldsymbol{\tau}) = \hat{\sigma}^N(\Omega^N_{\boldsymbol{a}'}, \boldsymbol{\tau}')$ then*

$$\frac{\mu^{\beta,\boldsymbol{\tau},N}(\Omega^N_{\boldsymbol{a}})}{\mu^{\beta,\boldsymbol{\tau}',N}(\Omega^N_{\boldsymbol{a}'})} = 1.$$

Proof. Call $\boldsymbol{\sigma}$ the commonly generated Bayesian strategy on the subsets $\Omega^N_{\boldsymbol{a}}$ and $\Omega^N_{\boldsymbol{a}'}$ respectively. From eq. (B.3) we immediately see that for any $\omega \in \Omega^N_{\boldsymbol{a}}$ and $\omega' \in \Omega^N_{\boldsymbol{a}'}$ we have

$$\rho(\omega, \boldsymbol{\tau}) = \tilde{\rho}_{[\boldsymbol{\sigma},\boldsymbol{m}]}(\omega),\ \rho(\omega', \boldsymbol{\tau}') = \tilde{\rho}_{[\boldsymbol{\sigma},\boldsymbol{m}]}(\omega').$$

We are done if we can show that

$$\sum_{\omega \in \Omega^N_{\boldsymbol{a}}} \tilde{\rho}_{[\boldsymbol{\sigma},\boldsymbol{m}]}(\omega) = \sum_{\omega' \in \Omega^N_{\boldsymbol{a}'}} \tilde{\rho}_{[\boldsymbol{\sigma},\boldsymbol{m}]}(\omega'),$$

since this is the operation performed by the algorithm defined in the proof of Proposition B.1. But this follow from the fact that the networks that can be formed on $\Omega^N_{\boldsymbol{a}}$ are isomorphic to the networks that can be formed on $\Omega^N_{\boldsymbol{a}'}$.[31] □

Note that the fact that networks can be mapped from $\Omega^N_{\boldsymbol{a}}$ to $\Omega^N_{\boldsymbol{a}'}$ is a consequence of semi-anonymity and our conditioning on $\mathcal{T}^N(\boldsymbol{m})$, so that the frequency of players of a certain type is constant. Equipped with this insight, we finally obtain a distribution over Bayesian strategies.

Theorem B.1. *Conditional on the type class $\mathcal{T}^N(\boldsymbol{m})$ the probability distribution on the set of Bayesian strategies Σ is given by*

$$\psi^{\beta,N}(\boldsymbol{\sigma}|\boldsymbol{m}) = \mathcal{K}^{\beta,N}(\boldsymbol{m})^{-1} \prod_{k=1}^{K} \nu_k^{\beta,N}(\boldsymbol{\sigma}|\boldsymbol{m}), \tag{B.4}$$

where for all $1 \leq k \leq K$

$$\nu_k^{\beta,N}(\boldsymbol{\sigma}|\boldsymbol{m}) := \frac{(Nm_k)!}{\prod_{a\in A}(Nm_k\sigma_k(a))!} \prod_{a=1}^{n} \Phi_k^a(\boldsymbol{\sigma},\beta,N)^{Nm_k\sigma_k(a)}. \tag{B.5}$$

The support of this probability distribution is the subset

$$\operatorname{supp}(\psi^{\beta,N}(\cdot|\boldsymbol{m})) = \Sigma^N(\boldsymbol{m}) = \{\boldsymbol{\sigma} \in \Sigma | Nm_k\sigma_k(a) \in \mathbb{N}, \forall a \in A, 1 \leq k \leq K\},$$

a finite inner approximation of the polyhedron of Bayesian strategies Σ.

[31] We call two networks $\boldsymbol{g}, \boldsymbol{g}'$ isomorphic if there exists a permutation of players that preserves adjacency, i.e. if π is a permutation then $g_{ij} = 1$ iff $g_{\pi(i),\pi(j)} = 1$.

