WORKING PAPERS

Mathias Staudigl

On a General class of stochastic co-evolutionary dynamics

February 2010

Working Paper No: 1001

DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

On a General class of stochastic co-evolutionary dynamics

Mathias Staudigl*

February 4, 2010

Abstract

This paper presents a unified framework to study the co-evolution of networks and play, using the language of evolutionary game theory. We show by examples that the set-up is rich enough to encompass many recent models discussed by the literature. We completely characterize the invariant distribution of such processes and show how to calculate stochastically stable states by means of a tree-characterization algorithm. Moreover, specializing the process a bit further allows us to completely characterize the generated random graph ensemble. This new result demonstrates a new and rather general relation between random graph theory and evolutionary models with endogenous interaction structures.

Keywords: Evolutionary game theory, Network co-evolution, Random graphs
JEL Classification Numbers: C02, C73, C45, D85

1 Introduction

Recently there has been an attempt to apply stochastic evolutionary game dynamics to models on the co-evolution of networks and play. Broadly speaking, one may divide these models in two classes. There is one

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

branch of literature which extends the mistakes model of Kandori et al. (1993) and Young (1993) to a random process of action adjustment and link creation/destruction. Jackson and Watts (2002), Goyal and Vega-Redondo (2005), Hojman and Szeidl (2006) are models in this direction, and we call them, due to their ancestry, "classical" models. Another type of models assume that the network is under a recurrent attack of unguided drift, which is interpreted as environmental volatility. Marsili et al. (2004), Ehrhardt et al. (2006; 2008a) are models in this direction, which we will call "volatility" models. The aim of this paper is to present an unified framework, that is rich enough to incorporate classical, as well as volatility models. We do so by presenting a rather general class of co-evolutionary models, called $\mathcal{M}^\beta$. In essence $\mathcal{M}^\beta$ is a family of perturbed Markov chains taking values on some finite state space Ω, which consists of all pairs of action profiles (α) and networks (g). We give an "axiomatic" definition of processes $\mathcal{M}^\beta$ which models the co-evolution of networks and play in an integrated way. At a heuristic level, the algorithm works as follows:

Suppose the system starts from some point $\omega = (\alpha, g)$. Departing from this state, the system may evolve via three possible routes. With some probability a randomly chosen individual gets the opportunity to change his action. This causes a change in the action profile α. With complementary probability the network changes, resulting in the creation of a new edge, or the destruction of an existing edge. The characterizing feature of the process is that the *behavioral rules*, describing how agents change their action, or how they create or delete links depend, in general, on the benefits of the bilateral interaction, which, in turn, is modeled by a game in normal form. This produces an interesting coupling between the evolution of the action profile α and the evolution of the network g. After one of these events, the process arrives at a new state, and the algorithm repeats these steps infinitely often.

The objective of this paper is to investigate the asymptotic properties of this stochastic algorithm. We assume that the rules defining the transition probabilities of $\mathcal{M}^{\beta}$ are governed by a *noise parameter* $\beta \in \mathbb{R}_{+}$, as is by now standard in stochastic evolutionary models.[1] For $\beta > 0$ the process will be ergodic, and the long-run predictions are given by its unique invariant distribution $\mu^{\beta} \in \Delta(\Omega)$. In principle, the invariant distribution contains all information one needs to deduce more specific information about the long-run behavior of the system, such as the *marginal* probability distribution over action configurations (the object studied in "classical" evolutionary game theory with fixed interaction structure) and the *conditional* probability distribution over networks.[2] Particularly interesting is the behavior of the invariant distribution as noise vanishes. This leads to the study of stochastically stable states, which is one of the most prominent selection criteria of evolutionary game theory. The traditional way to perform stochastic stability analysis is by viewing the Markov process as a weighted and directed graph and looking for paths with least resistance. Kandori et al. (1993) and Young (1993) pioneered this approach, by adapting tools developed by Freidlin and Wentzell (1998). The first contribution of this paper is the presentation of a tree-characterization algorithm to compute stochastically stable states in general co-evolutionary models. Thereby we obtain a selection criterion of recurrent classes of states consisting of profiles of actions and architectures of interaction, extending traditional models of evolutionary game theory where only the action profiles are considered as state variable. The "classical" models of Jackson and Watts (2002) and Goyal

[1] The term "noise" has been used by Blume (2003) to emphasize the random utility context of probabilistic behavioral rules. Other used terms have been "mistakes" or "mutations".

[2] Due to the coupling of the behavior dimension with the network dimension it would make no sense to study a marginal distribution over networks. Only a conditional distribution, i.e. the probability distribution over networks for a *fixed* action profile, makes sense in these models.

and Vega-Redondo (2005) are also concerned with this task. Our general model provides a systematic tool kit to find stochastically stable states in a transparent way. We show by means of two examples, a "volatility" model and a "classical" model based on Jackson and Watts (2002), that such a stochastic stability analysis is still tractable in co-evolutionary models.

The second, and truly original, contribution of this paper is the characterization of the generated *random graph ensemble*, conditional on a fixed profile of actions. For this characterization we impose 3 additional "axioms". We show that any stochastic process, satisfying the stated assumptions, will converge in the long run to the probability ensemble of so-called *inhomogeneous random graphs* (Söderberg, 2002, Bollobás et al., 2007). Inhomogeneous random graphs are a straightforward extension of the classical Erdös-Rényi model (Erdös and Rényi, 1960), by allowing edge success probabilities to be vertex specific. These models are very popular in the literature on random graphs, and to the best of our knowledge, this interesting connection between evolutionary game dynamics and random graph theory is novel. A co-evolutionary model with noise provides therefore a new and independent derivation of inhomogeneous random graphs.

The class of Markov chains we study is known in the literature on stochastic optimization as a "generalized Metropolis algorithm", and is rigorously surveyed by Catoni (1999; 2001). Beggs (2005) was among the first to recognize the close relationship between this class of random processes, and the stochastic dynamics used in evolutionary game theory. We also exploit this analogy and show that it provides a flexible language to study many models on the co-evolution of networks and play. To underline this, we devote a whole section to show that the models of Staudigl (2009b;a) fit perfectly into our framework. A minor modification of the process also allows us to study the model of Jackson and Watts (2002).

Related to our work is also the recent paper by Alós-Ferrer and Netzer (2007). However, these authors fix the behavioral rules of the agents at the outset, by assuming that strategy revisions are governed by a log-linear process, introduced by Blume (1993) into game theory. Moreover, their paper assumes an exogenously fixed interaction structure.

The rest of the paper is organized as follows. Section 2 introduces our theoretical framework. In Section 2.2 we derive a general form of the invariant distribution, and an algorithm to detect stochastically stable states. Section 3 presents a "classical model" and a "volatility" model. The characterization of the generated random graph ensemble is presented in Section 4. Section 5 concludes. Appendix B collect some well-known facts on stochastic stability analysis in a self-contained way.

2 A class of Markov processes

We consider a finite population of individuals $\mathcal{I} = \{1, 2, \dots, N\}$. Members of this set are also called agents or players. The set of all unordered pairs of individuals will be denoted by $\mathcal{I}^{(2)}$. The set of ordered pairs of a finite set Ω is denoted as $\Omega \times \Omega = \Omega^2$. In this paper we identify networks with simple and undirected graphs on the vertex set $\mathcal{I}$. Let us call $\mathcal{G}[\mathcal{I}]$ the set of all such graphs, members of which are pairs $G = (\mathcal{I}, \mathcal{E})$, where $\mathcal{E} = \mathcal{E}(G) \subseteq \mathcal{I}^{(2)}$ is the set of edges (links). Another convenient representation of a network is via a tuple $g = (g_{ij})_{1 \leq i < j \leq N} \in \{0, 1\}^{\mathcal{I}^{(2)}} \equiv \mathcal{G}[\mathcal{I}]$. If $g_{ij} = 1$ we say that individual i is connected to individual j, or j is a neighbor of i (and vice versa). Another terminology for connectedness will be that the edge (i, j) is active. If $g_{ij} = 0$ then i and j are not connected, or edge (i, j) is neutral. The neighbors of player i are contained in the set $\mathcal{N}^i(g) := \{j \in \mathcal{I} | g_{ij} = 1\}$. Call $\bar{\mathcal{N}}^i(g) := \mathcal{N}^i(g) \cup \{i\}$. The number of neighbors of player i defines his degree $\kappa^i(g) := |\mathcal{N}^i(g)|$. Given a network g and a subset of players $\mathcal{V} \subseteq \mathcal{I}$ denote the restriction of g on $\mathcal{V}$ as $g[\mathcal{V}]$, which is an element of $\mathcal{G}[\mathcal{V}]$. The complete network on the subset $\mathcal{V}$ is denoted by $g^c[\mathcal{V}]$. Hence, for every $g \in \mathcal{G}[\mathcal{I}]$ and a partition of $\mathcal{I}$ into sets $\mathcal{V}_1, \mathcal{V}_2$, we can write $g = g[\mathcal{V}_1] \oplus g[\mathcal{V}_2]$, where $\oplus$ is interpreted as the concatenation of two lists of binary valued functions (after possibly relabeling the players). In this notation $g' = g \oplus g^c[\{(i,j)\}] \equiv g \oplus (i,j)$ is the network obtained by adding the edge (i,j) to g. Analogously, $g' = g \ominus (i,j)$ is the network obtained from g by deleting edge (i,j). Denote by $e = \sum_{i,j>i} g_{ij}$ the number of edges in the network.

Each individual possesses a utility function u^i, describing her preferences over some finite set of common actions $\mathcal{A} = \{a_1, \ldots, a_q\}$.[3] This defines a *base game* $\Gamma = (\mathcal{I}, \mathcal{A}, (u^i)_{i \in \mathcal{I}})$.

The utility player i gets from choosing one of these actions depends on the behavior of his neighboring players. Let $\alpha = (\alpha^i)_{i \in \mathcal{I}} \in \mathcal{A}^{\mathcal{I}}$ denote an action profile of the population. A *population state* is a pair $\omega = (\alpha, g) \in \mathcal{A}^{\mathcal{I}} \times \mathcal{G}[\mathcal{I}] \equiv \Omega$. Given an action profile α, let $\alpha_i^a = (a, \alpha_{-i}) = (\alpha^1, \ldots, \alpha^{i-1}, a, \alpha^{i+1}, \ldots, \alpha^N)$. Utility of player i at state ω is defined as

$$\pi^i(\alpha, g) \equiv \pi^i(\omega) = \sum_{j \in \mathcal{N}^i(g)} u^i(\alpha^i, \alpha^j). \tag{2.1}$$

2.1 *Co-evolution with noise*

In the spirit of Young (1993) and Ellison (2000), we call a *co-evolutionary model with noise* a family of perturbed time-homogeneous Markov chains

$$\mathcal{M}^\beta = \left(\Omega, \mathcal{F}, \mathbb{P}, (X_n^\beta)_{n \in \mathbb{N}_0}\right), \ \beta \in \mathbb{R}_+,$$

where

- Ω is some finite or countable infinite set describing the state space of the system,
- $X^\beta = (X_n^\beta)_{n \in \mathbb{N}_0}$ is a family of Ω-valued random variables, indexed by a discrete time parameter n and a noise parameter β,

[3] In principle every individual could have his own action set. This would require more notation, and does not contribute anything to this paper.

- $\mathcal{F}$ is a σ-algebra on Ω (e.g. $\mathcal{F} = 2^{\Omega}$ the set of all subsets of Ω if the state space is finite),

- $\mathbb{P} : \mathcal{F} \to [0,1]$ a probability measure.

A realization $\{X_n^\beta = \omega\}$ defines an action profile α and a network g. The Markov property states that for any history $A_{n-1} = \{X_0^\beta, \ldots, X_{n-1}^\beta\}$ on which $\{X_{n-1}^\beta = \omega\}$ holds, the probability that the process visits state ω' in the next period depends only on ω, i.e.

$$\mathbb{P}(X_n^\beta = \omega' | A_{n-1}) = \mathbb{P}(X_n^\beta = \omega' | X_{n-1}^\beta = \omega) \equiv K^\beta(\omega, \omega'), \tag{2.2}$$

where $K^\beta : \Omega^2 \to [0,1]$ is the transition probability function of the stochastic process X^β. Denote by $\mathbf{K}^\beta := [K^\beta(\omega, \omega')]_{(\omega,\omega') \in \Omega^2}$ the transition matrix of the process X^β. Assume that these probabilities vary continuously with the noise parameter β. For $\beta \to 0$ we obtain the *unperturbed Markov chain* $\mathcal{M} = (\Omega, \mathcal{F}, \mathbb{P}, (X_n)_{n \in \mathbb{N}_0})$, with corresponding transition matrix $\mathbf{K}$. Denote by $\mathcal{L}_1, \ldots, \mathcal{L}_k$ the k-recurrent classes of the unperturbed process $\mathcal{M}$, and $\Re = \bigcup_{\sigma=1}^k \mathcal{L}_\sigma$ the union of the recurrent classes of $\mathcal{M}$. By the decomposition theorem $\Omega = \mathcal{Q} \cup \Re$, where $\mathcal{Q}$ is the class of transient states in the unperturbed process.
Given the current state $\{X_n^\beta = \omega\}$, the following 3 events may take place:

Action adjustment: With probability $q_1(\omega) \in [0,1]$ the action profile α changes. Let $\nu \geq 0$ denote the rate with which player i receives an action revision opportunity.[4] Define the *volume* of the action adjustment process as $N\nu$. The probability that player i gets a revision opportunity is defined as $1/N$. Denote by $b^{i,\beta}(\cdot|\omega)$ a *probabilistic behavioral rule* describing how player i selects an action, given the population state ω. Specifically, assume that this behavioral rule satisfies the two "axioms":

[4] Assuming that this rate is heterogeneous is possible, but this is the basic assumption made in the literature.

(A1) For all $i \in \mathcal{I}$ and $\beta > 0$, $b^{i,\beta}(\cdot|\omega)$ is a full support distribution on $\mathcal{A}$.

