# A Markov state modeling approach to characterizing the punctuated equilibrium dynamics of stochastic evolutionary games

M. Hallier*, C. Hartmann†

February 15, 2016

**Abstract**

Stochastic evolutionary games often share a dynamic property called punctuated equilibrium; this means that their sample paths exhibit long periods of stasis near one population state which are infrequently interrupted by switching events after which the sample paths stay close to a different population state, again for a long period of time. This has been described in the literature as a favorable property of stochastic evolutionary games. The methods used so far in stochastic evolutionary game theory, however, do not fully characterize these dynamics. We present an approach that aims at exposing the punctuated equilibrium dynamics by constructing Markov models on a reduced state space which approximate well this dynamic behavior. Besides having good approximation properties, the approach allows a simulation-based algorithm, which is appealing in the case of complex games.

*Free University Berlin, Berlin, Germany
†Free University Berlin, Berlin, Germany

# 1 Introduction

Stochastic evolutionary games often share the punctuated equilibrium property; this means that their sample paths exhibit long periods of stasis near one population state which are infrequently interrupted by switching events after which the sample paths stay close to a different population state, again for a long period of time. This punctuated equilibrium property has been characterized in the literature as a favorable property of stochastic evolutionary games (e.g., Young, 1998, 2006), not least because it might lead to a different perspective on modeling conventions and on the problem of equilibrium selection (Jaeger, 2008, 2012).

The methods used so far in stochastic evolutionary game theory, however, do not characterize the dynamics of the evolutionary games with respect to this property. More specifically, the analysis of evolutionary game models exclusively focuses on equilibrium selection, equilibria of mean dynamics, or the determination of stochastically stable states (Benaim and Weibull, 2003; Ellison, 2000; Foster and Young, 1990; Hofbauer and Sigmund, 2003; Kurtz, 1970; Sandholm, 2010, 2011; Weibull, 1995; Young, 1993, 1998, 2006). All of these approaches are not able to thorougly describe punctuated equilibrium dynamic behavior of the considered evolutionary processes.

This is in contrast to physics and chemistry where there has been much research in the last century on the mathematical description and analysis of this dynamic property[1] (for a short historical overview see, e.g., the introductory chapter of Bovier, 2009). We build on these existing approaches to present a novel approach to the analysis of stochastic evolutionary games. In particular, after introducing the necessary game theoretical notions (Section 2), we show how to construct Markov models of reduced complexity that approximate their essential dynamic behavior (Section 3). The basic idea behind these so-called *Markov state models* is to approximate the original Markov process by a Markov chain on a small finite state space. More specifically, a Markov state model is defined as a Markov chain whose state space consists of sets of population states near which the sample paths of the original Markov process reside for a long time and whose transition rates between these macrostates are given by the aggregate statistics of jumps between those sets of population states. An advantage of this approach in the context of complex models with large state spaces is that the transition probabilities between the macrostates can be estimated on the basis of simulated short-term trajectory data. Moreover, it has been shown that it is possible to construct Markov state models with good approximation properties if punctuated equilibrium dynamics characterize the system of interest. Thus, we can construct Markov state models that approximate the original stochastic evolutionary game on the long time scales; the approach therefore complements traditional approaches such as stochastic stability analysis, which studies infinite horizon behavior, or deterministic approximation results, which are valid only on short time intervals. In addition, the approach allows a simulation-based

[1] Called *metastability* in these fields.

algorithmic strategy to the construction of Markov state models for stochastic evolutionary games, which is of interest especially for complex games. One limitation of the approach, however, is that the results on the approximation quality of Markov state models depend on the original stochastic evolutionary game to be reversible. We discuss this limitation and give an outlook for further research in this direction (Section 4).

# 2 Definitions

A stochastic evolutionary game is defined by a *population game* and a *revision protocol*, which specifies the strategy updating process of the players in the population game.

## Population game

We consider only games played by a single, finite population, in which each player faces the same set of strategies. More specifically, let a *population game* consist of a population of $n$ players, a strategy set $S = \{1, \cdots, m\}$, and a a payoff function $F : \Delta_n^{m-1} \to \mathbb{R}^m$, where $\Delta_n^{m-1} = \{x \in \mathbb{R}^m_{\geq 0} : \sum_{j \in S} x_j = 1 \text{ and } n\, x \in \mathbb{Z}^m\}$ is the set of population states. The $j$-th component of $x \in \Delta_n^{m-1}$ represents the proportion of players choosing the strategy $j$ in the population and $F_i(x)$ represents the payoff to playing strategy $i$ when the population state is $x$.

## Revision Protocol

The basic idea of the strategy updating process of the players in a given population game is the following: at every moment in time, each agent has chosen a strategy in the strategy set $S$. At times $t = k\delta$, where $\delta = 1/n, k \in \mathbb{N}$, exactly one agent is randomly drawn (with equal probability for all players) to reconsider her strategy choice. We assume statistical independence between successive draws.

In this context, a *revision protocol*[2] formulates how players choose a strategy given a revision opportunity. It is a function $\rho : \Delta^{m-1} \times \mathbb{R}^m \to \mathbb{R}^{m \times m}_{\geq 0}$ with $\sum_{j=1}^m \rho_{ij}(x, \pi) = 1$ for each $i \in S$, all population states $x \in \Delta^{m-1}$, and all possible payoffs $\pi \in \mathbb{R}^m$. The revision protocol thus associates to each population state $x \in \Delta^{m-1}$ and payoffs $\pi \in \mathbb{R}^m$ a matrix of transition probabilities[3] $\rho(x, \pi) = (\rho_{ij}(x, \pi))_{i,j=1,\ldots,m}$ where $\rho_{ij}(x, \pi)$ represents the probability of the agent to switch from strategy $i$ to strategy $j$ given the current population state $x$ and payoffs $\pi$. For simplicity, we assume that all players display the same strategy-updating behavior, i.e., act according to the same revision protocol.