Proof. From Corollary B.2 we know that if the pair $(\boldsymbol{a}', \boldsymbol{\tau}')$ results from $(\boldsymbol{a}, \boldsymbol{\tau})$ by a permutation of the players labels, then $\mu^{\beta,\boldsymbol{\tau}',N}(\Omega^N_{\boldsymbol{a}'}) = \mu^{\beta,\boldsymbol{\tau},N}(\Omega^N_{\boldsymbol{a}})$. If $\boldsymbol{\tau} \in \mathcal{T}^N(\boldsymbol{m})$ and on $\Omega^N_{\boldsymbol{a}}$ the Bayesian strategy $\boldsymbol{\sigma}$ prevails, then there are Nm_k players of type θ_k and simple combinatorics tells us that there are $\frac{(Nm_k)!}{\prod_{a\in A}(Nm_k\sigma_k(a))!}$ ways to generate an action class with $Nm_k\sigma_k(a)$ players in subpopulation $1 \leq k \leq K$ who play action $a \in A$. This number of combinations is the same for all type profiles $\boldsymbol{\tau} \in \mathcal{T}^N(\boldsymbol{m})$. □

The closed-form for the invariant distribution over Bayesian strategies has another compact representation, which will turn out to be useful in studying the large population behavior of the measure. We introduce the functions

$$
\begin{aligned}
(1 \leq k \leq K): \quad f_k^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m}) &:= \sum_{a\in A} \sigma_k(a) \sum_{l\geq k} \log \Phi^a_{kl}(\boldsymbol{\sigma}, \beta, N), \\
f^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m}) &:= \sum_{k=1}^{K} m_k f_k^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m}).
\end{aligned} \tag{B.6}
$$

In terms of these maps we can write the distribution (B.4) as

$$
\psi^{\beta,N}(\boldsymbol{\sigma}|\boldsymbol{m}) = \mathcal{K}^{\beta,N}(\boldsymbol{m})^{-1} \prod_{k=1}^{K} \frac{(Nm_k)!}{\prod_{a\in A}(Nm_k\sigma_k(a))!} \exp\left(Nm_k f_k^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m})\right).
$$

It should be clear that the large population behavior of the law of Bayesian strategies depends on the convergence of the functions $\{f_k^{\beta,N}\}_{N\geq N^0}$. First we show that the sets $\Sigma^N(\boldsymbol{m}) \times \mathcal{L}_N$, for $\boldsymbol{m} \in \mathcal{L}_N$, approximate the continuous spaces $\Sigma \times \Delta(\Theta)$ arbitrarily well as $N \to \infty$. For our purposes it is enough to show this for the case where all types appear with a positive limiting frequency.[32] We hence have the following Lemma to prove.

Lemma B.3. *For every Bayesian strategy $\boldsymbol{\sigma} \in \Sigma$ and type distribution $\boldsymbol{m} \in \operatorname{int}\Delta(\Theta)$ there exists a sequence $\{(\boldsymbol{\sigma}^N, \boldsymbol{m}^N)\}_{N\geq N_0}$, with $\boldsymbol{\sigma}^N \in \Sigma^N(\boldsymbol{m}^N)$ and $\boldsymbol{m}^N \in \mathcal{L}_N$ for all $N \geq N_0$, such that $(\boldsymbol{\sigma}^N, \boldsymbol{m}^N) \to (\boldsymbol{\sigma}, \boldsymbol{m})$ as $N \to \infty$.*

Proof. The proof proceeds in two steps. First we show that we can find a sequence $\boldsymbol{m}^N \in \mathcal{L}_N$ that converges to $\boldsymbol{m}$ in total variation distance as $N \to \infty$.[33] Then we use this sequence to construct the sequence of Bayesian strategies.

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

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

[32]The reason for this is that we will later look at the convergence of the empirical measure $\boldsymbol{M}^N(\cdot)$ as $N \to \infty$. Since this measure is generated by a sequence of i.i.d. random variables with common law $\boldsymbol{q} \in \operatorname{int}\Delta(\Theta)$ we can prove a strong law of large numbers that states that $M^N(\cdot) \to \boldsymbol{q}$ almost surely (i.e. for all infinite sequences of types up to a set of $\mathrm{P}_{\boldsymbol{q}}$-measure zero).

[33]Any other norm will also do the job.