(A2) For all $i \in \mathcal{I}$ there exists a *cost function* $c_1^i : \Omega^2 \to \mathbb{R}_+$, satisfying

$$-\lim_{\beta \to 0} \beta \log b^{i,\beta}(a|\omega) = c_1^i(\omega, (\alpha_i^a, g)). \tag{2.3}$$

This can be alternatively written as

$$b^{i,\beta}(a|\omega) = \exp\left[-\frac{1}{\beta}(c_1^i(\omega, (\alpha_i^a, g) + o(1))\right]$$

where $o(1)$ represents terms that go to 0 as $\beta \to 0$.

As $\beta \to 0$ the probability that player i makes a costly decision converges to 0 at exponential rate. A costless transition will be made even in the zero noise limit. Observe that the revision processes of Kandori et al. (1993) and Blume (1993), or adaptive learning of Young (1998) satisfies all these assumptions.[5]

Link creation: With unconditional probability $q_2(\omega)$ the process allows the network to expand. For all $i \in \mathcal{I}$ define a *rate function* $\lambda^i : \Omega \to \mathbb{R}_+$, satisfying $\kappa^i(\omega) = N - 1 \Rightarrow \lambda^i(\omega) = 0$. The *volume* of the link creation process is defined as the sum of all rate functions $\bar{\lambda}(\omega) := \sum_{i \in \mathcal{I}} \lambda^i(\omega)$. The conditional probability that player i receives the chance to form a link is $\lambda^i(\omega)/\bar{\lambda}(\omega)$. Conditional on this event, player i computes a tuple $w^{i,\beta}(\omega) := (w_j^{i,\beta}(\omega))_{j \in \mathcal{I}}$, satisfying:

(L1) If $g_{ij} = 0$ and $\beta > 0$, then $\min\{w_j^{i,\beta}(\omega), w_i^{j,\beta}(\omega)\} > 0$. If $g_{ij} = 1$ or $i = j$, then $w_j^{i,\beta}(\omega) = w_i^{j,\beta}(\omega) = 0$ for all β,

[5](A1) and (A2) are the most basic assumptions. An appealing additional requirement would be

(A3) $(\forall i \in \mathcal{I}) : c_1^i(\omega, (\alpha_i^a, g)) > 0$ iff $a \notin \arg\max_{a' \in \mathcal{A}} \pi^i(\alpha_i^a, g)$.

which says that only suboptimal choices have positive transition costs. In this sense, players use noisy best response rules (see Sandholm, 2009). However, for the general discussion such an assumption is not necessary.

(L2) $(\forall i \in \mathcal{I})(\forall \omega \in \Omega) : \sum_{j \in \mathcal{I}} w_j^{i,\beta}(\omega) = 1$

(L3) $(\forall i, j \in \mathcal{I})(\forall \omega \in \Omega) : -\lim_{\beta \to 0} \beta \log w_j^{i,\beta}(\omega) = c_2^i(\omega, (\alpha, g \oplus (i, j)))$.

$c_2^i : \Omega^2 \to \mathbb{R}_+$ is again a cost function for player i. Condition (L1) says that all neutral edges have a positive probability of becoming created for $\beta > 0$. This is an irreducibility assumption. (L3) is a large deviation assumption on the link creation probability.

Let $\mathbf{W}^\beta(\omega) = \bar{\lambda}(\omega)^{-1} \operatorname{diag}[\lambda^1(\omega), \dots, \lambda^N(\omega)][w_j^{i,\beta}]_{i,j \in \mathcal{I}}$ denote the matrix of link creation probabilities at state ω.[6] The i-th row of this matrix is $\left(\lambda^i(\omega)/\bar{\lambda}(\omega)\right) w^{i,\beta}(\omega)$.[7] Next, define the symmetric matrix $\bar{\mathbf{W}}^\beta(\omega) := [\bar{w}_{ij}^\beta(\omega)]_{i,j \in \mathcal{I}} = \mathbf{W}^\beta(\omega) + \mathbf{W}^\beta(\omega)^\top$.[8] The scalar $\bar{w}_{ij}^\beta(\omega)$ is the conditional probability that the passive edge (i, j) is formed, starting from ω.

Link destruction: With unconditional probability $q_3(\omega)$ a link becomes destroyed. Let $\xi \geq 0$ denote the constant rate of link destruction.[9] A positive level of volatility will imply that, independent of β, there is always a chance that a link becomes destroyed. Additionally to this drift term, let us assign to each edge (i, j) a weight $v_{ij}^\beta(\omega)$. The higher the weight of an active edge, the larger will be the conditional probability that it becomes destroyed. Let $\mathbf{V}^\beta(\omega) = [v_{ij}^\beta(\omega)]_{1 \leq i,j \leq N}$ the $N \times N$ matrix of edge weights, satisfying:

(D1) $\mathbf{V}^\beta(\omega)$ is a symmetric matrix, and, for $\beta > 0$, $v_{ij}^\beta(\omega) > 0$ if $g_{ij} = 1$, and $v_{ij}^\beta(\omega) = 0$ for $g_{ij} = 0$,

---

[6] $\operatorname{diag}[x_1, \dots, x_n]$ is the $n \times n$ diagonal matrix having x_i as entry in its i-th principal diagonal and 0 off the principal diagonal.

[7] Note that the above conditions on the distribution $w^{i,\beta}$ requires that a completely connected individual puts weight 1 one himself. This causes no trouble because such players do not get a link creation opportunity by default. Hence the algorithm produces simple graphs, i.e. graphs that have no multiple connections and self-loops, as desired.

[8] $\mathbf{W}^\top$ is the transposition of $\mathbf{W}$.

[9] This is exactly the volatility parameter of Marsili et al. (2004), Ehrhardt et al. (2006; 2008a;b).

(D2) $\sum_{i,j>i} v_{ij}^{\beta}(\omega) = 1$,

(D3) $(\forall i \in \mathcal{I})(\forall \omega \in \Omega) : -\lim_{\beta \to 0} \beta \log v_{ij}^{\beta}(\omega) = c_3^{(i,j)}(\omega, (\alpha, g \ominus (i,j)))$.

(D1) says that edges (i,j) and (j,i) are treated symmetrically. This is a natural assumption for undirected graphs. Moreover, it requires that all currently active edges are destroyed with positive probability if $\beta > 0$. (D2) requires that, conditional on the event of link destruction, the expected number of destroyed edges is 1. (D3) is our large deviation assumption. The *volume* of the link destruction process is defined as $\bar{\xi}(\omega) := \xi f(\omega, \mathbf{V}^{\beta})$, where $f(\cdot,\cdot)$ is a bounded non-negative function, normalized by the condition $f(\omega, \mathbf{V}^{\beta}) = 0$ if the network is the empty graph at ω.[10]

Let $\omega = (\alpha, g)$ be the current population state. Define

$$\Lambda(\omega) = N\nu + \bar{\lambda}(\omega) + \bar{\xi}(\omega). \tag{2.4}$$

By the frequency interpretation of probabilities, one can interpret the number $N\nu =: \tau_a$ as the time scale of action adjustment events, and $\bar{\lambda}(\omega) + \bar{\xi}(\omega) =: \tau_g$ as the time scale of network evolution. The ratio $\tau = \tau_g/\tau_a$ measures how fast network evolution is, relative to action adjustment. If τ is much larger than 1, network evolution will proceed at a faster time scale than action adjustment. If τ is much smaller than 1 then action adjustment opportunities arrive much more frequently to the population. The probabilities $q_\sigma(\omega), \sigma = 1,2,3$, specifying the timing of evolution, are defined as

$$q_1(\omega) = \frac{N\nu}{\Lambda(\omega)},\ q_2(\omega) = \frac{\bar{\lambda}(\omega)}{\Lambda(\omega)},\ q_3(\omega) = 1 - q_1(\omega) - q_2(\omega). \tag{2.5}$$

[10]The reason why a positive rate of link destruction is needed is to exclude trivial stationary states where all players are completely connected. Of course, assuming $\xi > 0$ does not exclude the complete graph of being a stationary state. Henceforth assume that $\xi > 0$ and fixed, so that β is the only varying parameter.

The elements of the transition matrix $\mathbf{K}^{\beta}$ are then given by

$$K^{\beta}(\omega,\omega') = \begin{cases} q_1(\omega)\frac{1}{N}b^{i,\beta}(a|\omega) & \text{if } \omega' = (\alpha_i^a, g), \\ q_2(\omega)\bar{w}_{i,j}^{\beta}(\omega) & \text{if } \omega' = (\alpha, g \oplus (i,j)), \\ q_3(\omega)v_{ij}^{\beta}(\omega) & \text{if } \omega' = (\alpha, g \ominus (i,j)), \\ 0 & \text{otherwise.} \end{cases} \tag{2.6}$$

It is easy to verify that $\sum_{\omega' \in \Omega} K^{\beta}(\omega,\omega') = q_1(\omega) + q_2(\omega) + q_3(\omega) = 1$ for all $\omega \in \Omega$. By the irreducibility assumptions (A1), (L1) and (D1), the matrix $\mathbf{K}^{\beta}$ is irreducible for $(\beta, \xi) \gg (0,0)$. Further, it is easy to see that the chain is aperiodic. Since Ω is a finite set, ergodicity of the process X^{β} is guaranteed. Hence, provided $\beta > 0$, there exists a unique invariant distribution $\mu^{\beta} \in \Delta(\Omega)$. It is well known that for $\beta \to 0$ the process concentrates on a subset of $\Re$. To classify such states, we use the following definition of stochastic stability. [11]

Definition 1 (Sandholm (2009)). *Given a co-evolutionary model with noise $\mathcal{M}^{\beta}$, we call a state $\omega \in \Omega$ stochastically stable if*

$$\lim_{\beta \to 0} \beta \log \mu^{\beta}(\omega) = 0. \tag{2.7}$$

Let Ω^* denote the set of stochastically stable states.

2.2 *On trees, graphs and stochastic stability*

At every point of time the process may undertake one of three different transitions. The most appealing way to think about the stochastic dynamic is in terms of directed graphs, as done by Kandori et al. (1993), Young (1993), building on the work of Freidlin and Wentzell (1998). Every

[11] Most models using stochastic evolutionary dynamics call a state stochastically stable if it receives *positive* weight in the limit distribution. Definition 1 says that ω is stochastically stable if $\log \mu^{\beta}(\omega) \to a < 0$ as $\beta \downarrow 0$. This is a weaker requirement than the conventional stochastic stability criterion, since it may well be that the mass converges to 0 at a sub-exponential rate. See Sandholm (2009, ch. 12), for a detailed discussion.

co-evolutionary model with noise $\mathcal{M}^\beta$ can be analyzed via directed graphs of the form $T = (\Omega, \vec{E})$. The vertex set of such graphs is the state space and the edge set is a subset of Ω^2. A graph T will be called a *revision graph*, and we will henceforth identify every revision graph with its edge set $\vec{E}(T)$ by taking the vertex set always to be Ω.

Definition 2. *Given a co-evolutionary model with noise $\mathcal{M}^\beta$ and a revision graph T, define the **reach** of state $\omega \in \Omega$ under T as the set*

$$\mathcal{R}_T(\omega) := \{\omega' \in \Omega | (\exists \vec{e} \in \vec{E}(T)) : \vec{e} = (\omega, \omega')\}.$$

The reach of a state is thus the collection of states that the process may visit after one step under the revision graph T, starting from ω. The reach of a state ω can be subdivided as follows; call $\mathcal{R}_{T,1}(\omega)$ the set of states in the reach of ω that differ in the action configuration, $\mathcal{R}_{T,2}(\omega)$ the set of states that are reachable from ω by creation of a single link, and finally $\mathcal{R}_{T,3}(\omega)$ the set of states reachable from ω by deleting a single link. [12] We will work with the following special class of revision graphs. Their role has also been emphasized by Samuelson (1997), Catoni (1999), Beggs (2005) and Alós-Ferrer and Netzer (2007).

Definition 3. *Consider a non-empty set $\mathcal{X} \subset \Omega$. A revision graph T is called a $\mathcal{X}$-revision graph if it is an element of the class of graphs $\mathcal{T}(\mathcal{X})$, satisfying*

(i) $(\forall \omega \in \Omega) : |\mathcal{R}_T(\omega)| = \mathbb{1}_{\{\omega \notin \mathcal{X}\}}$,

(ii) T does not contain a cycle.

A labeled ω-revision tree (T_ω, ℓ) is a $\{\omega\}$-revision graph $T_\omega \in \mathcal{T}(\{\omega\}) \equiv \mathcal{T}_\omega$ together with a labeling function $\ell : \vec{E}(T_\omega) \to \mathcal{I}^{(2)}$ satisfying

(iii) for all edges $\vec{e}$, $\ell(\vec{e})$ returns the pair of players (i, j) involved in the transition modeled by the edge $\vec{e} \in \vec{E}(T_\omega)$. If $j = i$ then we interpret the pair (i, i) as i.

[12]Obviously $\mathcal{R}_T(\omega) = \mathcal{R}_{T,1}(\omega) \cup \mathcal{R}_{T,2}(\omega) \cup \mathcal{R}_{T,3}(\omega)$.