[2] The definition of a revision protocol given here differs from the one given, e.g., in Sandholm (2010) in that we consider time-discrete updating processes instead of time-continuous ones.

[3] The revision protocol is not to be confused with the transition matrix of the aggregate strategy updating process, see below.

## Aggregate Strategy Updating Process

Given a given population game with payoff function $F$ and revision protocol $\rho$, only one agent is drawn from the whole population (with uniform distribution) to reconsider its strategy choice. This means, on the aggregate level, that transitions between population states are only possible, i.e., have a probability greater 0, if they differ in at most one component by at most $1/n$. Moreover, the probability of drawing an agent that currently holds strategy $i \in S$ corresponds to the share $x_i$ of strategy $i$ in the current population state $x$, and the probability that an agent holding strategy $i$ changes to strategy $j$ when given the chance to reconsider the strategy choice is given by $\rho_{ij}(x, F(x))$. Assuming statistical independence, the strategy updating process on the population level is thus a time-discrete Markov chain $X = (X_t)_{t\in\mathbb{T}}$ on the set of population states $Z = \Delta_n^{m-1}$, where $\mathbb{T} = \{k\delta \mid k \in \mathbb{N}, \delta = 1/n\}$. Its transition matrix $P = (p_{xy})_{x,y\in\Delta_n^{m-1}}$ is given by

$$p_{xy} = \begin{cases} x_i\rho_{ij}(x, F(x)) & \text{if } y = x + \frac{1}{n}(e_j - e_i),\ i, j \in S, i \neq j, \\ 1 - \sum_{i\in S}\sum_{j\neq i} x_i\rho_{ij}(x, F(x)) & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

In the special case of games with only two strategies (i.e., $m = 2$), we can identify the state $x \in \Delta_n^1 \subset \mathbb{R}^2$ with $\chi = x_1$ (since $x_2 = 1 - x_1$). We can thus restrict our analysis of the chain $(X_t)_{t\in\mathbb{T}}$ to the state space $Z = \{0, \frac{1}{n}, \cdots, 1\}$ and we will write (with abuse of notation) $F(\chi)$ for $F(x)$ (i.e., $F : Z \to \mathbb{R}^2$) and $\rho(\chi, F(\chi))$ for $\rho(x, F(x))$ (i.e., $\rho : Z \times \mathbb{R}^2 \to \mathbb{R}^{2\times 2}_{\geq 0}$). Moreover, this implies that stochastic evolutionary games with two strategies are birth-and-death chains with transition matrix

$$P = \begin{pmatrix} 1-\alpha_0 & \alpha_0 & 0 & \cdots & & \\ \beta_1 & 1-(\beta_1+\alpha_1) & \alpha_1 & 0 & \cdots & \\ 0 & \beta_2 & 1-(\beta_2+\alpha_2) & \alpha_2 & 0 & \cdots \\ & & \vdots & & & \\ 0 & \cdots & & 0 & \beta_n & 1-\beta_n \end{pmatrix}, \quad (2.1)$$

where the parameters are

$$\alpha_j = \Big(1 - \frac{j}{n}\Big)\rho_{21}\Big(\frac{j}{n}, F\big(\frac{j}{n}\big)\Big) \quad \text{for } j = 0, \cdots, n-1 \quad (2.2)$$

$$\beta_j = \frac{j}{n}\rho_{12}\Big(\frac{j}{n}, F\big(\frac{j}{n}\big)\Big) \quad \text{for } j = 1, \cdots, n. \quad (2.3)$$

If $0 < \alpha_i, \beta_j < 1$ for each $i = 0, \cdots, n-1$, $j = 1, \cdots, n$, birth-and-death chains are an example of *reversible* Markov chains, that is, they fulfill the so-called *detailed balance* condition:

$$\mu(x)p_{xy} = \mu(y)p_{yx}, \quad (2.4)$$

where $\mu$ denotes the stationary distribution of the chain. Reversiblity constitutes an important property for the analysis of the approximation quality of Markov state models.

*Example.* Consider players in a population of size $n$ are randomly matched to play the 2x2 pure coordination game with payoff matrix $A$ given by

| | 1 | 2 |
|---|---|---|
| 1 | $a, a$ | $0, 0$ |
| 2 | $0, 0$ | $b, b$ |

where $a, b > 0$. Since there are only two strategies, we can represent the set of population states by $Z = \{0, \frac{1}{n}, \cdots, 1\}$; $\chi \in Z$ represents the proportion of players in the population playing strategy 1. For $\chi \in Z$, the payoff function $F$ is thus given by $F(\chi) = (a\chi, b(1-\chi))$. This game allows an interpretation of the strategies in term of currencies, e.g., strategy 1 represents "silver" and strategy 2 "gold". Because of this interpretation of the strategies in terms of currencies, we will call this specific population game with parameters $a$ and $b$ the *currency game*.

A prominent example of a revision protocol is the *best response with mutations* (BRM) revision protocol at *mutation rate* $\varepsilon > 0$ (Kandori et al., 1993; Young, 1993). A revising agent using the BRM revision protocol updates his strategy choice as follows: with probability $(1-\varepsilon)$ he chooses a best response $b \in B(x)$ to the current population state, while with probability $\varepsilon$ he chooses a strategy $s \in S$ at random (uniform distribution).

Now, let the birth-and-death chain $(X_t)_{t\in\mathbb{T}}$ with state space $Z = \{0, \frac{1}{n}, \cdots, 1\}$ be the stochastic evolutionary game that results from the currency game under the BRM revision protocol with noise parameter $\varepsilon$. Figure 1 gives an impression of characteristic sample paths for a resulting evolutionary process with parameters $a = b = 1, n = 11, \varepsilon = .3$. It shows the characteristic punctuated equilibrium behavior, i.e., the sample path usually stays either near the population state $\chi = 0$ or $\chi = 1$ for a long time while it switches infrequently to the other population state.