If $\boldsymbol{m}^N \in \mathcal{L}_N$ then each coordinate $m_k^N \in \{0, \frac{1}{N}, \ldots, \frac{N}{N}\}$. Thus, if $\boldsymbol{m} \in \Delta(\Theta)$ 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 $\boldsymbol{m}^N$ such that $||\boldsymbol{m}^N - \boldsymbol{m}||_{TV,\Theta} \leq \frac{K}{2N}$. Consequently, for a given $\delta > 0$ sufficiently small the set $\mathcal{N}^\delta(\boldsymbol{m}) := \{\boldsymbol{y} \in \Delta(\Theta) | \; ||\boldsymbol{y} - \boldsymbol{m}||_{TV,\Theta} < \delta\}$ is an open ball around $\boldsymbol{m}$ that contains $\boldsymbol{m}^N$ for all $N \geq N(\delta)$, where $N(\delta)$ is a suitably chosen integer. Hence $\boldsymbol{m}^N \to \boldsymbol{m}$ in total variation distance.

(ii) Given the sequence of empirical type distribution $(m^N)_{N \geq N_0}$ identified in item (i), let $\boldsymbol{\sigma}^N \in \Sigma^N(\boldsymbol{m}^N)$ 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 $\boldsymbol{m}^N \to \boldsymbol{m} \in \operatorname{int} \Delta(\Theta)$ it follows that for N sufficiently large there exists a $\epsilon > 0$ so that $m_k^N \geq \epsilon > 0$ for all $1 \leq k \leq K$. Hence, for $\delta > 0$ sufficiently small we may define a neighborhood $\mathcal{N}^\delta(\boldsymbol{\sigma}^N)$ as we did in point (i) and we observe that for $N \geq N(\delta)$, $\boldsymbol{\sigma}^N \in \mathcal{N}^\delta(\boldsymbol{\sigma})$. This completes the proof of the Lemma.

□

This gives us the security that we can always approximate a pair $(\boldsymbol{\sigma}, \boldsymbol{m}) \in \Sigma \times \Delta(\Theta)$ by a converging sequence of discrete distributions $(\boldsymbol{\sigma}^N, \boldsymbol{m}^N)$ which are measurable for the finite population process.

Lemma B.4. *For all $1 \leq k \leq K$ and along any sequence $\left\{(\boldsymbol{\sigma}^N, \boldsymbol{m}^N)\right\}_{N \geq N^0}$, $\boldsymbol{\sigma}^N \in \Sigma^N(\boldsymbol{m}^N), \boldsymbol{m}^N \in \mathcal{L}_N$, with limit $(\boldsymbol{\sigma}, \boldsymbol{m}) \in \Sigma \times \operatorname{int} \Delta(\Theta)$ we have*

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

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

$$f_k^\beta(\boldsymbol{\sigma}, \boldsymbol{m}) := \langle \boldsymbol{\sigma}_k, \boldsymbol{\theta}_k \rangle + \sum_{l \geq k} \frac{m_l}{1 + \delta_{kl}} \left\langle \boldsymbol{\sigma}_k, \boldsymbol{\varphi}_{kl}^\beta \boldsymbol{\sigma}_l \right\rangle. \tag{B.7}$$

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 $\varphi_{k,l}^{\beta,N}(a,b) = \frac{2\beta\exp(v(a,b)/\beta)}{\xi_{kl}^{\beta,N}}$. It follows from (LPB) that for all $1 \leq k,l \leq K$ and $a,b \in A$

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

which 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

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

and

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

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

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

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

Corollary B.1. *The sequence of functions $\{f^{\beta,N}\}_{N\geq N^0}$ converges almost everywhere to the limit function f^{β}.*

Proof. This follows from Lemma B.3 together with Lemma B.4. $\square$

In the preceding results all statements have been given for general sequences of type distributions. In the Bayesian interaction game the nature of the type assignment process is i.i.d with common law $\boldsymbol{q}$. Hence, asymptotically, the strong law of large numbers shows that type distributions resulting from the type assignment process must be close to $\boldsymbol{q}$. To be precise we have

Lemma B.5. *$\boldsymbol{M}^N \overset{a.s.}{\to} \boldsymbol{q}$ for $N \to \infty$.*