A $\mathcal{X}$-revision graph $T \in \mathcal{T}(\mathcal{X})$ joins every point in $\Omega \setminus \mathcal{X}$ to $\mathcal{X}$, without loops. In the main text of the paper we will only need the concept of labeled revision trees. The more general concept of and $\mathcal{X}$-revision graph will be used in Appendix B. For this class of revision graphs, conditions (i) and (ii) are a version of the standard graph-constructs of Freidlin and Wentzell (1998), and merely assert that T_ω is a tree with root ω. The distinguishing point in the definition of a labeled revision tree is exactly the labeling function, whose purpose will become clear later on.[13] For a given ω-revision tree $(T_\omega, \ell) \in \mathcal{T}_\omega$, define the set

$$\mathcal{S}_{T_\omega,\sigma} := \{\vec{e} = (\omega', \omega'') \in \vec{E}(T_\omega) | \omega'' \in \mathcal{R}_{T_\omega,\sigma}(\omega')\},\ \sigma \in \{1,2,3\},$$

which is the collection of all edges used on a transition of type $\sigma \in \{1,2,3\}$. By definition we have $\vec{E}(T_\omega) = \bigcup_{\sigma=1}^{3} \mathcal{S}_{T_\omega,\sigma}$.
Following Freidlin and Wentzell (1998) we can now completely characterize the invariant distribution of the co-evolutionary process. With a slight abuse of notation define the numbers

$$K^\beta(T_\omega) := \prod_{\vec{e} \in \vec{E}(T_\omega)} K^\beta(\vec{e}) = \prod_{\sigma=1}^{3} \prod_{\vec{e} \in \mathcal{S}_{T_\omega,\sigma}} K^\beta(\vec{e}),$$

$$\rho^\beta(\omega) := \sum_{(T_\omega,\ell) \in \mathcal{T}_\omega} K^\beta(T_\omega).$$

Theorem 2.1 (The Markov chain tree theorem)**.** *For $\beta > 0$, the unique invariant distribution of the co-evolutionary model with noise $\mathcal{M}^\beta$ is given by*

$$(\forall \omega \in \Omega) : \mu^\beta(\omega) = \frac{\rho^\beta(\omega)}{\sum_{\omega' \in \Omega} \rho^\beta(\omega')}. \tag{2.8}$$

Proof. See Section B in the appendix. This follows immediately from Freidlin and Wentzell (1998) (Lemma 3.1, Chapter 6). This representation holds for *every* irreducible Markov chain, and is not restricted to the current model. See Young (1998) or Sandholm (2009) for alternative elegant proofs of this fact. □

[13]For the current type of stochastic process, the labeling function is uniquely defined for a given revision tree T_ω. See Alós-Ferrer and Netzer (2007) for a process where this need not be the case.

Consider a state $\omega \in \Omega$ with revision tree (T_ω, ℓ). By construction of the transition probabilities, for every edge $\vec{e} \in \mathcal{S}_{T_\omega,\sigma}, \sigma = 1,2,3$ there exists a *derived* cost function $\hat{c}_\sigma : \Omega^2 \to \mathbb{R}_+ \cup \{+\infty\}$, such that

$$K^\beta(\vec{e}) = \exp\left[-\frac{1}{\beta}(\hat{c}_\sigma(\vec{e}) + o(1))\right], \ \sigma \in \{1,2,3\},$$

depending on the type of transition under the edge $\vec{e}$.[14] If the transition $\vec{e} \in \mathcal{S}_{T_\omega,\sigma}$ is not possible for $\beta > 0$, then set $\hat{c}_\sigma(\vec{e}) = \infty$. Define the derived costs of a revision tree (T_ω, ℓ) as

$$\hat{C}(T_\omega) = \sum_{\sigma=1}^{3} \sum_{\vec{e} \in \mathcal{S}_{T_\omega,\sigma}} \hat{c}_\sigma(\vec{e}), \tag{2.9}$$

so that $K^\beta(T_\omega) = \exp\left[\frac{1}{\beta}(\hat{C}(T_\omega) + o(1))\right]$. The *stochastic potential* of state ω is the lowest cost of reaching it, i.e.

$$\gamma(\omega) := \min_{(T_\omega,\ell) \in \mathcal{T}_\omega} \hat{C}(T_\omega). \tag{2.10}$$

We are now ready to present a fairly general result characterizing the low-noise behavior of the invariant distribution (see also Catoni, 1999, Beggs, 2005).

Proposition 2.1. *Consider a co-evolutionary model with noise $\mathcal{M}^\beta$ with derived cost functions $\hat{c} = (\hat{c}_1, \hat{c}_2, \hat{c}_3)$ and invariant distribution μ^β. Let $\gamma : \Omega \to \mathbb{R}_+$ be the potential function defined in eq. (2.10). For all $\omega \in \Omega$ we have*

$$-\lim_{\beta \to 0} \beta \log \mu^\beta(\omega) = \gamma(\omega) - \min_{\omega' \in \Omega} \gamma(\omega'). \tag{2.11}$$

Before proving this proposition we need some additional facts. Order the factors in the invariant measure ρ^β according to their leading terms as

[14] Derived cost functions will be used in this paper only for the link creation process. In the action revision process one would also need a derived cost function to account for the unlikelihood of a transition when one would apply the learning model of Alós-Ferrer and Netzer (2007) (see their concept of the waste of a labeled revision tree).

$\beta \to 0$. This leads to the low-noise expression

$$\rho^\beta(\omega) = \sum_{(T_\omega,\ell)\in\mathcal{T}_\omega} \exp\left[-\frac{1}{\beta}\left(\hat{C}(T_\omega) + o(1)\right)\right]$$
$$= B_\omega \exp(-\gamma(\omega)/\beta)(1 + o(1))$$

where B_ω is a real constant. For sufficiently small β, the invariant distribution can therefore be written as

$$\mu^\beta(\omega) = \frac{B_\omega \exp(-\gamma(\omega)/\beta)(1+o(1))}{\sum\limits_{\omega'\in\Omega} B_{\omega'} \exp(-\gamma(\omega')/\beta)(1+o(1))}. \tag{2.12}$$

The following simple fact is a useful intermediate result, as it identifies the leading term of the denominator in (2.12).

Lemma 2.1. *Given two finite sequences* $(f(1),\ldots,f(n))$, $(B_1,\ldots,B_n)$ *of non-negative real numbers, then*

$$\lim_{\beta\to 0} \frac{\log\left(\sum\limits_{i=1}^{n} B_i \exp(-f(i)/\beta)\right)}{\max\limits_{i=1}^{n} \log(B_i \exp(-f(i)/\beta))} = 1. \tag{2.13}$$

Proof. Without loss of generality, let $f(n) = \min\limits_{i=1}^{n} f(i)$. By absorbing states with equal values of $f(i)$ in the constant B_i we can, without loss of generality, assume that all values are different. The denominator is thus $\log(B_n \exp(-f(n)/\beta))$. Write the polynomial inside $\log(\cdot)$ in the numerator by collecting the terms of highest order, i.e.

$$\sum_{i=1}^{n} B_i \exp(-f(i)/\beta) = B_n \exp(-f(n)/\beta)\left(1 + \sum_{i=1}^{n-1} \frac{B_i}{B_n} \exp(-(f(i)-f(n))/\beta)\right)$$
$$= B_n \exp(-f(n)/\beta) r(\beta)$$

and $\beta \log r(\beta) \to 0$ as $\beta \to 0$. Hence, the ratio (2.13) can be written as

$$\frac{\beta \log B_n + \beta \log r(\beta) - f(n)}{\beta \log B_n - f(n)} \to 1, \text{as } \beta \to 0.$$

□

Proof of Proposition 2.1. Start from eq. (2.12). Take logarithms and multiply both sides by $-\beta$ to arrive at

$$-\beta \log \mu^{\beta}(\omega) = -\beta \log(B_{\omega} \exp(-\gamma(\omega)/\beta)) + \beta \log \left(\sum_{\omega' \in \Omega} B_{\omega'} \exp(-\gamma(\omega')/\beta) \right) + O(\beta).$$

The claim now follows from Lemma 2.1. □

This shows that a state is stochastically stable according to Definition 1 iff it is a state with minimal stochastic potential.

Corollary 2.1. $\Omega^* = \{\omega \in \Omega | \gamma(\omega) = \min_{\omega' \in \Omega} \gamma(\omega')\}$.

We see that the main difference between a co-evolutionary model with noise and a classical evolutionary model is the addition of two further cost functions, corresponding to the two added processes modeling the evolution of the network. Departing from here it is easy to see that all well-known results on stochastic stability are applicable. Referring to Appendix B for proofs of these facts, we just introduce some concepts in order to fix the notation.[15] Let $\mathcal{X}, \mathcal{X}'$ be some non-empty subsets of Ω. A (ω, ω')-path is a directed graph whose vertex set is a non-repeating sequence of states $\{\omega_1, \ldots, \omega_l\}$ such that $\omega_1 = \omega, \omega_l = \omega', \omega_t \notin \mathcal{X}', \forall t \in [2, l-1]$, and whose edges are the transitions $(\omega_i, \omega_{i+1}), 1 \leq i \leq l-1$. Denote by $\mathcal{P}_{\omega,\omega'}(\mathcal{X}, \mathcal{X}')$ the set of paths connecting ω to ω', and P one such path. To each ω, ω'-path there corresponds a labeling function ℓ, as in Definition 3. For $P \in \mathcal{P}_{\omega,\omega'}(\mathcal{X}, \mathcal{X}')$ let (P, ℓ) denote a (ω, ω')-*revision path*. Let $\mathcal{L}, \mathcal{L}'$ be two recurrent classes of the process $\mathcal{M}$, and denote the cost of transition from recurrent class $\mathcal{L}$ to $\mathcal{L}'$ by $C(\mathcal{L}, \mathcal{L}')$. The cost of a transition from recurrent class $\mathcal{L}$ to $\mathcal{L}'$ is

$$C(\mathcal{L}, \mathcal{L}') = \min_{\omega \in \mathcal{L}} \min_{\omega' \in \mathcal{L}'} \min_{(P,\ell): P \in \mathcal{P}_{\omega,\omega'}(\mathcal{L}, \mathcal{L}')} \hat{C}(P), \tag{2.14}$$

[15]See also Samuelson (1997), Young (1998) or Sandholm (2009) for textbook treatments of this, or Ellison (2000).

where $\hat{C}(P)$ is defined as in (2.9), applied to a (ω, ω')-revision path. In Appendix B we show that all states within one recurrent class are connected by a null cost path. This allows one to study revision graphs between recurrent classes. Therefore, we introduce the class of revision graphs $\hat{T} = (\{\mathcal{L}_1, \ldots, \mathcal{L}_k\}, \vec{E})$, where $\vec{E}(\hat{T}) \subseteq \{\mathcal{L}_1, \ldots, \mathcal{L}_k\}^2$. A $\mathcal{L}$-revision tree $\hat{T} \in \hat{\mathcal{T}}(\mathcal{L})$ is a revision graph in the sense of Definition 3, but acting on the recurrent sets of the unperturbed co-evolutionary process $\mathcal{M}$.[16] The costs of such a revision tree are defined as $C(\hat{T}) = \sum_{\vec{e} \in \vec{E}(\hat{T})} C(\vec{e})$, with $\vec{e} = (\mathcal{L}', \mathcal{L}'')$. Letting $\hat{\gamma} : \Re \to \mathbb{R}_+$ be a potential function on the set of recurrent classes, one can show (see Appendix B) that for all $\omega \in \mathcal{L}$

$$\gamma(\omega) = \hat{\gamma}(\mathcal{L}) = \min_{\hat{T} \in \hat{\mathcal{T}}(\mathcal{L})} C(\hat{T}). \tag{2.15}$$

3 Applications

In this section we apply the above general framework to some recent models. In both models we consider the base game $\Gamma = (\mathcal{I}, \{a_1, a_2\}, u)$, with normal form

	a_1	a_2
a_1	$(e-\phi, e-\phi)$	$(f-\phi, g-\phi)$
a_2	$(g-\phi, f-\phi)$	$(h-\phi, h-\phi)$

(3.1)

Assume that $h > e > f > g$ but $e + f > h + g$. This means that (a_1, a_1) is a risk-dominant strict Nash equilibrium, while (a_2, a_2) is a Pareto efficient strict Nash equilibrium. The number $\phi \geq 0$ is a fee two incident players have to pay in order to play the game. It does not alter the nature of the game, but possibly affects the way how players form their social network. There is also a mixed strategy equilibrium where a_1 is played with probability $x = \frac{h-f}{e-g+h-f} < 1/2$.

[16]This method of reducing the process to recurrent classes and monitoring transitions only between them has been proposed by Young (1993).

3.1 A volatility model

In Staudigl (2009b) a volatility model for general potential games is presented. Here we study a version of this model in the context of the symmetric coordination game (3.1) with $\phi = 0$. The co-evolutionary model with noise $\mathcal{M}^{\beta}$ is the following;

Action adjustment: Assume that

$$b^{i,\beta}(a_{\sigma}|\omega) = \frac{\exp(\pi^{i}(\alpha_{i}^{a_{\sigma}}, g)/\beta)}{\sum_{r=1}^{2} \exp(\pi^{i}(\alpha_{i}^{a_{r}}, g)/\beta)}, \quad \sigma = 1, 2. \tag{3.2}$$

This behavioral rule satisfies (A2) with cost function

$$\hat{c}_{1}(\omega, (\alpha_{i}^{a}, g)) = c_{1}^{i}(\omega, (\alpha_{1}^{a}, g)) = \max_{a' \in \mathcal{A}} \pi^{i}(\alpha_{i}^{a'}, g) - \pi^{i}(\alpha_{i}^{a}, g). \tag{3.3}$$

Link creation: Assume that $\lambda^{i}(\omega) = \lambda \mathbb{1}_{\{\kappa^{i}(\omega) < N-1\}}$, so that every incompletely connected player receives a link creation opportunity with rate $\lambda \geq 0$. Conditional on this event player i samples player j with probability

$$w_{j}^{i,\beta}(\omega) = \frac{\exp(u(\alpha^{i}, \alpha^{j})/\beta)}{\sum_{k \notin \bar{\mathcal{N}}^{i}(\omega)} \exp(u(\alpha^{i}, \alpha^{k})/\beta)}.$$

(L4) is satisfied with cost function

$$c_{2}^{i}(\omega, (\alpha, g \oplus (i,j)) = \max_{k \notin \bar{\mathcal{N}}^{i}(\omega)} u(\alpha^{i}, \alpha^{k}) - u(\alpha^{i}, \alpha^{j}).$$

Link destruction: Once a link is selected by the process (an event with rate ξ) it becomes destroyed at rate 1. Hence $v_{ij}^{\beta}(\omega) = \frac{g_{ij}}{e(\omega)}$. (D3) is satisfied with

$$\hat{c}_{3}(\omega, (\alpha, g \ominus (i,j))) = c_{3}^{(i,j)}(\omega, (\alpha, g \ominus (i,j))) \equiv 0.$$

The volume of this subprocess is given by $\bar{\xi} = \xi e(\omega)$, so that $f(\omega, \mathbf{V}^{\beta}) = e(\omega)$.[17]

[17]In Staudigl (2009a) the agents have idiosyncratic preferences over the actions, which