# 3 Markov State Models

Throughout what follows, let $(X_k)_{k\in\mathbb{N}}$ be an irreducible, reversible discrete-time Markov chain on a finite state space $Z = \{1, \cdots, l\}$ with transition matrix $P$ and let $\mu$ denote its unique stationary distribution.

The basic idea of a Markov state models is to approximate the original Markov process by a Markov chain on a small finite state space. Thus, more

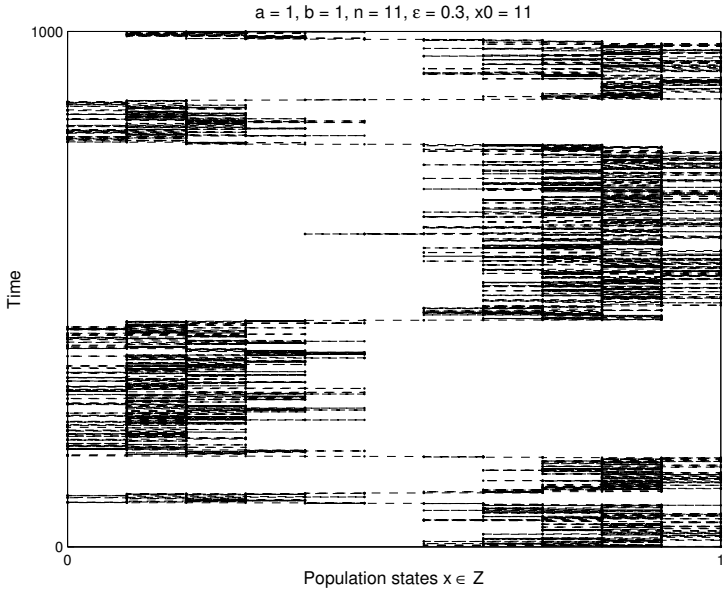


**Figure 1:** Typical sample path of the number of players holding strategy 1 in the evolutionary game defined by the currency game and the BRM revision protocol ($a = b = 1$, $n = 11$, $\varepsilon = .3$).

formally, our goal is to construct a Markov chain $(\hat{X}_k)_{k\in\mathbb{N}}$ on the state space $\hat{Z} = \{1, \cdots, m\}$ with $m$ considerably smaller than $l$ such that $(\hat{X}_k)$ captures the essential dynamics of the original Markov chain $(X_k)$. In general, each of the macrostates $i \in \hat{Z}$ corresponds to a subset of states $C_i \subset Z$, called *core sets*, where we assume that the $C_i$'s are pairwise disjoint. Roughly speaking, the idea is cluster the dynamics into core sets that may or may not partition the state space, but that represent the punctuated equilibrium dynamics in that

(i) the core sets carry most of the total statistical weight of the invariant distribution $\mu$ of the original Markov chain and

(ii) the process resides inside each core set for a long time (relative to the typical time scale of the original chain).

In Section 3.1, we first consider the special case in which the core sets $C_1, \cdots, C_m$ partition the state space $Z$, that is,

$$C_i \cap C_j = \emptyset \text{ for } i \neq j \quad \text{and} \quad \bigcup_{j=1}^{m} C_j = Z. \tag{3.1}$$

Using this special case, we demonstrate formally the basic idea of the Markov state modeling approach. Subsequently, in Section 3.2, we explain how to generalize the special case. Finally, in Section 3.3, we discuss how to actually identify suitable core sets.

### 3.1 Full partition of state space and least-squares approximation of Markov chains

In the special case in which these sets partition the state space $Z$, each macrostates $i$ directly represents the subset $C_i$ and we can define the reduced chain $(\hat{X}_k)$ on $\hat{Z} = \{1, \cdots, m\}$ with transition matrix $\hat{P} = (\hat{p}_{ij})$ by

$$\hat{p}_{ij} = \mathbb{P}[\tilde{X}_1 = j \mid \tilde{X}_0 = i], \tag{3.2}$$

where $(\tilde{X}_k)_{k\in\mathbb{N}}$ is the discrete-time process on $\hat{Z}$ that describes the dynamics of $(X_k)$ between the sets $C_1, \cdots, C_m$, i.e.,

$$\tilde{X}_k = i \Leftrightarrow X_k \in C_i. \tag{3.3}$$

Note that we have to differentiate between $(\hat{X}_k)$ and $(\tilde{X}_k)$ because $(\tilde{X}_k)$ is in general not Markovian.[4] However, we still want to approximate $(X_k)$ by a Markov chain which is why we consider $(\hat{X}_k)$.

*Example.* In our example of the currency game under BRM dynamics with parameters $a = b = 1, n = 11, \varepsilon = .3$, a reasonable partition of the set of population states into subsets such that there are only rare switches between them (see Figure 1) is $A = \{0, \cdots, 5/11\}$, $B = \{6/11, \cdots, 1\}$. The resulting matrix $\hat{P}$ is given by (rounded to four digits)

$$\hat{P} = \begin{pmatrix} .9989 & .0011 \\ .0011 & .9989 \end{pmatrix}. \tag{3.4}$$

In order to appreciate the approximation properties of the reduced chain $(\hat{X}_k)$ it is helpful to analyse the relation between the transition probabilities $p_{ij}$ and $\hat{p}_{ij}$. To this end, suppose that the original chain starts in equilibrium, $X_0 \sim \mu$, with $\mu = P^T\mu$ being the unique stationary distribution of $(X_k)$. Now, the transition probability $\hat{p}_{ij} = \mathbb{P}_\mu[X_1 \in C_j \mid X_0 \in C_i]$ in (3.2) can be recast as

$$\hat{p}_{ij} = \frac{\sum_{k,j} p_{kl}\chi_{C_i}(l)\chi_{C_j}(k)\mu(k)}{\sum_k \chi_{C_i}(k)\mu(k)}, \tag{3.5}$$

where $\chi_C \colon Z \to \{0, 1\}$ denotes the indicator function of a set $C \subset Z$, and the notation $\mathbb{P}_\mu$ indicates that $X_0$ is distributed according to $\mu$.