Proof. We take as metric on $\Delta(\Theta)$ again total variation distance[34], i.e. $||\boldsymbol{m} - \boldsymbol{q}||_{TV} := \frac{1}{2}\sum_{k=1}^K |m_k - q_k|$, $\forall \boldsymbol{m}, \boldsymbol{q} \in \Delta(\Theta)$. Recall that $\boldsymbol{q} \in \operatorname{int}\Delta(\Theta)$ is the common law of the types $\tilde{\tau}_i^{(N)}$. Around this point consider the countable family of open sets $\{B_{\boldsymbol{q},\epsilon}\}_{\epsilon \in \mathbb{Q}_+}$, where for all $\epsilon \geq 0$ $B_{\boldsymbol{q},\epsilon} := \{\boldsymbol{m} \in \Delta(\Theta) | \ ||\boldsymbol{m} - \boldsymbol{q}||_{TV} > \epsilon\}$, and The law induced by the empirical process $\{\boldsymbol{M}^N\}_{N \geq N_0}$ assigns mass to these sets as

$$\hat{P}_{\boldsymbol{q}}^N(B_{\boldsymbol{q},\epsilon}) = \mathrm{P}_{\boldsymbol{q}}(\{\boldsymbol{\tau}|\boldsymbol{M}^N(\boldsymbol{\tau}) \in B_{\boldsymbol{q},\epsilon}\})$$

From Sanov's Theorem it follows that

$$\lim_{N\to\infty} \frac{1}{N} \log \hat{P}_{\boldsymbol{q}}^N(B_{\boldsymbol{q},\epsilon}) = -\inf_{\boldsymbol{m} \in B_{\boldsymbol{q},\epsilon}} h(\boldsymbol{m}|\boldsymbol{q}),$$

where $h(\boldsymbol{m}|\boldsymbol{q}) := \sum_{k=1}^K m_k \log \frac{m_k}{q_k}$ is the relative entropy.[35] It is easily seen that $h(\cdot|\boldsymbol{q}) \geq 0$ with equality only at $\boldsymbol{q}$ (apply Jensen's inequality). Since $\boldsymbol{q} \notin B_{\boldsymbol{q},\epsilon}$ for all $\epsilon \geq 0$, it follows that for each such ϵ we can find a constant $c_\epsilon \in (0, \infty)$ such that

$$\hat{P}_{\boldsymbol{q}}^N(B_{\boldsymbol{q},\epsilon}) \leq e^{-Nc_\epsilon}. \tag{B.8}$$

We now translate the set $B_{\boldsymbol{q},\epsilon}$ into an event of nature's type assignment experiment. Consider the set

$$A_N(\epsilon) := \{\boldsymbol{\tau}|\boldsymbol{M}^N(\boldsymbol{\tau}) \in B_{\boldsymbol{q},\epsilon}\} = \{\boldsymbol{\tau}| \ ||\boldsymbol{M}^N(\boldsymbol{\tau}) - \boldsymbol{q}||_{TV} > \epsilon\}.$$

This event has $\mathrm{P}_{\boldsymbol{q}}$-probability $\hat{P}_{\boldsymbol{q}}^N(B_{\boldsymbol{q},\epsilon})$ by construction. Combined with equation (B.8) we see that

$$\sum_{N \geq N^0} \mathrm{P}_{\boldsymbol{q}}(A_N(\epsilon)) = \sum_{N \geq N^0} \hat{P}_{\boldsymbol{q}}^N(B_{\boldsymbol{q},\epsilon}) \leq \sum_{N \geq N^0} e^{-Nc_\epsilon} < \infty$$

Hence, by the first Borel-Cantelli Lemma[36] we conclude $\mathrm{P}_{\boldsymbol{q}}(\limsup_{N\to\infty} A_N(\epsilon)) = 0$ for all $\epsilon \in \mathbb{Q}_+$, which gives us almost sure convergence of the empirical process $\{\boldsymbol{M}^N\}_{N \geq N_0}$. □

In the main text of the paper we have argued that the "logit potential functions" play a vital role in deriving the large deviations principle for the family of measures $\{\psi^{\beta,N}(\cdot|\boldsymbol{M}^N)\}_{N \geq N^0}$. Recall that these function are defined as

$$\begin{aligned} (1 \leq k \leq K): \quad & \tilde{f}_k^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m}) := f_k^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m}) + \beta h(\sigma_k), \\ & \tilde{f}^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m}) := \sum_{k=1}^{k} m_k \tilde{f}_k^{\beta,N}(\boldsymbol{\sigma}, \boldsymbol{m}) \end{aligned} \tag{B.9}$$

We have to proof the following

[34] As in Lemma B.4 the particular choice of a metric is not important

[35] We set $0 \log 0 = 0$ and $0 \log \frac{0}{0} = 0$. Recall also that $\boldsymbol{q} \in \operatorname{int}\Delta(\Theta)$, and hence the relative entropy is well-defined on all of $\Delta(\Theta)$. Sanov's Theorem is related to Cramér's Theorem, but considers large deviations of empirical processes. See Dembo and Zeitouni (1998), chapter 2.