It remains to determine the derived cost function $\hat{c}_2$. When a link becomes created, a pair of players (i, j) is involved with $j > i$. Suppose this event is on the ω-revision tree (T_ω, ℓ), and call the edge of transition corresponding to this event $\vec{e}$. The labeling function returns the pair of players $\ell(\vec{e}) = (i, j)$. Let $\ell(\vec{e})^-$ be the player with the lower index involved in the transition $\vec{e}$, i.e. i, and $\ell(\vec{e})^+$ the player with the higher index, i.e. j.[18]

Lemma 3.1. *For every $\omega \in \Omega$ and $(T_\omega, \ell) \in \mathcal{T}_\omega$, the derived cost of a transition $\vec{e} \in \mathcal{S}_{T_\omega,2}$ is*

$$\hat{c}_2(\vec{e}) = \min\{c_2^{\ell(\vec{e})^-}(\vec{e}), c_2^{\ell(\vec{e})^+}(\vec{e})\}. \tag{3.4}$$

Proof. The probability that edge (i, j) becomes created is

$$\bar{w}_{ij}^\beta = \frac{\lambda}{\bar{\lambda}(\omega)}(w_j^{i,\beta}(\omega) + w_i^{j,\beta}(\omega)).$$

By the large deviation principle (L4), for small β we have

$$\begin{aligned} w_j^{i,\beta}(\omega) + w_i^{j,\beta}(\omega) = \exp&\left[-\frac{1}{\beta}(c_2^i(\omega, (\alpha, g \oplus (i, j))) + o(1))\right] \\ &+ \exp\left[-\frac{1}{\beta}(c_2^j(\omega, (\alpha, g \oplus (i, j))) + o(1))\right], \end{aligned}$$

and so we can apply Lemma 2.1, which gives us the desired result. □

Thus, for every $\omega \in \Omega$ the cost of a revision tree $(T_\omega, \ell) \in \mathcal{T}_\omega$ is $\hat{C}(T_\omega) = \sum_{\sigma=1}^2 \sum_{\vec{e} \in \mathcal{S}_{T_\omega,\sigma}} \hat{c}_\sigma(\vec{e})$.

is interpreted as the "type" of the agent. Link decay probabilities are then functions of the types of the involved players. Particularly, it is assumed that $v_{ij}^\beta(\omega) = \hat{\xi}_{ij}^\beta / \sum_{k>l} \hat{\xi}_{kl}^\beta g_{kl}$, for given functions $\{\hat{\xi}_{ij}^\beta\}$, which depend on the realized types of the agents and on $\beta > 0$. The corresponding volume is now $\bar{\xi}(\omega) = \sum_{j>i} \hat{\xi}_{ij}^\beta g_{ij}$.

[18] I thank Stefano DeMichelis for giving me the right hint for the following proof of the following Lemma.

3.1.1 Recurrent classes and stochastic stability

Define the set

$$\tilde{\Omega} = \{\omega \in \Omega | g_{ij} = 1 \Rightarrow \alpha^i = \alpha^j\}.$$

A network in this set has only edges between two coordinated players. It may have several connected components and, in particular, it may not be completely connected. Distinguished classes of states in $\tilde{\Omega}$ are the *global conformity sets*

$$\mathcal{L}_\sigma = \{\omega \in \tilde{\Omega} | (\forall i \in \mathcal{I}) : \alpha^i = a_\sigma\}, \ \sigma = 1, 2.$$

Let $\mathcal{L}_{1,2} = \tilde{\Omega} \setminus (\mathcal{L}_1 \cup \mathcal{L}_2)$, the *co-existence set*. The following Lemma characterizes the recurrent classes $\Re$ of the unperturbed process.

Lemma 3.2. *Consider the unperturbed co-evolutionary model $\mathcal{M}$ of Staudigl (2009b). We have $\Re = \tilde{\Omega}$.*

Proof. The proof proceeds by a fairly general constructive argument, which is presented in Appendix A. □

We see that the process allows for *global heterogeneity*, since there may be multiple connected components displaying different types of conventions. However, within every connected component we must have *local conformity*. Due to this large number of equilibria, we hope that the concept of stochastic stability gives us some hint which states are more likely to be observed in the long-run. The following proposition shows that this is not the case.

Proposition 3.1. *Consider the coordination game (3.1) with $\phi = 0$, and the co-evolutionary model with noise $\mathcal{M}^\beta$ of Staudigl (2009b). We have $\Omega^* = \Re$.*

Proof. Fix $\omega \in \mathcal{L}_1, \omega' \in \mathcal{L}_2$. We will construct a zero cost path $P \in \mathcal{P}_{\omega,\omega'}(\mathcal{L}_1, \mathcal{L}_2)$, which implies $C(\mathcal{L}_1, \mathcal{L}_2) = 0$. A symmetric argument shows that $C(\mathcal{L}_2, \mathcal{L}_1) = 0$, so that $\gamma(\omega) = \gamma(\omega') = 0$. From this it follows that $\gamma(\omega'') = 0$ for all $\omega'' \in \mathcal{L}_{1,2}$, since all paths from ω to ω' must pass through some state $\omega'' \in \mathcal{L}_{1,2}$.

Step 1: From ω apply a sequence of link destruction events. All this has zero costs and in finitely many steps we arrive at state $\hat{\omega} \in \mathcal{L}_1$ with the empty network.

Step 2: Give two randomly chosen players sequentially an action adjustment opportunity where they switch to a_2. This has zero costs, since a loner selects both actions with equal probability.

Step 3: Give one of the two players a link creation opportunity. Under **K** a link between them will be established. We are now at a state in $\mathcal{L}_{1,2}$.

Step 4: Give the remaining players action adjustment opportunities where they switch to a_2, and then a link creation opportunity. Iterate this until the process arrives at the desired state $\omega' \in \mathcal{L}_2$.

Steps 1-4 defines a path from ω to ω' having zero costs. Clearly all steps are reversible, i.e. steps 4-1 define a path from ω' to ω having zero costs. This demonstrates $\gamma(\omega) = \gamma(\omega') = 0$. □

3.2 *A classical model*

We discuss a slight variation of Jackson and Watts (2002). To get the most interesting scenario, we reduce the set of admissible parameters in requiring that $x > \frac{1}{N-1}$ and $\phi \in (g, e)$. This paper takes the mistakes model of Kandori et al. (1993) and Young (1993) as universal behavioral rule. Parameterizing noise as $\varepsilon = \exp(-1/\beta)$, and henceforth calling ε the noise parameter, allows us to study this behavioral rule. The co-evolutionary model with noise $\mathcal{M}^\varepsilon$ is as follows:

Action adjustment: Assume that each player receives with uniform probability $1/N$ the opportunity to change his action. Conditional on this event he selects action $a \in \mathcal{A}$ with probability

$$b^{i,\varepsilon}(a|\omega) = \begin{cases} 1 - \frac{\varepsilon}{2} & \text{if } \alpha^i \neq a \text{ and } \{a\} = \arg\max_{a' \in \mathcal{A}} \pi^i(\alpha_i^{a'}, g), \\ 1 - \frac{\varepsilon}{2} & \text{if } \alpha^i = a \text{ and } \{\alpha^i\} = \arg\max_{a' \in \mathcal{A}} \pi^i(\alpha_i^{a'}, g), \\ \frac{\varepsilon}{2} & \text{otherwise.} \end{cases}$$

This behavioral rule says that a player abandons his currently used action with relatively high probability, if there exists a strictly better action. Otherwise he sticks to his action and switches only with the relatively small probability ε. This behavioral rule satisfies condition (A2) with cost function

$$\hat{c}_1(\omega, (\alpha_i^a, g)) = c_1^i(\omega, (\alpha_i^a, g)) = \begin{cases} 0 & \text{if } \alpha^i \neq a \text{ and } \{a\} = \arg\max_{a' \in \mathcal{A}} \pi^i(\alpha_i^{a'}, g), \\ 0 & \text{if } \alpha^i = a \text{ and } \{\alpha^i\} = \arg\max_{a' \in \mathcal{A}} \pi^i(\alpha_i^{a'}, g), \\ 1 & \text{otherwise.} \end{cases}$$

Link creation: Jackson and Watts (2002) introduce a cooperative element into the link creation process. To capture this, we have to make a slight modification in the construction of our co-evolutionary model with noise. Let $\mathcal{D}(\omega)$ denote the set of neutral links at ω and $d(\omega)$ its cardinality. Instead of the individual players' rate functions, assume that the event of link creation arrives to the society at the constant rate $\bar{\lambda}(\omega) := \lambda d(\omega)$, where λ is a positive constant. Define the events

$$(1 \leq i, j \leq N) : A_j^i(\phi) := \{\omega \in \Omega | u(\alpha^i, \alpha^j) \geq \phi\}.$$

If $\omega \in A_j^i(\phi)$ then the edge (i, j) is profitable from the point of view of player i at ω. The number of mutually profitable neutral links is

$$m(\omega) = \sum_{(i,j) \in \mathcal{D}(\omega)} \mathbb{1}_{A_j^i(\phi) \cap A_i^j(\phi)}(\omega).$$

Following the spirit of pairwise stability (Jackson and Wolinsky, 1996), assume that a neutral link is set to be active with probability $1 - \varepsilon$ if both players mutually agree. With the small probability ε

assume that all links have a chance to be formed. The (conditional) probability that a neutral edge (i, j) will be added is

$$(\forall (i, j) \in \mathcal{D}(\omega)) : \bar{w}^{\varepsilon}_{ij}(\omega) := \begin{cases} \frac{1-\varepsilon}{m(\omega)} + \frac{\varepsilon}{d(\omega)} & \text{if } \mathbb{1}_{A^i_j(\phi) \cap A^j_i(\phi)}(\omega) = 1, \\ \frac{\varepsilon}{d(\omega)} & \text{otherwise.} \end{cases} \tag{3.5}$$

The term $\varepsilon/d(\omega)$ is the "background noise" of the system, and gives the uniform probability that a link will be formed. If edge (i, j) is neutral at ω, but both players are not hurt by the creation of the link, then they will independently agree to form it with the high probability $1 - \varepsilon$, which increases their chance of being formed.[19] Let $\bar{A}$ denote the complementary set of A. The cost function of this sub-process is

$$\hat{c}_2(\omega, (\alpha, g \oplus (i, j))) = \mathbb{1}_{\bar{A}^i_j(\phi) \cup \bar{A}^j_i(\phi)}(\omega).$$

Link destruction: With rate $\xi > 0$ links become destroyed. Conditional on this event, pick one edge $(i, j) \in \mathcal{E}(\omega)$ with uniform probability, and allow the incident players to re-evaluate the benefits arising from this connection. This leads to $\bar{\xi}(\omega) := \xi e(\omega)$. Denote by

$$\bar{m}(\omega) = \sum_{(i,j) \in \mathcal{E}(\omega)} \mathbb{1}_{\bar{A}^i_j(\phi) \cup \bar{A}^j_i(\phi)}(\omega)$$

the number of active links where at least one player benefits from the deletion of the link. If (i, j) is a link where at least one player is better off after its destruction, suppose that with large probability

[19] Jackson and Watts (2002) assume that a link is created with probability $1 - \varepsilon$ iff it is a strict Pareto improvement, i.e. at least one player is strictly better off after the connection has been established and no player is hurt from the creation of the link. We assume that a link is already formed if it is a weak Pareto improvement, i.e. no party is hurt by the formation of the link. Additionally, they assume that the error probability ε is not the same in the action evolution process as it is in the graph evolution process. However, it is required that the error probabilities go to zero at the same rate so not to get "twisted" equilibrium selection results, as argued by Bergin and Lipman (1996).

$1-\varepsilon$ it will be destroyed. With the small probability ε every link can be destroyed once it has been selected. This leads to the following version of link destruction probabilities

$$(\forall (i,j) \in \mathcal{E}(\omega)) : v^{\varepsilon}_{i,j}(\omega) = \begin{cases} \frac{1-\varepsilon}{\bar{m}(\omega)} + \frac{\varepsilon}{e(\omega)} & \text{if } \mathbb{1}_{\bar{A}^i_j(\phi) \cup \bar{A}^j_i(\phi)}(\omega) = 1, \\ \frac{\varepsilon}{e(\omega)} & \text{otherwise.} \end{cases}$$

The cost function of this process is given by

$$\hat{c}_3(\omega, (\alpha, g \ominus (i,j))) = \mathbb{1}_{A^i_j(\phi) \cap A^j_i(\phi)}(\omega).$$

3.2.1 Recurrent classes and stochastically stable states

Define $\mathcal{I}_1(\omega) = \{i \in \mathcal{I} | \alpha^i = a_1 \text{ on } \omega\}$, and for every $2 \leq n \leq N-2$,

$$\mathcal{L}^n_{1,2} = \{\omega \in \Omega | g = g^c[\mathcal{I}_1(\omega)] \oplus g^c[\mathcal{I}_2(\omega)] \;\&\; |\mathcal{I}_1(\omega)| = n\},$$
$$\mathcal{L}_{1,2} = \bigcup_{n=2}^{N-2} \mathcal{L}^n_{1,2}.$$

Let

$$\mathcal{L}_\sigma = \{\omega \in \Omega | (\forall i \in \mathcal{I}) : \alpha^i = a_\sigma \;\&\; g = g^c[\mathcal{I}]\},\ \sigma = 1,2.$$

Lemma 3.3. *Let $\mathcal{M}$ be the unperturbed co-evolutionary process of Jackson and Watts (2002) with $\phi \in (g, e)$. Then*

$$\Re = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \mathcal{L}_{1,2}.$$

Proof. The algorithm in appendix A shows that in finite time there are no links between agents playing different actions. Call ω_m the state at which the algorithm stops. In the unperturbed model, with probability 1, only links which are mutually profitable are formed and links which harm at least one player are destroyed. At ω_m no player has an incentive to change his action. If a link creation event takes place, with conditional probability 1 only an edge is formed if the selected pair is in the same action class $\mathcal{I}_\sigma(\omega_m), \sigma = 1,2$. Moreover, these links never become destroyed. A link destruction event at ω_m leaves the state invariant with probability 1, since at this state no edges between players from different action classes exist. For the same reason, an action adjutment event leaves ω_m invariant, with probability 1.If $\mathcal{I}_\sigma(\omega_m) = \emptyset$ for a $\sigma = 1, 2$, then the process arrives at a state where global conformity prevails. Otherwise the process leads to a state in the co-existence set $\mathcal{L}_{1,2}$. □

To select among the recurrent sets, we now perform an analysis via stochastic stability.