Let us assume that the indicator functions $\chi_{C_1}, \ldots, \chi_{C_m}$ form a partition of unity, i.e., $\sum_i \chi_{C_i}(x) = 1$, which is the case if the $C_1, \ldots, C_m$ partition our state space $Z$. The last equation can then be interpreted as the orthogonal projection onto the span of the functions $\chi_{C_1}, \ldots, \chi_{C_m}$ with respect to the $\mu$-weighted scalar product

$$\langle f, g\rangle_\mu = \sum_{k\in Z} f(k)g(k)\mu(k) \tag{3.6}$$

[4] The reason why $(\tilde{X}_k)$ is not Markovian is called the *recrossing problem*. This name refers to the issue that transitions between the subsets of state space are much more likely at the boundaries of the sets. This is, however, not such a big issue if the original process shows a strongly punctuated equilibrium dynamic behavior since the probability to be at the boundary is in this case negligible.

on $\mathbb{R}^l$. A compact way to write (3.5) thus is

$$\hat{p}_{ij} = \frac{\langle P\chi_{C_i}, \chi_{C_j}\rangle_\mu}{\langle \chi_{C_i}, \chi_{C_i}\rangle\mu}\,, \tag{3.7}$$

which shows that the corresponding transition matrix $\hat{P} = (\hat{p}_{ij})_{i,j\in\hat{Z}}$ is in fact the orthogonal projection of $P = (p_{ij})_{i,j\in Z}$ onto $\mathrm{span}\{\chi_{C_1}, \ldots, \chi_{C_m}\}$, understood as a linear subspace of $\mathbb{R}^l$ endowed with the weighted scalar product $\mu$.

By being an orthogonal projection, $\hat{P}$ is the best approximation of $P$ onto the space spanned by the indicator functions of the core sets $C_1, \ldots, C_m$ in the sense of least squares where the weighting with the invariant measure $\mu$ arises naturally as a consequence of the fact that the Markov chain is initialized in its stationary distribution so as to make the macroscopic transition probabilities time-independent. Further notice that, since the $\chi_{C_i}$ are non-negative and form a partition of unity, the projected transition matrix $\hat{P}$ is a stochastic matrix and inherits many important properties of the original transition matrix $P$:

1. If $P$ is irreducible and aperiodic, then so is $\hat{P}$.
2. $\hat{P}$ has a unique invariant distribution $\hat{\mu} = (\hat{\mu}(i))_{i\in\hat{Z}}$ that equals the marginal distribution of the core sets $C_1, \ldots, C_m$:

$$\hat{\mu}(i) = \langle \chi_{C_i}, \chi_{C_i}\rangle_\mu = \mu(C_i)\,, \quad i \in \hat{Z}. \tag{3.8}$$

3. If $P$ is reversible with respect to $\mu$, then $\hat{P}$ is reversible with respect to $\hat{\mu}$.

A further advantage of (3.7) is that it tells us that, given a long realization of the original Markov chain $(X_t)$ of length $T$, the expression

$$\hat{p}_{ij}^{(T)} = \frac{\sum_{t=1}^{T} \chi_{C_i}(X_t)\chi_{C_j}(X_{t+1})}{\sum_{t=1}^{T}(\chi_{C_i}(X_t))^2} \tag{3.9}$$

is an unbiased estimator of the macroscopic transition probabilities $\hat{p}_{ij}$. By the assumption that $\mu$ is unique and $Z$ is finite, the law of large numbers implies that $\hat{p}_{ij}^{(T)}$ converges almost surely to $\hat{p}_{ij}$ as $T \to \infty$ for every initial value $X_0 = 0$.

## 3.2 General core sets and sparse least-squares approximation

The case of a full partition of state space demonstrates the basic idea of Markov state models as a coarse-grained Markov chain that can be obtained by projection onto suitable ansatz functions. In the general case, however, the sets $C_1, \cdots, C_m$ do not necessarily partition the state space $Z$; thus, the approach has to be modified since already the definition of the process $\tilde{X}$ in Eq. (3.3) is not well defined anymore.

In order to construct a reduced Markov chain that best approximates our original Markov chain in this case, the idea is to replace the set of indictator

functions by a clever "mollification", forming a partition of unity and having support outside the sets $C_1, \ldots, C_m$. One such choice is the set of probabilistic ansatz functions $\{q_1, \ldots, q_m\}$, so-called *committor functions*, defined by

$$q_i(x) = \mathbb{P}_{\delta_x}(\tau^0_{C_i} < \tau^0_{C\setminus C_i}), \tag{3.10}$$

where $C = \cup_i C_i$, $\delta_x$ is the point mass on $x$, and $\tau^k_A = \inf\{k' \geq k | X_{k'} \in A\}$ denotes the $k$-th *hitting time* for $k \geq 0$. In words, the committor function $q_i : Z \to [0,1]$ is the function that gives for a state $x \in Z$ the probability that the Markov chain $(X_k)$ will visit the set $C_i$ first rather than $C \setminus C_i$. By construction, each $q_i$ is equal to one on $C_i$, equal to zero on the other sets $C_j$, $j \neq i$, and interpolates between these values outside the sets $C_1, \ldots, C_m$. Moreover, since $(X_t)$ is irreducible and positive recurrent (by $Z$ being finite), the process terminates after finite time with probability one by hitting one of the sets $C_i$, independently of the initial condition $X_0 = x$, and as a consequence the $q_i$ sum up to one and form a partition of unity.

