Applied Probability Trust (17 October 2014)

ANALYSIS OF A STOCHASTIC APPROXIMATION ALGORITHM FOR COMPUTING QUASI-STATIONARY DISTRIBUTIONS

J. BLANCHET,* *Columbia University*

P. GLYNN,** *Stanford University*

S. ZHENG,*** *Columbia University*

Keywords: Stochastic approximations, quasi-stationary distribution, ODE method.

Abstract

This paper analyzes the convergence properties of an iterative Monte Carlo procedure proposed in the Physics literature for estimating the quasi-stationary distribution on a transient set of a Markov chain (see, [10, 9, 12]). In contrast to existing linear algebra methods, this approach eliminates the need for explicit transition matrix manipulations in order to compute the principal eigenvector. Our paper analyzes the procedure proposed in the physics literature by casting it as a stochastic approximation algorithm ([26, 19]). Using this connection we are able to not only verify the consistency of the estimator but also provide a rate of convergence in the form of a Central Limit Theorem. We provide a simple example showing that convergence might occur very slowly and indicate how this issue can be alleviated by using averaging.

2010 Mathematics Subject Classification: Primary 60J22
Secondary 60J10

1. Introduction

The motivation for the algorithm that we study here comes from the physics literature. The papers [9, 10, 12] propose and study an iterative Monte Carlo procedure to estimate the quasi-stationary distribution of an interacting particle system (IPS),

* Postal address: Columbia University, Department of Industrial Engineering & Operations Research
* Email address: jose.blanchet@columbia.edu
** Postal address: Stanford University, Department of Management Science & Engineering
** Email address: glynn@stanford.edu
*** Postal address: Columbia University, Department of Industrial Engineering & Operations Research
JohnZ622@gmail.com

[20]. As we shall review, a quasi-stationary distribution can be computed via the left principal eigenvector of the substochastic matrix corresponding to non-absorbing states. Consequently, it is natural to use numerical linear algebra methods for computing such principal eigenvector. Nevertheless, the application of these methods, and related Monte Carlo variations, such as [13, 17, 18, 23], become difficult when the underlying matrix is large. In particular, the application of these methods is specially prohibitive in IPS. In contrast, the approach described in [9, 10, 12] only uses a small portion of the underlying matrix in each iteration and, thus, the method can be executed in the setting of huge state spaces.

The Fleming-Viot method, [6, 14, 22], provides a powerful alternative method for computing quasi-stationary distributions. It consists of N particles evolving according to suitably dynamics in continuous time. As both time, t, and the number of particles, N, tend to infinity, the empirical measure of the positions of the particles at time t converges almost surely to the underlying quasi-stationary distribution. A significant advantage of the method is that it can be run in parallel. A disadvantage is that, for a fixed value t, if only N is send to infinity, the method will not converge to the underlying quasi-stationary distribution.

In contrast to the Fleming-Viot method, the method that we analyze here is asymptotically unbiased as the number of iterations, N, tends to infinity. Moreover, as we shall show, a small modification can be made to the method to ensure convergence at rate $N^{-1/2}$. Note that such rate of convergence is impossible to achieve in the Fleming-Viot method because of the presence of the bias appearing by truncating the evolution at time t.

The method suggested in [9, 10, 12], it turns out, is equivalent to a class of algorithms studied in the context of urn processes, [1, 2, 24]. So, the convergence of the sampling procedure has been rigorously established in the urn process literature. Moreover, in [2], some results on rates of convergence have been obtained. These results involve a Central Limit Theorem (CLT) for the inner product of the following two quantities: any non-principal eigenvector (or linear combinations thereof) and the estimated quasi-stationary vector (i.e. the estimated principal eigenvector). One of our contributions in this paper is the development of a multidimensional CLT for the estimated quasi-stationary vector. Therefore, we can obtain a CLT for the inner product between the estimated quasi-stationary vector and any other vector (not necessarily one which must be represented as the linear combination of non-principal eigenvectors). More generally, our main contributions are as follows:

- Our paper recognizes the algorithm in [9, 10, 12] as a stochastic approximation algorithm (Section 3.2).
- Using the stochastic approximation connection we prove the convergence of the underlying estimator and provide sufficient conditions for a CLT (Theorem 3.1)

which is not restricted to inner products involving linear combinations of non-principal eigenvectors.