Proposition 3.2. *Let $\mathcal{M}^\beta$ be the co-evolutionary model with noise of Jackson and Watts (2002). Then for $\phi \in (g, e)$ and $x \geq \frac{1}{N-1}$*

$$\Omega^* = \mathcal{L}_1 \cup \mathcal{L}_2.$$

Proof. We will explicitly calculate the potentials of the three recurrent classes and show that, under the stated parameter assumptions, $\hat{\gamma}(\mathcal{L}_1) = \hat{\gamma}(\mathcal{L}_2) < \hat{\gamma}(\mathcal{L}^n_{1,2})$ for all $2 \leq n \leq N-2$.
Let $\omega \in \mathcal{L}_1, \omega' \in \mathcal{L}_2$. Observe that, under the assumption $x \geq \frac{1}{N-1}$, at least two agents must change their action to enter $\mathcal{L}_{1,2}$ via an unperturbed move, i.e. $C(\mathcal{L}_\sigma, \mathcal{L}_{1,2}) = 2, \sigma = 1, 2$.[20] Further observe that for $3 \leq j \leq N-3, C(\mathcal{L}^j_{1,2}, \mathcal{L}^{j\pm1}_{1,2}) = 1$, since a single deviator reduces/increases the set of a_1 players, and applying then link creation/destruction leads to some

[20] To see this, note that $C(\mathcal{L}_1, \mathcal{L}^2_{1,2}) = 2$ by definition of risk dominance. To get $C(\mathcal{L}_2, \mathcal{L}^{N-2}_{1,2}) = 2$ suppose that one player deviates from $\omega' \in \mathcal{L}_2$ and plays a_1. The network remains unchanged. Apply the action adjustment process in the next period to a current a_2 player. This player will switch to a_1 iff $e + (N-2)f - \phi(N-1) > g + (N-2)h + \phi(N-1)$, or iff $x < 1/(N-1)$. Since we assume that $x \geq 1/(N-1)$ another tremble is needed to make a_1 a best response.

state in the desired recurrent class. It follows that

$$
\begin{aligned}
\hat{\gamma}(\mathcal{L}_1) &= C(\mathcal{L}_2, \mathcal{L}_{1,2}^{N-2}) + \sum_{j=2}^{N-3} C(\mathcal{L}_{1,2}^{j+1}, \mathcal{L}_{1,2}^{j}) + C(\mathcal{L}_{1,2}^{2}, \mathcal{L}_1) \\
&= 2 + (N-3-2+1) = N-2 \\
\hat{\gamma}(\mathcal{L}_2) &= C(\mathcal{L}_1, \mathcal{L}_{1,2}^{2}) + \sum_{j=2}^{N-3} C(\mathcal{L}_{1,2}^{j}, \mathcal{L}_{1,2}^{j+1}) + C(\mathcal{L}_{1,2}^{N-2}, \mathcal{L}_2) \\
&= N-2 \\
\hat{\gamma}(\mathcal{L}_{1,2}^{n}) &= C(\mathcal{L}_1, \mathcal{L}_{1,2}^{2}) + C(\mathcal{L}_2, \mathcal{L}_{1,2}^{N-2}) \\
&\quad + \sum_{j=2}^{n-1} C(\mathcal{L}_{1,2}^{j}, \mathcal{L}_{1,2}^{j+1}) + \sum_{j=2}^{N-n-1} C(\mathcal{L}_{1,2}^{N-j}, \mathcal{L}_{1,2}^{N-j-1}) \\
&= 2 + 2 + (n-1-2+1) + (N-n-1-2+1) = N
\end{aligned}
$$

□

4 A micro-founded model for inhomogeneous random graphs.

The theory of random graphs provides in essence 2 classes of models; the "randomly grown graphs", mostly using a version of preferential attachment (Barabási and Albert, 1999), and generalized random graphs (Newman, 2003). Under some additional assumptions on the structure of the Markov chain (2.6), we are able to characterize the induced ensemble of random graphs for general behavioral rules.

Let us add the following two assumptions on the structure of the transition probabilities:

(L4) $(\forall i \in \mathcal{I}) : \lambda^i(\omega) = \lambda \mathbb{1}_{\{\kappa^i(\omega) < N-1\}}$;

(L5) $(\forall i.j \in \mathcal{I}) : w_j^{i,\beta}(\omega) = \hat{w}_j^{i,\beta}(\alpha)(1 - g_{ij})$, where $\hat{w}_j^{i,\beta}(\cdot)$ satisfies (L1)-(L3).

(L4) defines the volume of the link creation subprocess as $\bar{\lambda}(\omega) = \lambda \sum_{i \in \mathcal{I}} \mathbb{1}_{\{\kappa^i(\omega) < N-1\}}$. In the link destruction process we demand additionally

(D4) $(\forall i, j \in \mathcal{I}) : v^{\beta}_{ij}(\omega) = \frac{\hat{v}^{\beta}_{ij}(\alpha) g_{ij}}{f(\omega, \mathbf{V}^{\beta})}$.

(D2) tells us that $f(\omega, \mathbf{V}^{\beta}) = \sum_{j>i} \hat{v}^{\beta}_{ij}(\alpha) g_{ij}$.

Using these additional assumptions we will derive a *random graph process*, modeling the evolution of the network for a fixed action profile α.[21] Let $(\tilde{G}^{\beta}_n)_{n=0}^{\infty}$ denote the random graph process with transition probabilities $K^{\beta}_{2,3} : \mathcal{G}[\mathcal{I}] \times \mathcal{G}[\mathcal{I}] \to [0,1]$, defined as

$$\begin{aligned} K^{\beta}_{2,3}(g, g') &= \mathbb{P}(\tilde{G}^{\beta}_{n+1} = g' | \tilde{G}^{\beta}_n = g) \\ &= \mathbb{P}(X^{\beta}_{n+1} = (\alpha, g') | X^{\beta}_n = (\alpha, g), \text{Network evolution}) \\ &= \frac{1}{\mathbb{P}(\text{Network evolution} | X^{\beta}_n = (\alpha, g))} \mathbb{P}(X^{\beta}_{n+1} = (\alpha, g') | X^{\beta}_n = (\alpha, g)) \\ &= \frac{1}{q_2(\alpha, g) + q_3(\alpha, g)} K^{\beta}((\alpha, g), (\alpha, g')) \end{aligned}$$

Using (2.5), (2.6) and (L4), (L5), (D4), the transition probabilities are given by

$$K^{\beta}_{2,3}(g, g') = \frac{1}{q_2(\alpha, g) + q_3(\alpha, g)} \begin{cases} \lambda(\hat{w}^{i,\beta}_j(\alpha) + \hat{w}^{j,\beta}_i(\alpha)) & \text{if } g' = g \oplus (i,j), \\ \xi \hat{v}^{\beta}_{ij}(\alpha) & \text{if } g' = g \ominus (i,j), \\ 0 & \text{otherwise.} \end{cases}$$

This chain is irreducible but no longer aperiodic. It serves as a jump chain of the *continuous-time* random graph process $(\tilde{G}^{\beta}(t))_{t \geq 0}$ with generator[22]

$$\eta^{\beta}_{2,3}(g \to g') = (q_2(\alpha, g) + q_3(\alpha, g))(K^{\beta}_{2,3}(g, g') - \delta_{g,g'}) \tag{4.1}$$

where $\delta_{g,g'} = 1$ iff $g = g'$, and 0 otherwise. This continuous time process allows us to identify the invariant distribution of the original

[21] An interpretation of such a process can be given by assuming that action adjustment is a relatively fast process compared to network evolution. In this case, it makes sense to assume that the profile α reaches a temporary stationary state for a given network g, and when evolution shapes the network the profile α is fixed.

[22] See Norris (1997).

Markov chain in a simple way. Let **Id** denote the identity matrix on $\mathcal{G}[\mathcal{I}]$, and define the matrix $\boldsymbol{\eta}^{\beta}_{2,3} := \left[\eta^{\beta}_{2,3}(g \to g')\right]_{g,g' \in \mathcal{G}[\mathcal{I}]}$. Additionally, call $\hat{q}(g) := q_2(\alpha, g) + q_3(\alpha, g)$, and $\hat{\boldsymbol{q}} := [\hat{q}_2(g)]_{g \in \mathcal{G}[\mathcal{I}]}$. The generator of the continuous-time process $(\tilde{G}^{\beta}(t))_{t \geq 0}$ is defined by $\boldsymbol{\eta}^{\beta}_{2,3} = \hat{\boldsymbol{q}}(\mathbf{K}^{\beta}_{2,3} - \mathbf{Id})$. A measure ν on $\mathcal{G}[\mathcal{I}]$ is said to be invariant under the generator $\boldsymbol{\eta}^{\beta}_{2,3}$ if $\nu \boldsymbol{\eta}^{\beta}_{2,3} = \mathbf{0}$.

Lemma 4.1. *The following are equivalent:*

(a) ν is invariant under $\boldsymbol{\eta}^{\beta}_{2,3}$,

(b) $\mu \mathbf{K}^{\beta}_{2,3} = \mu$ where $\mu(g) = \nu(g)\hat{q}(g)$.

Proof. Define the measure $\mu(g) := \nu(g)\hat{q}(g)$ for all $g \in \mathcal{G}[\mathcal{I}]$. For all g, g' we have $\eta^{\beta}_{2,3}(g \to g') = \hat{q}(g)(K^{\beta}_{2,3}(g, g') - \delta_{g,g'})$. Thus,

$$\begin{aligned}\sum_{g \in \mathcal{G}[\mathcal{I}]} \mu(g)(K^{\beta}_{2,3}(g, g') - \delta_{g,g'}) &= \sum_{g \in \mathcal{G}[\mathcal{I}]} \nu(g)\hat{q}(g)(K^{\beta}_{2,3}(g, g') - \delta_{g,g'}) \\ &= \sum_{g \in \mathcal{G}[\mathcal{I}]} \nu(g)\eta^{\beta}_{2,3}(g \to g').\end{aligned}$$

□

The next proposition characterizes the invariant distribution of the continuous-time random graph process.

Proposition 4.1. *Consider the random graph process $(\tilde{G}^{\beta}(t))_{t \geq 0}$ with generator (4.1). Its unique invariant distribution is the product measure*

$$\hat{\mu}^{\beta}_{2,3}(g|\alpha) = \prod_{i=1}^{N} \prod_{j>i} p^{\beta}_{ij}(\alpha)^{g_{ij}}(1 - p^{\beta}_{ij}(\alpha))^{1-g_{ij}} \tag{4.2}$$

with the edge-success probability

$$(\forall i, j \in \mathcal{I}) : p^{\beta}_{ij}(\alpha) = \frac{\lambda(\hat{w}^{i,\beta}_{j}(\alpha) + \hat{w}^{j,\beta}_{i}(\alpha))}{\lambda(\hat{w}^{i,\beta}_{j}(\alpha) + \hat{w}^{j,\beta}_{i}(\alpha)) + \xi \hat{v}^{\beta}_{ij}(\alpha)}. \tag{4.3}$$

Proof. The Markov process $(\tilde{G}^{\beta}(t))_{t\geq 0}$ is irreducible for $\beta > 0$ and reversible by the symmetry assumption (D1). Solving the detailed balance conditions

$$\hat{\mu}_2^{\beta}(g|\alpha)\eta_{2,3}^{\beta}(g \to g \oplus (i,j)) = \hat{\mu}_{2,3}^{\beta}(g \oplus (i,j)|\alpha)\eta_{2,3}^{\beta}(g \oplus (i,j) \to g)$$

for all $g \in \mathcal{G}[\mathcal{I}]$ gives us

$$\hat{\mu}_{2,3}^{\beta}(g|\alpha) = \frac{1}{Z_{2,3}^{\beta}(\alpha)} \prod_{i=1}^{N} \prod_{j>i} \left(\frac{\lambda}{\xi} \frac{\hat{w}_j^{i,\beta}(\alpha) + \hat{w}_i^{j,\beta}(\alpha)}{\hat{v}_{ij}^{\beta}(\alpha)} \right)^{g_{ij}}.$$

Let $\bar{w}_{ij}^{\beta}(\alpha) := \hat{w}_j^{i,\beta}(\alpha) + \hat{w}_i^{j,\beta}(\alpha)$ and define $\theta_{ij}^{\beta}(\alpha) := \log\left(\frac{\lambda}{\xi}\frac{\bar{w}_{ij}^{\beta}(\alpha)}{\hat{v}_{ij}^{\beta}(\alpha)}\right)$. Further, define the Hamiltonian $H^{\beta}(g|\alpha) := \sum_{i,j>i} \theta_{ij}^{\beta}(\alpha) g_{ij}$, so that $\hat{\mu}_{2,3}^{\beta}(g|\alpha) = \frac{\exp(H^{\beta}(g|\alpha))}{\sum_{g' \in \mathcal{G}[\mathcal{I}]} \exp(H^{\beta}(g'|\alpha))}$. The constant of normalization can then be written as

$$Z_{2,3}^{\beta}(\alpha) = \sum_{g' \in \mathcal{G}[\mathcal{I}]} \exp(H^{\beta}(g'|\alpha)) = \sum_{i,j>i} \sum_{g_{ij}=0}^{1} \left(\prod_{i,j>i} \exp(\theta_{ij}^{\beta}(\alpha) g_{ij}) \right) = \prod_{i,j>i} (1 + \exp(\theta_{ij}^{\beta})).$$

The probability that edge (i,j) is active in the long run is

$$\begin{aligned} p_{ij}^{\beta}(\alpha) &= \sum_{g' \in \mathcal{G}[\mathcal{I}]: g'_{ij}=1} \hat{\mu}_{2,3}^{\beta}(g'|\alpha) = \sum_{g' \in \mathcal{G}[\mathcal{I}]} g'_{ij} \hat{\mu}_2^{\beta}(g'|\alpha) = \frac{\partial \log Z_{2,3}^{\beta}(\alpha)}{\partial \theta_{ij}^{\beta}(\alpha)} = \frac{\exp(\theta_{ij}^{\beta}(\alpha))}{1 + \exp(\theta_{ij}^{\beta}(\alpha))} \\ &= \frac{\lambda \bar{w}_{ij}^{\beta}(\alpha)}{\lambda \bar{w}_{ij}^{\beta}(\alpha) + \xi \hat{v}_{ij}^{\beta}(\alpha)}. \end{aligned}$$

Collecting terms and doing some simple manipulations gives the desired result. □

This strong result gives a full characterization of the induced ensemble of random graphs for volatility models such as Marsili et al. (2004), Ehrhardt et al. (2008a;b). It further establishes an interesting and surprising connection with the inhomogeneous random graph models proposed by Söderberg (2002), Park and Newman (2004) or Bollobás et al. (2007).