The analysis of the microscopic dynamics that is carried out in Sarich (2011) shows that the reduced Markov chain on $\hat{Z}$ can be defined in terms of the *quasi*-transition matrix $\hat{P}W^{-1}$, where the matrices $\hat{P}$ and $W$ are given by

$$\hat{P}(i,j) = \mathbb{P}(\tau^{k+1}_{C_j} < \tau^{k+1}_{C\setminus C_j} \mid \tilde{X}_k = i), \tag{3.11}$$

$$W(i,j) = \mathbb{P}(\tau^{k}_{C_j} < \tau^{k}_{C\setminus C_j} \mid \tilde{X}_k = i), \tag{3.12}$$

where $(\tilde{X}_k)$ is the *milestoning process* defined by

$$\tilde{X}_k = i \Leftrightarrow X_{\sigma(k)} = i, \text{ where } \sigma(k) = \max\{t \leq k \mid X_t \in C\}. \tag{3.13}$$

Equation (3.13) means that the milestoning process remains in state $i$ as long as the original Markov chain $(X_k)$ last visited core set $i$ (see Figure 2). Thus, in words, $W(i,j)$ for $j \neq i$ gives the probability that the Markov chain next hits $C_j$ while being in a state in $Z \setminus C$ at some time $k$ and last came from core set $C_i$, where $C = \cup_{j=1}^m C_j$. Similarly, $\hat{P}(i,j)$ gives the probability that the next core set hit is $C_j$ conditional on having hit the core set $C_i$ last at some time $k$. Moreover, each macrostate $i \in \hat{Z}$ is associated with the respective committor function $q_i$ on the core set $C_i$ and can thus be interpreted as representing the *affiliation* with set $C_i$. Note that while the definition of the quasi-transition matrix of our Markov state model by $\hat{P}W^{-1}$ might not seem intuitively obvious, it reduces to the matrix $\hat{P}$ defined in Eq. (3.2) in the case of a full partition of state space.

We call $\hat{P}W^{-1}$ a quasi-transition matrix since $\hat{P}W^{-1}$ is not always a stochastic matrix (even though $P$ and $W$ are). We only know that its rows sum up to one since this is the case for both $\hat{P}$ and $W$, and thus also for $W^{-1}$ as well as $\hat{P}W^{-1}$. In the example given here as well as in the examples studied in Hallier (2015) the entries of $\hat{P}W^{-1}$ are also non-negative, but in general the entries can be negative as has been pointed out in Sarich (2011).

It is possible to show, however, that $\hat{\mu}$ defined by

$$\hat{\mu}(j) = \sum_{i \in Z} \mu(i) q_j(i) \tag{3.14}$$

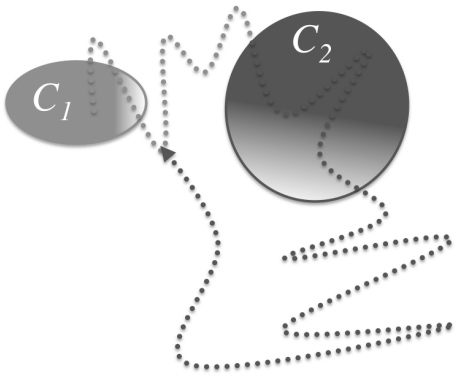


**Figure 2:** Illustration of the milestoning process for two core sets $C_1$ and $C_2$.

is the unique ergodic stationary distribution of $(\hat{X}_k)$. Unlike in the case of a full state space partition, the matrix $\hat{P}W^{-1}$ does not trivially inherit all properties of the original chain, such as irreducibility, aperiodicity and reversibility; but $(\hat{X}_k)$ is reversible with respect to $\hat{\mu}$ if $\hat{P}$ and $W^{-1}$ commute.

Despite the apparent lack of structure preservation of the sparse core set approximation shows excellent spectral approximation properties, in that the dominant eigenvalues of the original chain are generally well approximated. The latter implies that the projected transition matrix can be used to accurately estimate transition rates between the core sets as well as mean residence times, and hence residence times and rates for the punctuated equilibria. Moreover, both matrices $W$ and $\hat{P}$ can be estimated from trajectory data in the following way: given a realization $(x_0, \cdots, x_K)$ of $(X_k)$ of length $K$, we can estimate

$$W^{*,K}(i,j) = \begin{cases} \frac{R_{ij}^K}{r_i^K} & \text{if } j \neq i, \\ 1 - \sum_{j \neq i} W^{*,K}(i,j) & \text{otherwise,} \end{cases} \tag{3.15}$$

$$\hat{P}^{*,K}(i,j) = \frac{R_{ij}^{+,K}}{r_i^K}, \tag{3.16}$$

where $R_{ij}^K$ denotes the number of times where the chain came from core set $C_i$, is in a state in $Z \setminus C$ and hits $C_j$ next, $r_i^K$ is the total number of time steps the trajectory was in $i$; that is, $\tilde{X}_k = i$, and $R_{ij}^{+,K}$ denotes the number of times where the chain came from core set $C_i$ and hit $C_j$ next.