- We recognize common scenarios based on the spectral gap of the underlying substochastic matrix (Section 4) where the CLT of the estimator fails and convergence is slow.
- More importantly, using Polyak-Ruppert averaging, [25], we suggest an improved algorithm (Section 4.2.1) which exhibits a valid CLT under no additional restriction on the spectral gap of the underlying substochastic matrix.
- We provide an estimator which allows us to compute the variance in the CLT for quasi-stationary expectations, see Section 4.2.2.

We concentrate on discrete-time Markov chains. The adaptation of our results to continuous-time Markov chains is relatively straightforward and it is given in [27]. The convergence of the estimator in general state-space chains is also studied in [27].

The vanilla version of the algorithm analyzed here (without averaging) has independently been studied in [3]. In contrast to [3], our focus is more algorithmic. In particular, our emphasis on exploring the close connection of our algorithm to stochastic approximation leads naturally to a Polyak-Ruppert averaging variant that exhibits an optimal square root convergence rate under modest side conditions, independent of the problem data, in contrast to the original algorithm that often displays sub-square root convergence rates. Given our algorithmic focus, we also discuss estimation of a variance parameter that enters the construction of confidence intervals that naturally arise in connection with our proposed Monte Carlo estimation procedure.

We refer the readers to [27] for numerical experiments showing the dramatic improvement of the algorithmic variations that we introduce here, in particular, averaging; see also [4].

The rest of the paper is organized as follows. In Section 2, we review some background material on quasi-stationary distributions. In Section 3 we provide a quick overview of stochastic approximation methods and we also sketch the proof of convergence (the full proof is given in Sections 5.1 and 5.2). In Section 4.2.1 discusses an improved version of the algorithm using projection along with averaging. All the proofs and technical results are elaborated in Section 5.

2. Quasi-stationary Distribution: Basic Notions

An early reference for quasi-stationary distributions is the work in [7]. Let $\{X_n : n \geq 0\}$ be a discrete-time, finite-state-space, Markov chain $\{X_n : n \geq 0\}$ with transition matrix P. Let us assume that 0 is an absorbing state and that $1, ..., d$ are non-absorbing. We can partition P as

$$\mathcal{P} = \left[\begin{array}{cc} 1 & 0 \\ \alpha & Q \end{array} \right],$$

where the upper left corner represents the fact that $1 = P(X_1 = 0 \mid X_0 = 0)$. We then have that $\alpha(i) = P(X_1 = 0 \mid X_0 = i)$ and Q is a substochastic transition matrix of size d by d.

Now, given a probability distribution π supported on $\{1, ..., d\}$ we can define, for all $n \geq 1$,

$$\mu_j^{\pi}(n) = P(X_n = j \mid X_0 \sim \pi, X_1, ..., X_{n-1} \notin \{0\}) = \frac{\pi' Q^n \mathbf{e}_j}{\pi' Q^n \mathbf{e}},$$

where $\{\mathbf{e}_j\}_{j=1}^d$ is the standard basis for $\mathbb{R}^d$ and $\mathbf{e}$ is a vector with all coordinates equal to one. We use the notation $X_0 \sim \pi$ inside the conditioning to indicate that X_0 follows the distribution π. Throughout the rest of the paper we use " ′ " to denote transposition. Then, we can provide the definition of a quasi-stationary distribution.