Any co-evolutionary model with noise, satisfying the set of assumptions (A1)-(A2), (L1)-(L5) and (D1-D4) will generate an inhomogeneous random graph, with edge-success probabilities (4.3).

Corollary 4.1. *The unique invariant distribution of the discrete-time random graph process* $(G^{\beta}_n)_{n\geq 0}$ *is*

$$\frac{\hat{q}(g)\nu^{\beta}(g|\alpha)}{\sum_{g'\in\mathcal{G}[\mathcal{I}]}\hat{q}(g')\nu^{\beta}(g'|\alpha)}$$

with

$$\nu^{\beta}(g|\alpha) := \prod_{i=1}^{N}\prod_{j>i}\left(\frac{\lambda}{\xi}\frac{\hat{w}^{i,\beta}_j(\alpha)+\hat{w}^{j,\beta}_i(\alpha)}{\hat{v}^{\beta}_{ij}(\alpha)}\right)^{g_{ij}}.$$

Proof. This follows form Lemma 4.1. □

5 Conclusion

This paper presented a general framework for studying co-evolutionary models with noise. We gave a complete characterization of the invariant distribution of such a model, which is a joint probability distribution on the set of action profiles and the set of networks. By means of two examples, a volatility model akin to Ehrhardt et al. (2008b), Staudigl (2009b) and a classical model based on Jackson and Watts (2002), we have shown how the unified approach is useful to make a systematic investigation of co-evolutionary models. Beside presenting a unified formalism to perform the by now important equilibrium selection technique of stochastic stability, we have demonstrated that a co-evolutionary model with noise generates an inhomogeneous random graph ensemble for the long run interaction structure of the population. The main result in this direction provides a closed form solution for the probability measure of this graph ensemble, and presents the general form for edge-success probabilities. Based on this novel insight, there are many new questions arise.

First, the edge success probabilities depend only on the behavioral rules the agents are assumed to employ. It would be interesting to see what differences between networks of this ensemble arise by assuming different behavioral rules. For instance, do best-responding agents tempt to self-organize in more structured and efficient network topologies as imitative agents? What role plays the underlying noise structure of the model (meant here as the interplay between behavioral noise β and overall network volatility ξ)? Second, the literature on social and epidemic diffusion (see e.g. Morris, 2000, Alós-Ferrer and Weidenholzer, 2008, Pastor-Satorras and Vespignani, 2001) have emphasized the importance of the network architecture in order to understand the phenomenon of contagion. In particular, notions of network clustering and cohesiveness have turned out to be important. We do not yet know the statistics produced by a co-evolutionary model. Third, in the context of volatility models Ehrhardt et al. (2008a) find three interesting dynamic effects; Resilience, Equilibrium co-existence and phase transitions (i.e. a discontinuous switch in the connectivity of the network by a slight change of the parameters affecting the edge success probability). Under what parameter configurations are these phenomena reproduced in the framework of a co-evolutionary model? The recent work by Bollobás et al. (2007) studies inhomogeneous random graphs and detects also a phase transition in network connectivity by exploring the size of components with a branching process approximation.

A Proof of Lemma 3.2

We first show that if $\omega \in \Re$, then there is no positive probability path under $\mathbf{K}$ that leads out of this set. Under ω every connected pair of players is coordinated. Let i be a current a_1 player. Every player in the component to which i belongs must then also play a_1.[23] Hence, every graph corresponding to $\omega \in \Omega^*$ must consist of finitely many components, each characterized by behavioral confor-

[23]If j would be a player in the component who plays a_2 he cannot be linked with a player who is path connected with i.

mity. By definition, applying **K** to such a state will not lead to a state outside Ω^*.

Now consider a state $\omega \notin \Re$. To show that such a state is transient under **K**, we have to find a positive probability path under **K** that leads to some state $\omega' \in \Re$, but no path from $\Re$ can be constructed that goes to ω. It is easy to see that once the unperturbed process is in $\Re$ there is no positive probability path that leads the process out of it, so one direction of the proof is already shown. For the other direction, the following algorithm constructs an (ω, ω')-path in finitely many steps;

Let $\omega_0 = \omega$ be our initial state. The set of uncoordinated edges $\mathcal{E}(\mathcal{I}_1(\omega_0), \mathcal{I}_2(\omega_0)) \neq \emptyset$, by hypothesis.[24] Let $t = 0, 1, 2, \ldots, m$ measure the number of iterations of the algorithm. Start from $t = 0$. The algorithm generates a sequence $\{\omega_t\}_{t=0}^m$, where the transition from ω_t to ω_{t+1} proceeds as follows:

Step 1: Pick the first edge from this set. Let one of the two involved players receive an action adjustment opportunity where he switches only to an action that gives him a strictly larger payoff compared to ω_t.[25] If this player changes his action, delete the edge from the list of uncoordinated edges, and call the resulting state ω_{t+1}. Then repeat Step 1. If the player does not change his action, go to Step 2.

Step 2: Give the other player an action adjustment opportunity as in Step 1. If he changes his action, delete the edge from the list of uncoordinated edges and call the resulting state ω_{t+1}. Then repeat Step 1. If the player does not change his action, go to Step 3.

Step 3: Delete the edge by a targeted link destruction event.[26] Call the resulting state ω_{t+1} and note that the set of uncoordinated edges decreased by 1. Go to Step 4.[27]

---

[24]This is the set of links that connect players from $\mathcal{I}_1(\omega)$ to players in $\mathcal{I}_2(\omega)$. $\mathcal{I}_\sigma(\omega)$ is the set of a_σ-players at ω.

[25]In 2×2 games with finite populations this choice rule is generically equivalent to demanding that a play switches to a best-response.

[26]Note that this is always a zero-cost step.

[27]An intermediate stage could be added to the algorithm, where we apply **K** to ω_{t+1} by letting the involved players create a link. This will lead to the creation of maximally

Step 4: Order the edges in $\mathcal{E}(\mathcal{I}_1(\omega_{t+1}), \mathcal{I}_2(\omega_{t+1}))$ in some way. If this set is empty exit the algorithm. Otherwise, go to Step 1.

B The Markov chain tree theorem and set-valued cost functions

To prove (2.8) we will make use of some general results from the theory of Markov chains and simulated annealing. Norris (1997) and Grimmett and Stirzaker (2001) are good references for the theory of finite Markov chains, and Catoni (1999; 2001) collects the relevant results from simulated annealing. Let $\omega \in \Omega, x, y, z \in \Omega \setminus \{\omega\}$ and $\mathcal{X} \subset \Omega$ a nonempty set. Denote by $\mathbf{K}^{\beta,n}$ the n-fold Matrix product of $\mathbf{K}^{\beta}$. The interpretation of this operation is that $K^{\beta,n}(x,y) = \mathbb{P}(X_n^{\beta} = y | X_0^{\beta} = x)$. Let $\omega \in \Omega$ be an arbitrary fixed state and define its *first passage time* as the random variable

$$\tau(\omega) := \inf\{n \geq 1 | X_n^{\beta} = \omega\}.$$

Since ω is recurrent we have $\mathbb{P}(\tau(\omega) < \infty | X_0^{\beta} = z) = 1$ for all z. Hence, the process returns to state ω almost surely, independent from where it takes off. Suppose we start the process from y and want to keep track of the number of times the chain visits x before it returns to ω. Phrased in probabilistic terms this amounts to calculate

$$\mathbb{E}\left[\sum_{n=0}^{\infty} \mathbb{1}_{\{X_n^{\beta}=x\} \cap \{\tau(\omega) \geq n\}} | X_0^{\beta} = y\right]. \tag{B.1}$$

The graph description of finite Markov chains is useful to calculate this seemingly complicated expression. Recall that a $\mathcal{X}$-revision graph is an element of the set of graphs $\mathcal{T}(\mathcal{X})$, connecting every point in $\Omega \setminus \mathcal{X}$ to a point in $\mathcal{X}$, without loops. If $\mathcal{X} = \{\omega, x\}$ then $\mathcal{T}(\{\omega, x\})$ contains all graphs which connect points from $\Omega \setminus \{\omega, x\}$ in a unique way either to ω or x. Denote by $\mathcal{T}_{y,x}(\mathcal{X})$ the set of $\mathcal{X}$-graphs which contain a path $\{\omega_1, \ldots, \omega_l\}$ such that $\omega_1 = y, \omega_l = x$ and $\omega_t \notin \mathcal{X}$, for all $2 \leq t \leq l-1$. If $y = x$ define $\mathcal{T}_{x,x}(\mathcal{X}) = \mathcal{T}(\mathcal{X})$. If $y \in \mathcal{X}$ set $\mathcal{T}_{y,x}(\mathcal{X}) = \emptyset$. It is intuitive that (B.1) should be proportional to the probability

2 coordinated links.

of graphs $T \in \mathcal{T}_{y,x}(\{\omega, x\})$. However, we also require to return to ω, so not all possible paths are allowed. We have to condition on the ω-trees, since these are the paths that lead in a unique way to ω. This heuristic argument suggests that (B.1) can be calculated as

$$\frac{\sum_{T\in\mathcal{T}_{y,x}(\{\omega,x\})} K^\beta(T)}{\sum_{T_\omega\in\mathcal{T}_\omega} K^\beta(T_\omega)} = \frac{1}{\rho^\beta(\omega)} \sum_{T\in\mathcal{T}_{y,x}(\{\omega,x\})} K^\beta(T). \tag{B.2}$$

Lemma 3.1 of Catoni (1999) gives a rigorous proof of this heuristic.[28]

Lemma B.1 (Lemma 3.1, Catoni (1999))**.** *Let $\bar{\mathbf{K}}^\beta$ denote the matrix $\mathbf{K}^\beta$ restricted to the set $\Omega \setminus \{\omega\}$. Then*

$$\begin{aligned}\sum_{n=0}^{\infty} \bar{K}^{\beta,n}(y,x) &= \mathbb{E}\left[\sum_{n=0}^{\infty} \mathbb{1}_{\{X_n^\beta = x\}\cap\{\tau(\omega)>n\}} | X_0^\beta = y\right] \\ &= \frac{\sum_{T\in\mathcal{T}_{y,x}(\{\omega,x\})} K^\beta(T)}{\rho^\beta(\omega)}.\end{aligned}$$

Before proving this result, we need the following simple observation.

Lemma B.2. *For all $y, x \neq \omega$, we have $\lim_{n\to\infty} \bar{K}^{\beta,n}(y,x) = 0$.*

Proof.

$$\lim_{n\to\infty} \bar{K}^{\beta,n}(y,x) \leq \mathbb{P}(\tau(\omega) = \infty | X_0^\beta = y) = 0$$

since ω is a recurrent state. □

As a consequence we see that $(\mathbf{Id} - \bar{\mathbf{K}}^\beta)$ is invertible. This is interesting, because for all $y, x \neq \omega$

$$\begin{aligned}(\mathbf{Id} - \bar{\mathbf{K}}^\beta)^{-1}(y,x) &= \sum_{n=0}^{\infty} \bar{K}^{\beta,n}(y,x) \\ &= \sum_{n=0}^{\infty} \mathbb{E}\left(\mathbb{1}_{\{X_n^\beta = x\}\cap\{\tau(\omega)>n\}} | X_0^\beta = y\right).\end{aligned}$$

Hence, this gives us the expected number of times the process visits x before hitting ω, which is the quantity we want to compute in Lemma B.1.

[28]The proof, which is taken from Catoni (1999), extends literally to the case where the singleton is replaced by a non-empty subset $\mathcal{X}$.