*Example.* Consider again the stochastic evolutionary game given in the last section. Let $(X_t)_{t \in \mathbb{T}}$ be the birth-and-death chain with state space $z = \{0, \frac{1}{11}, \cdots, 1\}$. We simulated the Markov chain with $a = 1, b = 1, \varepsilon = .3, n = 11$ to get a trajectory of population states $(x_t)$ for $t = 1/11, \ldots, 50000$ and thus of length $5.5 \cdot 10^5$. If we consider, for instance, the core sets $C_1 = \{0, 1/11\}$ and $C_2 = \{10/11, 1\}$,

the estimated transition matrix $\hat{P}^* W^{*,-1}$ is given by

$$\hat{P}^* W^{*,-1} = \begin{pmatrix} .9993 & .0007 \\ .0007 & .9993 \end{pmatrix}. \tag{3.17}$$

An analysis of the approximation error made (Hallier, 2015) shows that this is a better approximation of the underlying stochastic evolutionary game than the full partition Markov state model we considered above. In the example given, it shows that the approximation error is considerably smaller for the core set model, especially in the case of smaller population sizes and larger values of the noise parameter $\varepsilon$. This is precisely the case where the other approximation approaches of stochastic evolutionary games cannot be applied. This example thus demonstrates that core set Markov state models fill a gap and constitute an important complement to existing approximation approaches.

*Remark* 3.1. Adopting a slightly more abstract point of the described clustering framework proves useful in deriving systematic error bounds for the Markov chain approximation. The idea is to view the clustering of the original Markov chain as a projection onto a linear subspace of the Hilbert space

$$H = \left\{ f \colon Z \to \mathbb{R} \, : \, \sum_{k \in Z} (f(k))^2 \mu(k) < \infty \right\} \tag{3.18}$$

of square summable functions endowed with the weighted scalar product $\langle \cdot, \cdot \rangle_\mu$. For example, in the first mentioned case, the $\mu$-orthogonal projection of a function $f \in H$ onto $\operatorname{span}\{\chi_{C_1}, \ldots, \chi_{C_m}\}$ can be understood as the best approximation of $f$ by functions that are measurable with respect to the partition $\{C_i \colon i = 1, \ldots m\}$, measured in the natural norm on $H$ that is induced by $\langle \cdot, \cdot \rangle_\mu$. In other words, the macroscopic transition probabilities $\hat{p}_{ij}$ are the conditional expectation, and hence least squares approximation, of the microscopic transition probabilities $p_{ij}$, given only information about the macrostates $C_i$.

In the Markov state modelling approach due to Schütte and co-workers (Deuflhard and Weber, 2005; Huisinga and Schmidt, 2006; Schütte and Sarich, 2013), the object of interest that is amenable to systematic approximations is the *transfer operator*, a family of linear operators $\mathcal{P}_t : \mathcal{M}_f^+ \to \mathcal{M}_f^+$ that map any finite non-negative Borel measure to a finite non-negative Borel measures. For our purposes, the transfer operator can be defined by

$$\mathcal{P}_t \nu = \left(P^T\right)^t \nu \,, \; t \in \mathbb{N}_0 \,, \tag{3.19}$$

where $P$ is the transition matrix of the original Markov chain $(X_t)$ and $\nu$ is a counting density, understood as a non-negative column vector in $\mathbb{R}^l$. Note that $\mathcal{P}_t$ preserves the $L^1$ norm, i.e., if $\nu$ is a probability density, then so is $\mathcal{P}_t \nu$ for all $t$. Given an initial distribution of $X_t$ at $t = 0$, the transfer operator can be used to determine the probability distribution of $X_t$ at a later time $t > 0$, i.e., it

transfers probability distributions in time and thus encodes information about the law of the process $(X_t)$. Within the finite-state Markov chain framework considered in this paper, approximating the transfer operator amounts to finding a suitable low-rank approximation of the transition matrix $P$.

Understanding the approximation of the original dynamics as a projection of the associated transfer operator, it is possible to show that the approximation error made with respect to the propagation of probability distributions as well as in terms of the dominant eigenvalues, which directly relate to the longest timescales of the original Markov chain, crucially depends on the projection error made; a small projection error implies a good approximation quality of our Markov state models (Hallier, 2015; Sarich, 2011).

### 3.3 Identification of core sets

While we outlined above how to construct Markov state models given core sets $C_1, \cdots, C_m$, the question remains of how to actually identify suitable core sets. One approach is to use the results on the relationship between the approximation quality and the projection error. For full partition models, the projection error is as small as possible if the dominant right eigenvectors of the transfer matrix $P$ are as constant as possible on the sets of the partition. This relationship has been exploited by approaches that partition state space by clustering algorithms (as has, for example, been done in the molecular dynamics context by Krivov and Karplus, 2004; Noé et al., 2007; Rao and Caflisch, 2004). Similarly, in terms of core set Markov state models, finding core sets $C_1, \cdots, C_m$ so that the projection error is as small as possible can be interpreted as a fuzzy clustering problem (Djurdjevac, 2012; Sarich, 2011).

In the case of stochastic evolutionary games with a noise parameter that determines the punctuated equilibrium behavior, we might also use the information given by its stationary distribution to identify possible core sets. More specifically, if the system under investigation depends on a noise parameter $\varepsilon > 0$ in such a way that for smaller $\varepsilon$ the punctuated equilibrium behavior increases, that is, for smaller $\varepsilon$ the sample paths of the process stay even longer in certain subsets of the population state space and the switches between such subsets get more rare – just as in our example. In such a case, we can identify the set $C = \cup_{i=1}^{m} C_i$ by comparing the stationary distribution $\mu^*$ of the system with noise level $\varepsilon^*$ with its propagated distribution $\mathcal{P}_\theta \mu^* = (P^T)^\theta \mu^*$ under a decreased noise level $\varepsilon < \varepsilon^*$ for a chosen timescale $\theta > 0$ (see Remark 3.1). Note that the $\mathcal{P}_t \mu^*$ of $\mu^*$ converges to the stationary distribution $\mu$ associated with the stochastic evolutionary game at the lower noise level $\varepsilon$. Moreover, both stationary distributions $\mu$ and $\mu^*$ have the same form in the sense of local minima and maxima, but the stationary distribution $\mu^*$ with increased noise intensity is less peaked. Now, the basic idea of the identification strategy is that a population state $x$ belongs to the core set region $C$ if it gets more attractive in the stochastic evolutionary game with the decreased noise parameter $\varepsilon$, i.e., if

$$\mu^*(x) < (\mathcal{P}_\theta \mu^*)(x). \tag{3.20}$$

*Example.* In our example with parameters $a = 1$, $b = 1$, $\varepsilon^* = .3$, $\varepsilon = .15$, $n = 11$, this identification approach leads to the set $C = \{0, 1/11, 10/11, 1\}$ for all $\theta \in \mathbb{T}$, which suggests the core sets $C_1 = \{0, 1/11\}$ and $C_2 = \{10/11, 1\}$, see also Figure 3.