Definition 1. If there is a distribution π over the transient states $\{1, ..., d\}$ such that $\bar{\mu}^{\pi}(n) := (\bar{\mu}_j^{\pi}(n) : 1 \leq j \leq n)$ is independent of m, then we call $\mu^{\pi} := \mu^{\pi}(n)$ a quasi-stationary distribution.

Under the assumption that the substochastic matrix Q is irreducible (although not necessarily aperiodic) it is straightforward to see (from the Perron-Frobenius Theorem, [16]) that there exists a unique quasi-stationary distribution which can be computed via the solution to the principal eigenvector problem

$$\mu_*' Q = \lambda_* \mu_*',$$

where $\lambda_* > 0$. The paper [22] discusses the existence of quasi-stationary distributions for infinite spaces.

Assumption 1. *Throughout the rest of the paper we shall assume that Q is irreducible and that $Q^n \to 0$ as $n \to \infty$.*

3. Stochastic Approximations Analysis of the Algorithm

3.1. Basic Notions of Stochastic Approximations

We focus on one of the simplest martingale-difference-type noise driving the stochastic approximation sequence which takes the form

$$\theta_{n+1} = \theta_n + \varepsilon_n W_n, \tag{1}$$

for $n \geq 0$, where $\{\varepsilon_n\}_{n \geq 0}$ is a step-size sequence (user defined) of non-negative numbers and the n-th noise observation, W_n, depends on the whole history $\mathcal{F}_n =$

$\sigma\{(\theta_k, W_{k-1}) : 1 \le k \le n\}$. The initial condition θ_0 is given. We require the existence of a function $g(\cdot)$ such that

$$E(W_n|\mathcal{F}_n) = g(\theta_n).$$

The step-size sequence is such that $\sum \varepsilon_n = \infty$ but $\sum \varepsilon_n^2 < \infty$ and under mild regularity conditions, to be reviewed momentarily in our setting, we have that θ_n converges almost surely to the stable attractors of the ODE

$$\dot{\theta}(t) = g(\theta(t)),$$

see, for example, Theorem 5.2.1 in [19].

3.2. The Precise Algorithm in Stochastic Approximation Form

In simple words the algorithm proceeds as follows. Suppose that we have d bins (one for each element in the underlying transient set). At the beginning of the n-th iteration we have a certain distribution of balls across the bins and we select an initial position according to such distribution. For example, if $d = 2$ and there are 3 balls in the first bin and 5 balls in the second bin, then state 1 is selected with probability 3/8 and state 2 is selected with probability 5/8. The first iteration then proceeds by running a Markov chain starting from the selected state $i \in \{1, ..., d\}$ according to the underlying dynamics, until absorption (i.e. until hitting state 0) and we call such a trajectory a tour. We count the number of times state j is visited during such tour (for $j \in \{1, ..., d\}$, note for example that the i-th state is visited at least once) and update the distribution of balls across bins by adding these counts accordingly. So, for example if $d = 2$ and during the tour state 1 was visiting 2 times, while state 2 was visited 4 times, then the distribution of balls at the beginning of the $(n+1)$-th iteration will be (5=3+2, 9=5+4). The output of the algorithm is the normalized distribution of balls (in order to obtain a probability vector) after many iterations as an estimate of the quasi-stationary distribution.

We now explain how this procedure can be described in terms of a stochastic approximation recursion.

Notation:

- μ_n is the sequence of probability vectors of the transient set $\{1, ..., d\}$ obtained at the n-th iteration of the algorithm. This vector will store the cumulative empirical measure up to, and including, the n-th iteration of the algorithm. We use $\mu_n(x)$ to denote the particular value at the transient state x.
- $\{X_k^{(n)}\}_{k\ge 0}$ is the Markov chain ran during the n-th iteration of the algorithm. These Markov chains are conditionally independent (given the values X_0^n). The n-th Markov chain has an initial position drawn from the vector μ_n.
- We define $\tau^{(n)} = \inf\{k \ge 0 : X_k^{(n)} = 0\}$. (Recall that 0 is the underlying absorbing state.)