Proof of Lemma B.1. For all $y, x \neq \omega$, let us define

$$M(y,x) := \frac{1}{\rho^{\beta}(\omega)} \sum_{T \in \mathcal{T}_{y,x}(\{\omega,x\})} K^{\beta}(T).$$

Define the Kronecker-delta function as $\delta_{y,x} = 1$ if $y = x$ and 0 otherwise. We have to show that for all $y, x \neq \omega$

$$\sum_{z \neq \omega} (\delta_{y,z} - \bar{K}^{\beta}(y,z)) M(z,x) = \delta_{y,x}.$$

This can be written as

$$\sum_{z \neq y} K^{\beta}(y,z) M(y,x) = \delta_{y,x} + \sum_{z \in \Omega \setminus \{\omega,y\}} K^{\beta}(y,z) M(z,x)$$

$$\Leftrightarrow \sum_{z \neq y} K^{\beta}(y,z) \sum_{T \in \mathcal{T}_{y,x}(\{\omega,x\})} K^{\beta}(T) = \delta_{y,x} \rho^{\beta}(\omega) + \sum_{z \in \Omega \setminus \{\omega,y\}} K^{\beta}(y,z) \sum_{T \in \mathcal{T}_{z,x}(\{\omega,x\})} K^{\beta}(T)$$

Define the sets $\mathcal{C}_1 := \{(z,T) | z \neq y,\ T \in \mathcal{T}_{y,x}(\{\omega,x\})\}$ and $\mathcal{C}_2 := \{(z,T) | z \in \Omega \setminus \{\omega,y\},\ T \in \mathcal{T}_{z,x}(\{\omega,x\})\}$, so that we can equivalently write

$$\sum_{(z,T) \in \mathcal{C}_1} K^{\beta}(y,z) K^{\beta}(T) = \delta_{y,x} \rho^{\beta}(\omega) + \sum_{(z,T) \in \mathcal{C}_2} K^{\beta}(y,z) K^{\beta}(T). \tag{B.3}$$

Let us consider the case $y = x$ first, so that $\mathcal{C}_1$ is defined by the revision graphs in $\mathcal{T}(\{\omega,x\})$. Then $\mathcal{C}_2 \subset \mathcal{C}_1$, since every $\{\omega,x\}$-revision tree that contains an (z,x)-path is a $\{\omega,x\}$-revision graph. The converse, of course, need not apply. Define the map

$$\varphi : \mathcal{C}_1 \setminus \mathcal{C}_2 \to \mathcal{T}_{\omega}, (z,T) \mapsto \varphi(z,T) = (\Omega, \vec{E}(T) \cup \{(x,z)\}).$$

This operation takes an $\{\omega,x\}$-revision tree, not containing a (z,x)-path, and adds the edge (x,z). Thus, from the point z we have to arrive at ω in a unique way. By adding the edge (x,z) we create an ω-revision tree. If we can show that φ is bijective, then we can move between $\mathcal{C}_1 \setminus \mathcal{C}_2$ and $\mathcal{T}_{\omega}$ without losing any information. For $T' = \varphi(z,T) \in \mathcal{T}_{\omega}$, the inverse mapping is

$$\varphi^{-1}(T') = (\varphi_1^{-1}(T'), \varphi_2^{-1}(T')) = (\mathcal{R}_{T'}(x), \vec{E}(T') \setminus \{(x, \mathcal{R}_{T'}(x))\}) = (z, \vec{E}(T') \setminus \{(x,z)\}).$$

The left-hand side of eq. (B.3) turns then to[29]

$$
\begin{aligned}
\sum_{(z,T)\in\mathcal{C}_1} K^\beta(x,z)K^\beta(T) &= \sum_{(z,T)\in\mathcal{C}_1\setminus\mathcal{C}_2} K^\beta(x,z)K^\beta(T) + \sum_{(z,T)\in\mathcal{C}_2} K^\beta(x,z)K^\beta(T) \\
&= \sum_{T'\in\mathcal{T}_\omega} K^\beta(x,\varphi_1^{-1}(T'))K^\beta(\varphi_2^{-1}(T')) + \sum_{(z,T)\in\mathcal{C}_2} K^\beta(x,z)K^\beta(T) \\
&= \sum_{T'\in\mathcal{T}_\omega} K^\beta(x,\mathcal{R}_{T'}(x))\frac{K^\beta(T')}{K^\beta(x,\mathcal{R}_{T'}(x))} + \sum_{(z,T)\in\mathcal{C}_2} K^\beta(x,z)K^\beta(T) \\
&= \rho^\beta(\omega) + \sum_{(z,T)\in\mathcal{C}_2} K^\beta(x,z)K^\beta(T)
\end{aligned}
$$

what coincides with the right-hand side of this equation.
Now, consider the case $y \neq x$. Define the map $\varphi : \mathcal{C}_1 \to \mathcal{C}_2$ by

$$
\varphi(z,T) = \begin{cases} (z,T) & ,\text{if } T \in \mathcal{T}_{z,x}(\{\omega,x\}) \\ (\mathcal{R}_T(y), (\vec{E}(T) \cup \{(y,z)\}) \setminus \{(y,\mathcal{R}_T(y))\}) & ,\text{if } T \notin \mathcal{T}_{z,x}(\{\omega,x\}). \end{cases}
$$

φ maps the pair (z,T) onto itself if T contains an (z,x)-path. If such a path does not exist, then it connects y with z, deletes the (unique) outgoing edge from y, and shifts the initial vertex of the path from y to its unique neighbor under T, $\mathcal{R}_T(y)$. Since there exists a path connecting y with x, the (unique) neighbor of y on T is also connected with x. Hence, we have constructed a revision tree $T' \in \mathcal{T}_{\mathcal{R}_T(y),x}(\{\omega,x\})$, with $\mathcal{R}_T(y) \in \Omega \setminus \{\omega,y\}$.[30] If we can show that φ is bijective, then $\mathcal{C}_1 = \mathcal{C}_2$ follows and we are done. We claim

$$
\varphi^{-1}(z,T) = \begin{cases} (z,T) & \text{if } T \in \mathcal{T}_{y,x}(\{\omega,x\}), \\ (\mathcal{R}_T(y), (\vec{E}(T) \cup \{(y,z)\}) \setminus \{(y,\mathcal{R}_T(y))\}) & \text{if } T \notin \mathcal{T}_{y,x}(\{\omega,x\}). \end{cases}
$$

Then $\varphi^{-1}(\varphi(z,T)) = (z,T)$ for all $(z,T) \in \mathcal{C}_1$. To see this, start with $(z,T) \in \mathcal{T}_{z,x}(\{\omega,x\})$. Then $\varphi(z,T) = (z,T) \in \mathcal{C}_2$ and $T \in \mathcal{T}_{y,x}(\{\omega,x\})$, hence $\varphi^{-1}(\varphi(z,T)) = (z,T)$. In the case where $T \notin \mathcal{T}_{z,x}(\{\omega,x\})$, let us call $\varphi(z,T) = (z',T') \in \mathcal{C}_2$. Then $T' \notin \mathcal{T}_{y,x}(\{\omega,x\})$, and consequently

$$
\begin{aligned}
\varphi(z',T') &= (\mathcal{R}_{T'}(y), (\vec{E}(T') \cup \{(y,z')\}) \setminus \{(y,\mathcal{R}_{T'}(y))\}) \\
&= (z, (\vec{E}(T') \cup \{(y,\mathcal{R}_T(y))\}) \setminus \{(y,z)\}) \\
&= (z,T).
\end{aligned}
$$

[29]Define $0 \cdot \infty = 0$.

[30]If $\mathcal{R}_T(y) = x$ then we get the pair (x,T) with $T \in \mathcal{T}(\{\omega,x\})$ which lies in $\mathcal{C}_2$ for $z = x$.

□

The expected time spent in some state x before the system returns to ω is given by

$$v_x(\omega) = \mathbb{E}\left(\sum_{n=0}^{\infty} \mathbb{1}_{\{X_n^\beta = x\}\cap\{\tau(\omega)\geq n+1\}} | X_0^\beta = \omega\right). \tag{B.4}$$

Intuitively, this is the average length of ω-cycles on which x is visited.

Lemma B.3. *Let $v(\omega)$ denote the vector whose elements are defined by (B.4). Then*

(i) $v_\omega(\omega) = 1$;

(ii) $v(\omega)\mathbf{K}^\beta = v(\omega)$;

(iii) $v(\omega)$ *is bounded and positive.*

Proof. (i) By definition.

(ii) By the Markov property and time-homogeneity,

$$\begin{aligned}
v_x(\omega) &= \sum_{n=1}^{\infty} \mathbb{P}(X_n^\beta = x, \tau(\omega) \geq n | X_0^\beta = \omega) \\
&= \sum_{n=1}^{\infty} \sum_{\omega' \in \Omega} \mathbb{P}(X_{n-1}^\beta = \omega', X_n^\beta = x, \tau(\omega) \geq n | X_0^\beta = \omega) \\
&= \sum_{n=1}^{\infty} \sum_{\omega' \in \Omega} \mathbb{P}(X_{n-1}^\beta = \omega', \tau(\omega) \geq n | X_0^\beta = \omega) K^\beta(\omega', x) \\
&= \sum_{n=0}^{\infty} \sum_{\omega' \in \Omega} \mathbb{P}(X_n^\beta = \omega', \tau(\omega) - 1 \geq n | X_0^\beta = \omega) K^\beta(\omega', x) \\
&= \sum_{\omega' \in \Omega} v_{\omega'}(\omega) K^\beta(\omega', x)
\end{aligned}$$

(iii) Suppose there exists a state x such that $v_x(\omega) = 0$. Then, for all $n \geq 1$,

$$0 = \sum_{\omega' \in \Omega} v_{\omega'}(\omega) K^{\beta,n}(\omega', x) = K^{\beta,n}(\omega, x) + \sum_{y \neq \omega} v_y(\omega) K^{\beta,n}(y, x)$$

and so $K^{\beta,n}(\omega, x) = 0$, contradicting irreducibility. Essentially the same argument can be used to see that $v_x(\omega) < \infty$ for all x.

□

The *expected return time* to ω is $\bar{v}(\omega) = \sum_{\omega' \in \Omega} v_{\omega'}(\omega)$. This is a measure of the average length of ω-cycles. A state ω is called *positive recurrent* if $\bar{v}(\omega) < \infty$.

Lemma B.4. *Let $\mathbf{K}^\beta$ be irreducible and recurrent. Then $\mathbf{K}^\beta$ has an invariant distribution μ^β such that $\mu^\beta(\{\omega\}) = \mu^\beta(\omega) = 1/\bar{v}(\omega)$.*

Proof. Since Ω is finite, there exists a positive recurrent state $\omega \in \Omega$. Now, since the process is irreducible, all states are positive recurrent. Then $\bar{v}(\omega) = \sum_{\omega' \in \Omega} v_{\omega'}(\omega) < \infty$. Since $v(\omega)$ defines an invariant measure for $\mathbf{K}^\beta$, $\mu^\beta = (1/\bar{v}(\omega))v(\omega)$ is an invariant distribution for $\mathbf{K}^\beta$, satisfying $\mu^\beta(\omega) = 1/\bar{v}(\omega)$. □

Using this Lemma, observe that

$$
\begin{aligned}
\mu^\beta(\omega) &= \left(\sum_{\omega' \in \Omega} v_{\omega'}(\omega) \right)^{-1} = \left(1 + \sum_{x \neq \omega} v_x(\omega) \right)^{-1} \\
&= \left[1 + \sum_{n=1}^{\infty} \mathbb{P}(X_n^\beta \neq \omega, \tau(\omega) \geq n+1 | X_0^\beta = \omega) \right]^{-1} \\
&= \left[1 + \sum_{n=1}^{\infty} \sum_{y \neq \omega} \mathbb{P}(X_1^\beta = y, \tau(\omega) \geq n+1 | X_0^\beta = \omega) \right]^{-1} \\
&= \left[1 + \sum_{y \neq \omega} K^\beta(\omega, y) \sum_{n=1}^{\infty} \mathbb{P}(\tau(\omega) \geq n | X_0^\beta = y) \right]^{-1} \\
&= \left[1 + \sum_{y \neq \omega} K^\beta(\omega, y) \mathbb{E}(\tau(\omega) | X_0^\beta = y) \right]^{-1}.
\end{aligned}
$$

We have for all $y \neq \omega$ the identity

$$
\begin{aligned}
\mathbb{E}(\tau(\omega) | X_0^\beta = y) &= \mathbb{E}\left(\sum_{n=0}^{\infty} \mathbb{1}_{\{X_n^\beta \neq \omega\} \cap \{\tau(\omega) > n\}} | X_0^\beta = y \right) \\
&= \mathbb{E}\left(\sum_{x \neq \omega} \sum_{n=0}^{\infty} \mathbb{1}_{\{X_n^\beta = x\} \cap \{\tau(\omega) > n\}} | X_0^\beta = y \right) \\
&= \sum_{x \neq \omega} \sum_{n=0}^{\infty} \bar{K}^{\beta,n}(y, x).
\end{aligned}
$$

The last equality follows from Lemma B.1. Using this identity gives

$$\begin{aligned}
\mu^{\beta}(\omega) &= \left[1 + \frac{1}{\rho^{\beta}(\omega)} \sum_{y,x\neq\omega} K^{\beta}(\omega,y) \sum_{T\in\mathcal{T}_{y,x}(\{\omega,x\})} K^{\beta}(T)\right]^{-1}. \\
&= \frac{\rho^{\beta}(\omega)}{\rho^{\beta}(\omega) + \sum_{x\neq\omega}\sum_{y\neq\omega} K^{\beta}(\omega,y)\sum_{T\in\mathcal{T}_{y,x}(\{\omega,x\})} K^{\beta}(T)} \\
&= \frac{\rho^{\beta}(\omega)}{\rho^{\beta}(\omega) + \sum_{x\neq\omega}\sum_{T_x\in\mathcal{T}_x} K^{\beta}(T_x)}
\end{aligned}$$

which is eq. (2.8).

We now provide some justifications for the cost functions (2.14). The results presented here are due to Beggs (2005), who in turn builds on the work of Catoni (1999). The clue is to consider a modified Markov chain, which monitors only transitions in a suitably chosen subset $\mathcal{X} \subset \Omega$. Therefore, for $m \in \mathbb{N}_0$, define the stopping times of successive visitations of the set $\mathcal{X}$ as $\tau_{-1}(\mathcal{X}) \equiv 0$, $\tau_m(\mathcal{X}) := \inf\{n \geq \tau_{m-1}(\mathcal{X}) + 1 | X_n^{\beta} \in \mathcal{X}\}$. The Markov chain $Z_m^{\beta} := X_{\tau_m(\mathcal{X})}^{\beta}$ records all visitations of X^{β} to the set $\mathcal{X}$.

Lemma B.5. *Let $\mathcal{X} \subset \Omega$ be a non-empty set. $(Z_m^{\beta})_{m\geq 0}$ is an irreducible, recurrent and time-homogenous Markov chain on $\mathcal{X}$. Its unique invariant distribution is given by $\mu^{\beta}(\cdot|\mathcal{X})$ and its transition probabilities are for all $y, x \in \mathcal{X}$*

$$\mathbb{P}(Z_{m+1}^{\beta} = x | Z_m^{\beta} = y) = K^{\beta}(y,x) + \sum_{z\in\Omega\backslash\mathcal{X}} K^{\beta}(y,z) Q_{\Omega\backslash\mathcal{X}}^{\beta}(y \to x|z) \tag{B.5}$$

where

$$Q_{\Omega\backslash\mathcal{X}}^{\beta}(y \to x|z) := \frac{\sum_{T\in\mathcal{T}_{z,x}(\mathcal{X})} K^{\beta}(T)}{\sum_{T\in\mathcal{T}(\mathcal{X})} K^{\beta}(T)}.$$

is the expected number of visitations to z before the restricted process reaches x.