As in this example, the clustering of $C$ into core sets is usually straightforward as the core sets are dynamically well separated. An advantage of the just sketched approach is that the necessary quantities can be estimated from trajectory data as well. Thus, it allows for a simulation-based approach to the construction of Markov state models. For more details, see Hallier (2015).

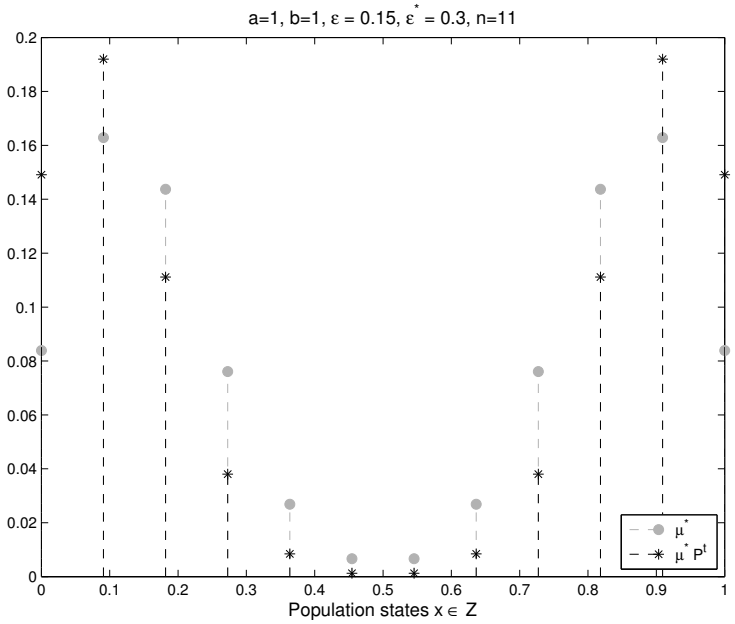


**Figure 3:** Weights of the stationary distribution $\mu^*$ for the stochastic evolutionary game of our running example with parameters $a = 1, b = 1, \varepsilon^* = .3, n = 11$ and its propagation $\mathcal{P}^t(\mu^*)$ under the stochastic evolutionary game with parameters $a = 1, b = 1, \varepsilon = .15, n = 11, t = 10/11$.

# 4 Discussion and Outlook

We presented the Markov state modeling approach to extract the aggregated long term dynamics of stochastic evolutionary games. The approach is especially interesting for large, complex games in order to see the wood for the trees. In essence, Markov state models approximate the original Markov chain on a reduced state space. The transition probabilities between the macrostates can be estimated on the basis of short-term trajectory data. While this basic idea behind Markov state models has informally been used before (e.g., Kandori

et al., 1993), it is for the first time that the approach is formally considered and its approximation quality can be assessed.

Apparent advantages of a reduced state space are that it is easier to compute eigenvalues and eigenvectors as well as other properties such as waiting times. In addition, as our example demonstrates, the approximation of stochastic evolutionary games with punctuated equilibrium dynamics by Markov state models is especially interesting for small population sizes and larger mistake probabilities as traditional approximation techniques such as deterministic approximation or stochastic stability analysis are not applicable or are of limited informative value.

One limitation is, however, that the approach and its analysis depends on the original Markov chain that captures the aggregated dynamics of the stochastic evolutionary game to be reversible. This is the case for the example we presented. More generally, stochastic evolutionary games that result from population game with two strategies and full support revision protocols as well as those that result from finite-population games with clever agents under a logit choice revision protocol can be shown to have reversible dynamics (see Sandholm, 2010, Chapter 11.5.3, in addition, provides more general conditions on revision protocols under which finite-population potential games with clever agents result in reversible dynamics).

In general, it will be difficult to say whether it is reasonable to assume that a stochastic evolutionary game results in a reversible Markov chain. One reason for this difficulty is that, if we estimate the transition matrix from simulated trajectory data, it does not need to fulfill the detailed balance equation, even if the underlying Markov chain is reversible (Noé, 2008; Prinz et al., 2011). In the context of molecular dynamics, however, it was possible to derive approximative models that can be proven to be reversible although the original model is not. An example is the diffusion model, which represents an approximation to the Langevin model in the limit of high friction (see, e.g., Schütte and Sarich, 2013, Chapter 2 and references therein). As a future research question, it seems worthwhile to explore whether similar results can be obtained for stochastic evolutionary games; that is, whether there are approximations of certain stochastic evolutionary games that can be shown to be reversible.