[36] See for instance Dudley (2002, Theorem 8.3.4).

Theorem B.2. *Let $(\boldsymbol{m}^N)_{N\geq N^0}$ be a sequence of type distributions which converges to the prior distribution $\boldsymbol{q}$. The family $\{\psi^{\beta,N}(\cdot|\boldsymbol{m}^N)\}_{N\geq N_0}$, generated under an admissible semi-anonymous volatility mechanism, satisfies a large deviations principle with rate function $r^\beta(\boldsymbol{\sigma},\boldsymbol{q}) := \max_{\sigma'\in\Sigma}\tilde{f}^\beta(\boldsymbol{\sigma}',\boldsymbol{q}) - \tilde{f}^\beta(\boldsymbol{\sigma},\boldsymbol{q})$, for all $\boldsymbol{\sigma}\in\Sigma$, in the sense that*

$$\lim_{N\to\infty}\frac{\beta}{N}\log\psi^{\beta,N}(\boldsymbol{\sigma}^N|\boldsymbol{m}^N) = -r^\beta(\boldsymbol{\sigma},\boldsymbol{q}) \tag{B.10}$$

for every sequence $\{\boldsymbol{\sigma}^N\}_{N\geq N^0}$ such that $\boldsymbol{\sigma}^N\in\Sigma^N(\boldsymbol{m}^N),\forall N\geq N^0$ and $\boldsymbol{\sigma}^N\to\boldsymbol{\sigma}$.

Proof. By Lemma B.5 we can pick a sequence $(\boldsymbol{m}^N)_{N\geq N^0}$ that converges to $\boldsymbol{q}$ with probability 1. Let us denote by $\boldsymbol{e}_1 = (\boldsymbol{e}_1(1),\dots,\boldsymbol{e}_K(1))$ the Bayesian strategy where all players of all types play action 1, i.e. for each $1\leq k\leq K$, $\boldsymbol{e}_k(1)$ is the unit vector of $\mathbb{R}^n$ with 1 in its first component and zero in its $n-1$ remaining components. Of course $\boldsymbol{e}_1\in\Sigma^N(\boldsymbol{m}^N)$ for all N. Then for all $\boldsymbol{\sigma}\in\Sigma^N(\boldsymbol{m}^N)$

$$\frac{\psi^{\beta,N}(\boldsymbol{\sigma}|\boldsymbol{m}^N)}{\psi^{\beta,N}(\boldsymbol{e}_1|\boldsymbol{m}^N)} = \prod_{k=1}^{K}\frac{(Nm_k)!}{\prod_{a\in A}(Nm_k^N\sigma_k^N(a))!}\exp\left[Nm_k^N\left(f_k^{\beta,N}(\boldsymbol{\sigma}^N,\boldsymbol{m}^N) - f_k^{\beta,N}(\boldsymbol{e}_1,\boldsymbol{m}^N)\right)\right]. \tag{B.11}$$

Taking logarithms and multiplying by $\frac{\beta}{N}$ gives us

$$\begin{aligned}\frac{\beta}{N}\log\frac{\psi^{\beta,N}(\boldsymbol{\sigma}^N|\boldsymbol{m}^N)}{\psi^{\beta,N}(\boldsymbol{e}_1|\boldsymbol{m}^N)} =& \frac{\beta}{N}\sum_{k=1}^{K}\log\left(\frac{(Nm_k^N)!}{\prod_{a\in A}(Nm_k^N\sigma_k^N(a))!}\right)\\ &+\sum_{k=1}^{K}m_k^N\left(f_k^{\beta,N}(\boldsymbol{\sigma}^N,\boldsymbol{m}^N) - f_k^{\beta,N}(\boldsymbol{e}_1,\boldsymbol{m}^N)\right)\end{aligned}$$