We are interested in analyzing the recursion

$$\mu_{n+1}(x) = \frac{\left(\sum_{k=0}^{n} \tau^{(k)}\right) \mu_n(x) + \left(\sum_{k=0}^{\tau^{(n+1)}-1} I\left(X_k^{(n+1)} = x \mid X_0^{(n)} \sim \mu_n\right)\right)}{\left(\sum_{k=0}^{n+1} \tau^{(k)}\right)}, \quad (2)$$

for all $x \in \{1, ..., d\}$, where the notation

$$I\left(X_k^{(n+1)} = x \mid X_0^{(n)} \sim \mu_n\right)$$

described the indicator of the event $\{X_k^{(n+1)} = x\}$ and we emphasize that $X_0^{(n)}$ is sampled using the distribution μ_n. We may select the initial probability distribution μ_0 supported on $\{1, ..., d\}$ in an arbitrary way.

We transform μ_n into a more familiar stochastic approximation form by writing

$$\mu_{n+1}(x) = \mu_n(x) + \frac{1}{n+1}\left(\frac{\sum_{k=0}^{\tau^{(n+1)}-1}\left(I\left(X_k^{(n+1)} = x \mid X_0^{(n)} \sim \mu_n\right) - \mu_n(x)\right)}{\left(\sum_{j=0}^{n+1} \tau^{(j)}\right)/(n+1)}\right).$$

Compared to the standard form in (1) we recognize that $\varepsilon_n = 1/(n+1)$, however, if we attempt to make a direct translation into (1) we see that the denominator is a bit problematic because its conditional expectation (given the whole history of the algorithm up to the end of the n-th iteration) is not only a function of μ_n. To address this issue, we add another variable, T_n, leading to the recursions (assuming $T_0 = 0$)

$$T_{n+1} = T_n + \frac{1}{n+2}\left(\tau^{(n+1)} - T_n\right) = \frac{1}{n+1}\sum_{j=0}^{n} \tau^{(j)}, \quad (3)$$

$$\mu_{n+1}(x) = \mu_n(x) + \frac{1}{n+1}\left(\frac{\sum_{k=0}^{\tau^{(n+1)}-1}\left(I\left(X_k^{(n+1)} = x \mid X_0^{(n)} \sim \mu_n\right) - \mu_n(x)\right)}{T_n + \tau^{(n+1)}/(n+1)}\right).$$

In order to provide a more succinct notation let us define

$$Y_n(\mu', T)(x) := \frac{\sum_{k=0}^{\tau-1}\left(I\left(X_k = x \mid X_0 \sim \mu\right) - \mu(x)\right)}{T + \tau/(n+1)}, \quad (4)$$

$$Z(\mu', T) := (\tau - T),$$

where $\{X_l : l \geq 0\}$ denotes a generic Markov chain with transition matrix P, X_0 is distributed according to μ (supported on $\{1, ..., n\}$), and τ corresponds to the first hitting time to 0 of the chain $\{X_l : l \geq 0\}$. We also write

$$Y(\mu', T)(x) := \frac{\sum_{k=0}^{\tau-1}\left(I\left(X_k = x \mid X_0 \sim \mu\right) - \mu(x)\right)}{T}.$$

Note that Y is time homogeneous where as Y_n is not. Then, we can rewrite the stochastic approximation recursion in distribution via

$$\mu_{n+1}(x) = \mu_n(x) + \frac{1}{n+1} Y_n(\mu_n, T_n)(x),$$
$$T_{n+1} = T_n + \frac{1}{n+2} Z(\mu_n, T_n).$$

If we let $\theta_n = (\mu_n', T_n)$ we now have a setting very close to that described in (1), except for the fact that $g(\cdot)$ is time homogeneous (i.e. of the form $g_n(\cdot)$).