Beyond that, we would like an approach that applies also to non-reversible Markov dynamics. Notice that it is not difficult to derive a construction of a matrix representation of the core set Markov state models for given core sets in the case of non-reversible Markov chains (see, e.g., Djurdjevac, 2012; Djurdjevac et al., 2010). However, we neither have results with respect to their approximation quality nor an approach to the identification of core sets for non-reversible Markov chains. One fundamental problem is that the eigenvalues and eigenvectors of the transfer matrices corresponding to non-reversible Markov chains need not be real anymore. In this case, the interpretation of the spectral information is unclear. Up to now, there are few approaches that apply also to non-reversible Markov chains (Eckhoff, 2002; Gaudillière and Landim, 2011; Horenko, 2011; Sarich and Schütte, 2014). A graph-theoretical framework for constructing reversible surrogates of non-reversible dynamics, based on a cycle

decomposition of the underlying Markov chain, has been suggested in Banisch (2015), however the applicability to evolutionary games is yet open.

The identification of Markov state models for general stochastic evolutionary games is therefore an open problem and will be a topic of future research.

# References


Banisch, R. (2015). *Markov Processes Beyond Equilibrium and Optimal Control*. PhD thesis, Fachbereich Mathematikund Informatik, FU Berlin.

Benaim, M. and Weibull, J. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.

Bovier, A. (2009). Metastability. In Kotecký, R., editor, *Methods of Contemporary Mathematical Statistical Physics*, Lecture Notes in Mathematics 1970. Springer-Verlag Berlin Heidelberg.

Deuflhard, P. and Weber, M. (2005). Robust perron cluster analysis in conformation dynamics. *Linear Algebra and its Applications*, 398:161184.

Djurdjevac, N. (2012). *Methods for analyzing complex networks using random walker approaches*. PhD thesis, Fachbereich Mathematikund Informatik, FU Berlin.

Djurdjevac, N., Sarich, M., and Schütte, C. (2010). On Markov state models for metastable processes. In *Proceedings of the International Congress of Mathematicians, Hyderabad, India*.

Eckhoff, M. (2002). The low lying spectrum of irreversible, infinite state Markov chains in the metastable regime. Preprint.

Ellison, G. (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, 67(1):17–45.

Foster, D. and Young, P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232.

Gaudillière, A. and Landim, C. (2011). A Dirichlet principle for non reversible Markov chains and some recurrence theorems. *Probability Theory and Related Fields*, 158:55–89.

Hallier, M. (2015). *Formalization and Metastability Analysis of Agent-Based Evolutionary Models*. PhD thesis, Fachbereich Mathematikund Informatik, FU Berlin.

Hofbauer, J. and Sigmund, K. (2003). Evolutionary game dynamics. *Bulletin (New Series) of the American Mathematical Society*, 40(4):479–519.

Horenko, I. (2011). Parameter identification in nonstationary Markov chains with external impact and its application to computational sociology. *SIAM Multiscale Modeling & Simulation*, 9(4):1700–1726.

Huisinga, W. and Schmidt, B. (2006). Metastability and dominant eigenvalues of transfer operators. In Leimkuhler, B., Chipot, C., Elber, R., Laaksonen, A., Mark, A., Schlick, T., Schütte, C., and Skeel, R., editors, *New Algorithms for Macromolecular Simulation*, volume 49 of *Lecture Notes in Computational Science and Engineering*, pages 167–182. Springer.

Jaeger, C. (2008). Mathematics and the social sciences: The challenge ahead. Introductory contribution to the Dahlem Conference: "Is There a Mathematics of Social Entities"?, December 2008, Berlin, Germany. Unpublished manuscript.

Jaeger, C. C. (2012). Scarcity and coordination in the global commons. In Jaeger, C. C., Hasselmann, K., Leipold, G., Mangalagiu, D., and Tábara, J. D., editors, *Reframing the Problem of Climate Change: From Zero Sum Game to Win-Win Solutions*, pages 85–101. Routledge.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Krivov, S. V. and Karplus, M. (2004). Hidden complexity of free energy surfaces for peptide (protein) folding. *Proceedings of the National Academy of Sciences of the United States of America*, 101:1476614770.

Kurtz, T. (1970). Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7:49–58.

Noé, F. (2008). Probability distributions of molecular observables computed from Markov models. *The Journal of Chemical Physics*, 128:244103.

Noé, F., Horenko, I., Schütte, C., and Smith, J. C. (2007). Hierarchical analysis of conformational dynamics in biomolecules: Transition networks of metastable states. *The Journal of Chemical Physics*, 125:155102.

Prinz, J.-H., Wu, H., Sarich, M., Keller, B., Senne, M., Held, M., Chodera, J. D., Schütte, C., and Noé, F. (2011). Markov models of molecular kinetics: Generation and validation. *The Journal of Chemical Physics*, 134(174105).

Rao, F. and Caflisch, A. (2004). The protein folding network. *Journal of Molecular Biology*, 342(1):299306.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.

Sandholm, W. H. (2011). Stochastic evolutionary game dynamics: foundations, deterministic approximation, and equilibrium selection. Sigmund, Karl (ed.), Evolutionary game dynamics. American Mathematical Society short course,

January 4–5, 2011, New Orleans, Lousiana, USA. Providence, RI: American Mathematical Society (AMS). Proceedings of Symposia in Applied Mathematics 69, 111-141 (2011).

Sarich, M. (2011). *Projected Transfer Operators*. PhD thesis, Fachbereich Mathematikund Informatik, FU Berlin.

Sarich, M. and Schütte, C. (2014). Utilizing hitting times for finding metastable sets in non-reversible Markov chains. Technical Report 14-32, Konrad-Zuse-Zentrum für Informationstechnik Berlin.

Schütte, C. and Sarich, M. (2013). *Metastability and Markov State Models in Molecular Dynamics: Modeling, Analysis, Algorithmic Approaches*, volume 24 of *Courant Lecture Notes*. American Mathematical Society.

Weibull, J. W. (1995). *Evolutionary Game Dynamics*. MIT Press, Cambridge.

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

Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press.

Young, H. P. (2006). Social dynamics: Theory and applications. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics: Agent-Based Computational Economics*, volume 2, chapter 22, pages 1081–1108. North-Holland.