To handle the limit of the combinatorial terms the following simple version of Stirling's formula $n!\cong\sqrt{2\pi n}\left(\frac{n}{e}\right)^n$ will be sufficient for us. Some tedious, but straightforward algebra, gives us

$$\frac{1}{N}\log\left(\frac{(Nm_k^N)!}{\prod_{a\in A}(Nm_k^N\sigma_k^N(a|\boldsymbol{m}^N))!}\right) = m_k^N\left[h(\sigma_k^N) + O(1/N)\right].$$

From Lemma B.4 we know that $f^{\beta,N}(\boldsymbol{\sigma}^N,\boldsymbol{m}^N)\to f^\beta(\boldsymbol{\sigma},\boldsymbol{q})$ along the converging sequence $\{(\boldsymbol{\sigma}^N,\boldsymbol{m}^N)\}_{N\geq N^0}$. It follows that

$$\lim_{N\to\infty}\frac{\beta}{N}\frac{\psi^{\beta,N}(\boldsymbol{\sigma}^N|\boldsymbol{m}^N)}{\psi^{\beta,N}(\boldsymbol{e}_1|\boldsymbol{m}^N)} = \tilde{f}^\beta(\boldsymbol{\sigma},\boldsymbol{q}) - \tilde{f}^\beta(\boldsymbol{e}_1,\boldsymbol{q}) \tag{B.12}$$

where $\tilde{f}^\beta(\cdot,\cdot)$ is the logit potential function defined in (B.9). Next, let $\boldsymbol{\sigma}_*^N$ be a maximizer of the function

$$\tilde{f}^{\beta,N}(\sigma^N,\boldsymbol{m}^N) := \sum_{k=1}^{K}m_k^N\left[f_k^{\beta,N}(\sigma^N,\boldsymbol{m}^N)+\beta h(\boldsymbol{\sigma}_k^N)\right],\ \boldsymbol{\sigma}^N\in\Sigma^N(\boldsymbol{m}^N)$$

for all $N\geq N_0$. Then, by uniform convergence, we have $\tilde{f}^{\beta,N}(\boldsymbol{\sigma}_*^N,\boldsymbol{m}^N)\to\tilde{f}^\beta(\boldsymbol{\sigma}_*,\boldsymbol{q})$ as $N\to\infty$, and the limit point is a maximizer of $\tilde{f}^\beta(\cdot,\boldsymbol{q})$. It follows from Sandholm (2010, Theorem 12.2.2) that

$$\lim_{N\to\infty}\frac{\beta}{N}\log\psi^{\beta,N}(\boldsymbol{\sigma}_*^N|\boldsymbol{m}^N) = 0. \tag{B.13}$$

This completes the proof, since for all $\sigma^N \to \sigma \in \Sigma$ we know from (B.11) that

$$\begin{aligned}\lim_{N\to\infty}\frac{\beta}{N}\log\psi^{\beta,N}(\sigma^N|\boldsymbol{m}^N) &= \lim_{N\to\infty}\Big[\frac{\beta}{N}\log\frac{\psi^{\beta,N}(\sigma^N|\boldsymbol{m}^N)}{\psi^{\beta,N}(\boldsymbol{e}_1|\boldsymbol{m}^N)} \\ &\quad -\frac{\beta}{N}\log\frac{\psi^{\beta,N}(\sigma_*^N|\boldsymbol{m}^N)}{\psi^{\beta,N}(\boldsymbol{e}_1|\boldsymbol{m}^N)}+\frac{\beta}{N}\log\psi^{\beta,N}(\sigma_*^N|\boldsymbol{m}^N)\Big] \\ &= \tilde{f}^\beta(\sigma,\boldsymbol{q})-\tilde{f}^\beta(\sigma_*,\boldsymbol{q}) \text{ |by (B.13)} \\ &= -r^\beta(\sigma,\boldsymbol{q}).\end{aligned}$$

□

References

Aumann, R. J., 1987. Correlated equilibrium as an expression of bayesian rationality. Econometrica 55, No.1, 1–18.

Ballester, C., Calvó-Armengol, A., Zenou, Y., 2006. Who's who in networks. wanted: The key player. Econometrica 74 (5), 1403–1417.

Binmore, K., Samuelson, L., 1997. Muddling through: Noisy equilibrium selection. Jounal of Economic Theory 74, 235–265.

Binmore, K., Samuelson, L., Vaughn, R., 1995. Musical chairs: Modelling noisy evolution. Games and Economic Behavior 11, 1–35.

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.