Proof. That the restricted process is a Markov chain with these properties can be proved quite easily. See Proposition 7.2.1 in Catoni (2001). For the second claim, note that the strong Markov property (see Norris, 1997), applied to the stopping times $\tau_m(\mathcal{X})$, implies that

$$\begin{aligned}
\mathbb{P}(Z^\beta_{m+1} = x | Z^\beta_m = y) &= \mathbb{P}(X^\beta_{\tau_{m+1}(\mathcal{X})} = x | X^\beta_{\tau_m(\mathcal{X})} = y) = \mathbb{P}(X^\beta_{\tau_1(\mathcal{X})} = x | X^\beta_0 = y) \\
&= K^\beta(y,x) + \sum_{n=2}^{\infty} \sum_{z \notin \mathcal{X}} \mathbb{P}(X^\beta_1 = z, X^\beta_s \notin \mathcal{X}\ \forall s \in (1, n-1], X^\beta_n = x | X^\beta_0 = y) \\
&= K^\beta(y,x) \\
&+ \sum_{n=1}^{\infty} \sum_{z,\omega \notin \mathcal{X}} K^\beta(y,z) \mathbb{P}(X^\beta_n = \omega, \tau(\mathcal{X}) \geq n | X^\beta_0 = z) K^\beta(\omega, x) \\
&= K^\beta(y,x) + \sum_{z,\omega \notin \mathcal{X}} K^\beta(y,z) \frac{\sum_{T \in \mathcal{T}_{z,\omega}(\mathcal{X} \cup \{\omega\})} K^\beta(T)}{\sum_{T \in \mathcal{T}(\mathcal{X})} K^\beta(T)} K^\beta(\omega, x) \\
&= K^\beta(y,x) + \sum_{z \notin \mathcal{X}} K^\beta(y,z) \frac{\sum_{T \in \mathcal{T}_{z,x}(\mathcal{X})} K^\beta(T)}{\sum_{T \in \mathcal{T}(\mathcal{X})} K^\beta(T)} \\
&= K^\beta(y,x) + \sum_{z \in \Omega \setminus \mathcal{X}} K^\beta(y,z) Q^\beta_{\Omega \setminus \mathcal{X}}(y \to x | z)
\end{aligned}$$

where we have used in the fourth line an extended version of Lemma B.1. □

We will apply this result to derive the set-valued cost functions (2.14). Let $\mathcal{L}_1, \ldots, \mathcal{L}_k$ denote the recurrent classes of the unperturbed model $\mathcal{M}$ and $\Re = \bigcup_{i=1}^k \mathcal{L}_i$ the union of recurrent classes. The literature often refers to the sets $\mathcal{L}_i$ as limit sets. From each limit set we make an arbitrary selection $x_i \in \mathcal{L}_i, 1 \leq i \leq k$, and define $\mathcal{X} := \{x_1, \ldots, x_k\}$. Note that $\mathcal{X}$ contains the absorbing states (i.e. the singleton recurrent sets). For $y, x \in \mathcal{X}$, let

$$\hat{C}^{\mathcal{X}}(y,x) := -\lim_{\beta \to 0} \beta \log \mathbb{P}(Z^\beta_{m+1} = x | Z^\beta_m = y)$$

be the cost function of the restricted process $(Z^\beta_m)_{m \geq 0}$. Further, define $\hat{c}^*(\omega) := \min_{y \in \Omega \setminus \{\omega\}} \hat{c}(\omega, y)$ the least cost transition from some state $\omega \in \Omega$ (omitting the type of transition).

Lemma B.6. *Let $\mathcal{X} = \{x_1, \ldots, x_k\}, x_i \in \mathcal{L}_i, 1 \leq i \leq k$. Then, for all $y, x \in \mathcal{X}$, the costs of transiting from y to x are given by*

$$c^{\mathcal{X}}(y,x) = \min_{P \in \mathcal{P}_{y,x}(\bar{\mathcal{X}} \cup \{y\}, \mathcal{X})} \hat{C}(P), \tag{B.6}$$

where for any path P, $\hat{C}(P) := \sum_{\vec{e} \in \vec{E}(P)} \hat{c}(\vec{e})$, and $\bar{\mathcal{X}} = \Omega \setminus \mathcal{X}$.

Proof. The proof is based on Lemma 2.1 and the transition probability of the restricted process (Z_m^β) found in Lemma B.5. We know that $K^\beta(y,x) = \exp\left[-\frac{1}{\beta}(\hat{c}(y,x) + o(1))\right]$. For a given point $z \in \Omega \setminus \mathcal{X}$, we have to find an asymptotic bound for $K^\beta(y,z)Q^\beta_{\Omega\setminus\mathcal{X}}(y \to x|z)$. For sufficiently small β this probability can be written as

$$\exp\left[-\frac{1}{\beta}(\hat{c}(y,z) + o(1))\right] \frac{\sum_{T\in\mathcal{T}_{z,x}(\mathcal{X})} \exp\left[-\frac{1}{\beta}(\hat{C}(T) + o(1))\right]}{\sum_{T\in\mathcal{T}(\mathcal{X})} \exp\left[-\frac{1}{\beta}(\hat{C}(T) + o(1))\right]}.$$

Taking logarithms, and multiplying by $-\beta$ gives us

$$\hat{c}(y,z) - \beta \log[Q^\beta_{\Omega\setminus\mathcal{X}}(y \to x|z)]. \tag{B.7}$$

The second terms is governed by

$$\log\left[\sum_{T\in\mathcal{T}_{z,x}(\mathcal{X})} \exp\left(-\frac{1}{\beta}(\hat{C}(T) + o(1))\right)\right] - \log\left[\sum_{T\in\mathcal{T}(\mathcal{X})} \exp\left(-\frac{1}{\beta}(\hat{C}(T) + o(1))\right)\right].$$

Lemma 2.1 tells us that for $\beta \downarrow 0$ this number is asymptotically equivalent to

$$\max_{T\in\mathcal{T}_{z,x}(\mathcal{X})} \exp(-\hat{C}(T)/\beta) - \max_{T\in\mathcal{T}(\mathcal{X})} \exp(-\hat{C}(T)/\beta).$$

(B.7) boils then down to

$$\hat{c}(y,z) + \min_{T\in\mathcal{T}_{z,x}(\mathcal{X})} \hat{C}(T) - \min_{T\in\mathcal{T}(\mathcal{X})} \hat{C}(T). \tag{B.8}$$

Call $T^*_{z,x} \in \mathcal{T}_{z,x}(\mathcal{X})$ a least cost $\mathcal{X}$-graph containing an (z,x)-path, and $T^*_{\mathcal{X}} \in \mathcal{T}(\mathcal{X})$ a least cost $\mathcal{X}$-revision graph. Call P^* the (z,x) path used on $T^*_{z,x}$. We claim that all edges in $T^*_{z,x}$, which are not on the path P^*, are also used under $T^*_{\mathcal{X}}$. This follows from the fact that $\mathcal{T}_{z,x}(\mathcal{X}) \subset \mathcal{T}(\mathcal{X})$. The only difference between the graphs $T^*_{z,x}$ and $T^*_{\mathcal{X}}$ are the edges on the path $P^* = \{\omega_1, \ldots, \omega_l\}, \omega_1 = z, \omega_l = x, \omega_t \notin \mathcal{X}, \forall t = 2, \ldots, l-1$. The edge (ω_{t-1}, ω_t) need not be globally optimal, so that this edge causes supplementary costs $\hat{c}(\omega_{t-1}, \omega_t) - \hat{c}^*(\omega_{t-1})$. The term $\hat{c}^*(\omega_{t-1})$ is the cost of the edge leaving ω_{t-1} under $T^*_{\mathcal{X}}$. Hence, for any $z \notin \mathcal{X}$ we can pin down the costs of a transition from y to x, via z, as

$$\hat{c}(y,z) + \min_P \left\{ \sum_{t=2}^{l} [\hat{c}(\omega_{t-1}, \omega_t) - \hat{c}^*(\omega_{t-1})] : P = \{\omega_1, \ldots, \omega_l\} \in \mathcal{P}_{z,x}(\Omega \setminus \mathcal{X}, \mathcal{X}) \right\}.$$

Call this $\hat{C}^{\mathcal{X}}(y \to x|z)$. It follows that

$$c^{\mathcal{X}}(y,x) = \min\left\{\hat{c}(y,x), \min_{z\in\Omega\setminus\mathcal{X}} \hat{C}^{\mathcal{X}}(y \to x|z)\right\}.$$

Next, we claim that if ω is used on the optimal path P^*, then $\hat{c}^*(\omega) = 0$. To see this, observe that by definition of such paths, ω is either a transient state, or it is a recurrent state, not contained in the selection $\mathcal{X}$. In the first case, $\hat{c}^*(\omega) = 0$, since any transient state can be appended to a zero-cost path leading to some recurrent state. In the second case we also have $\hat{c}^*(\omega) = 0$, since ω cannot be absorbing, hence communicates with another state in the same recurrent class. Hence, if ω_1 is the first state on the optimal path P^* then $\hat{c}^*(\omega_1) = 0$, and iteration gives $\hat{c}^*(\omega_{t-1}) = 0$ for all $2 \leq t \leq l-1$. Consequently

$$\begin{aligned} c^{\mathcal{X}}(y,x) &= \min\left\{\hat{c}(y,x), \min_{z\in\Omega\setminus\mathcal{X}} \min_{P\in\mathcal{P}_{z,x}(\Omega\setminus\mathcal{X},\mathcal{X})} \left(\hat{c}(y,z) + \hat{C}(P)\right)\right\} \\ &= \min_{P\in\mathcal{P}_{y,x}(\tilde{\mathcal{X}}\cup\{y\},\mathcal{X})} \hat{C}(P). \end{aligned}$$

□

This Lemma gives us the costs of a transition between two recurrent states y, x. If $y \in \mathcal{L}$ and $x \in \mathcal{L}'$, then we can extend the above argument to a setwise cost functions, measuring the difficulty of a transition from recurrent class $\mathcal{L}$ to recurrent class $\mathcal{L}'$. Let $\omega \in \mathcal{L}, \omega' \in \mathcal{L}'$. There is a zero-cost path connecting y with ω, and a zero-cost path connecting x with ω'. Hence, the least cost of moving from $\mathcal{L}$ to $\mathcal{L}'$ is exactly (2.14). This in turn shows that the least cost of reaching a state $\omega \in \mathcal{L}$ coincides with the minimal cost needed to reach the limit set $\mathcal{L}$ from all other limit sets, justfying eq. (2.15). Hence, if ω is stochastically stable, so must all states in the same recurrent class.

Corollary B.1.

$$\Omega^* = \bigcup\{\mathcal{L}|(\exists\omega \in \mathcal{L}) : \gamma(\omega) = \min_{\omega'\in\Omega} \gamma(\omega')\}. \tag{B.9}$$

One can also use the argument in Lemma B.6 to establish a connection with the radius/co-radius formulas of Ellison (2000). I refer to Beggs (2005) for further discussion.

References

Alós-Ferrer, C., Netzer, N., 2007. The logit-response dynamics, 2007, Forthcoming: Games and Economic Behavior.

Alós-Ferrer, C., Weidenholzer, S., 2008. Contagion and efficiency. Jounal of Economic Theory 143, Issue 1, 251–274.

Barabási, A., Albert, R., 1999. Emergence of scaling in random networks. Science 286, 509–512.

Beggs, A., 2005. Waiting times and equilibrium selection. Economic Theory 25, 599–628.

Bergin, J., Lipman, B., 1996. Evolution with state-dependent mutations. Econometrica 64, 943–956.

Binmore, K., Samuelson, L., 1999. Evolutionary drift and equilibrium selection. Review of Economic Studies 66, 363–393.

Blume, L., 1993. Statistical mechanics of strategic interaction. Games and Economic Behavior 5, 387–424.

Blume, L. E., 2003. How noise matters. Games and Economic Behavior 44, 251–271.

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

Catoni, O., 1999. Simulated annealing algorithms and Markov chains with rare transitions. Séminaire de Probabilités XXXIII. Lecture Notes in Mathematics, Vol. 1079.

Catoni, O., 2001. Statistical Learning Theory and Stochastic Optimization. Lecture Notes in Mathematics, Vol. 1851.

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

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

Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2008b. Networks emerging in a volatile world. EUI Working papers.

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.

Erdös, P., Rényi, A., 1960. On the evolution of random graphs. Publications of the Mathematical Institute of the Hungarian Academy of Sciences 5, 17–61.

Freidlin, M., Wentzell, A., 1998. Random Perturbations of Dynamical Systems, 2nd Ed. Springer Verlag, New York.

Goyal, S., Vega-Redondo, F., 2005. Network formation and social coordination. Games and Economic Behavior 50, 178–207.

Grimmett, G., Stirzaker, D., 2001. Probability and Random Processes. Oxford University Press, USA; 3 edition.

Hojman, D. A., Szeidl, A., 2006. Endogenous networks, social games, and evolution. Games and Economic Behavior 55, Issue 1, 112–130.

Jackson, M., Watts, A., 2002. On the formation of interaction networks in social coordination games. Games and Economic Behavior 41, 265–291.

Jackson, M., Wolinsky, A., 1996. A strategic model of social and economic networks. Journal of Economic Theory 71, 44–74.

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

Marsili, M., Vega-Redondo, F., Slanina, F., 2004. The rise and fall of the networked society: A formal model. PNAS 101, 1439–1442.

Morris, S., 2000. Contagion. Review of Economic Studies 67, 57–78.

Newman, M., 2003. The structure and function of complex networks. SIAM Review 45, no.2, 167–256.

Norris, J., 1997. Markov Chains. Cambridge Series in Statistical and Probabilistic Mathematics.

Park, J., Newman, M. E. J., 2004. The statistical mechanics of networks. Physical Review E 70, 066117–1–15.

Pastor-Satorras, R., Vespignani, A., 2001. Epidemic spreading in scale-free networks. Physical Review E 86, 14.

Samuelson, L., 1997. Evolutionary Games and Equilibrium Selection. MIT Press, Cambridge Massachusetts.

Sandholm, W., 2009. 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., 2009a. Co-evolutionary dynamics and bayesian interaction games, mimeo, University of Vienna.

Staudigl, M., 2009b. Potential games in volatile environments, mimeo, University of Vienna.

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

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