Bollobás, B., Janson, S., Riordan, O., 2007. The phase transition in inhomogeneous random graphs. Random Structures and Algorithms vol. 31, Issue 1, 3–122.

Currarini, S., Jackson, M. O., Pin, P., 2009. An economic model of friendship: Homophily, minorities and segregation. Econometrica 77, No. 4, 1003–1045.

Dembo, A., Zeitouni, O., 1998. Large Deviation Techniques and Applications. Springer Verlag (2nd edition).

Dudley, R. M., 2002. Real Analysis and Probability. Cambridge University Press.

Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2006a. Diffusion and growth in an evolving network. International Journal of Game Theory 34, 383–397.

Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2006b. Phenomenological models of socio-economic network dynamics. Physical Review E 74, 036106.

Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2008. Emergence and resilience of social networks: a general theoretical framework. Annales d' économie et de statistique 86, 1–13.

Ellison, G., 2000. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. Review of Economic Studies 67, 17–45.

Ely, J., Sandholm, W. H., 2005. Evolution in bayesian games i: Theory. Games and Economic Behavior 53, 83–109.

Frank, O., Strauss, D., 1986. Markov graphs. Jounal of the American Statistical Association Vol.81, no. 395, 832–842.

Fudenberg, D., Levine, D. K., 1998. The theory of learning in games. The MIT Press, Cambridge Massachusetts.

Golub, B., Jackson, M. O., 2010. How homophily affects diffusion and learning in networks. Mimeo, http://arxiv.org/abs/0811.4013 [physics.soc-ph].

Goyal, S., 2007. Connections: An introductions to the economics of networks. Princeton University Press.

Hofbauer, J., Sandholm, W. H., 2002. On the global convergence of stochastic fictious play. Econometrica 70 (6), 2265 – 2294.

Hofbauer, J., Sandholm, W. H., 2007. Evolution in games with randomly disturbed payoffs. Journal of Economic Theory 132, 47–69.

Hofbauer, J., Sigmund, K., 1988. The Theory of Evolution and Dynamical Systems. Cambridge University Press.

Horst, U., Scheinkman, J. A., 2006. Equilibria in systems of social interactions. Jounal of Economic Theory 130, Issue 1, 44–77.

Jackson, M. O., 2008a. Average distance, diameter, and clustering in social networks with homophily. In: Internet and Network Economics 4th International Workshop, WINE 2008. Springer Berlin / Heidelberg.

Jackson, M. O., 2008b. 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.

Jackson, M. O., Yariv, L., 2008. Diffusion, strategic interaction, and social structure. In: Benhabib, J., Bisin, A., Jackson, M. O. (Eds.), Handbook of Social Economics. Elsevier.

Kajii, A., Morris, S., 1997. The robustness of equilibria to incomplete information. Econometrica 65, no.6, 1283–1309.

Kalai, E., 2004. Large robust games. Econometrica 72 (6), 1631–1665.

Kirman, A., Oddou, S., Weber, S., 1986. Stochastic communication and coalition formation. Econometrica 54 (1), 129–138.

König, M. D., Staudigl, M., 2010. Centrality based network formation of boundedly rational agents with limited information. Mimeo.

Mailath, G. J., Samuelson, L., Shaked, A., 1997. Correlated equilibrium and local interactions. Economic Theory 9, 551–556.

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., 1997. Interaction games: A unified analysis of incomplete information, local interaction and random matching. CARESS Working Paper 97-02.

Morris, S., Ui, T., 2005. Generalized potentials and robust sets of equilibria. Journal of Economic Theory 124, 45–78.

Robins, G., Pattison, P., Kalish, Y., Lusher, D., 2007. An introduction to exponential random graph (p*) models for social networks. Social Networks 29, 173–191.

Sandholm, W. H., 2010. Evolutionary dynamics and Population games. Forthcoming: The MIT Press.

Söderberg, B., 2002. General formalism for inhomogeneous random graphs. Physical Review E 66, 066121– 1– 6.

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

Staudigl, M., 2010b. Potential games in volatile environments. Forthcoming: Games and Economic Behavior.

Stroock, D., 2005. An Introduction to Markov Processes. Graduate Texts in Mathematics 230, Springer Berlin Heidelberg New York.

Ui, T., 2001. Robust equilibria of potential games. Econometrica Vol. 69, No. 5, 1373–1380.

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.

European University Institute
European University Press