We have the following remarks:

- As noted earlier, the term Y_n is not time homogeneous because of the presence of the term $\tau/(n+1)$. It is not difficult to argue (as we shall do in Lemma 5.1) that such term is asymptotically negligible because $\tau^{(n)} = O(\log(n))$ almost surely.
- Note that each μ_n during the course of the algorithm is a probability vector, that is, $\mu_n \in H := \{x \in \mathbb{R}_+^d : \mathbf{e}'x = 1\}$. So, the boundedness requirement in Theorem 5.2.1 of [19] holds automatically at least for μ_n although we shall need to argue boundedness for the coordinate T_n.
- A similar algorithm can be defined for continuous-time Markov chains by keeping track of the amount of time spent in each transient state. The details are omitted for the purpose of saving space, but the description is given in [27].

3.2.1. *Convergence Result: Consistency and CLT* We now state the main result of this section.

Theorem 3.1. *Suppose that μ_0 is any probability vector supported on $\{1, ..., n\}$ and pick $T_0 \geq 1$. Then $(\mu_n', T_n) \to \theta_* := (\mu_*', 1/(1-\lambda_*))$ with probability one, where the left principal eigenvector μ_*' of Q is normalized so that $\mu_*'\mathbf{e} = 1$. Finally, if $\bar{\lambda}$ is any non-principal eigenvalue of Q (i.e. $\bar{\lambda} \neq \lambda_*$) and*

$$\mathrm{Re}\left(\frac{1}{1-\bar{\lambda}}\right) < \frac{1}{2}\left(\frac{1}{1-\lambda_*}\right), \tag{5}$$

then

$$n^{1/2}(\mu_n - \mu_*) \Rightarrow N(0, V_0),$$

for some V_0, explicitly characterized by equation (20).

Sketch of Proof of Theorem 3.1: The full proof is given in Sections 5.1 and 5.2, but we outline the main idea here. The technique uses the ODE method (Theorem 5.2.1 in [19]), which requires us to examine the asymptotic behavior of the following coupled dynamic system:

$$\dot{\mu}(t) = \frac{1}{T(t)} E\left(\sum_{k=0}^{\tau-1} \left(I(X_k = \cdot \mid X_0 \sim \mu(t)) - \tau\mu(t)\right)\right) = \frac{1}{T(t)}\left(\mu(t)' R - (\mu(t)' R\mathbf{e})\mu(t)'\right)',$$
$$\dot{T}(t) = E(\tau \mid X_0 \sim \mu(t)) - T(t) = \mu(t)' R\mathbf{e} - T(t),$$

where $R = (I - Q)^{-1}$. In Section 5.1 we are able to show using Duhamel's principle that for a given initial position in the probability simplex H, the solution to a suitably reduced dynamical system (obtained by ignoring $\dot{T}(t)$ and assuming that $T(t) = 1$ in the evolution of $\mu(t)$) exists and converges as $t \to \infty$ to its stationary point. This stationary point is the unique solution to the eigenvalue problem $\mu_*' R = \rho_* \mu_*'$, $\mu_*' \mathbf{e} = 1$, and $\mu_* \geq 0$, where $\rho_* = 1/(1 - \lambda_*)$. The uniqueness of the solution of this eigenvalue problem follows from Perron-Frobenius' theorem. The complete dynamical system (given above for $\mu(t)$ and $T(t)$) is a time change of the reduced one, so we can connect them via a simple transformation.

Thus, applying Theorem 5.2.1 from [19] we can conclude that μ_n converges to the quasi-stationary distribution for all initial configurations $(\mu_0, T_0) \in H \times [1, \infty)$.

For the CLT we invoke Theorem 10.2.1 in [19]. Because the recursion in (3) uses step size $\varepsilon_n = 1/(n+1)$, we need to verify that the Jacobian matrix of the ODE vector field, evaluated at the stability point, has spectral radius less than $-1/2$. As we show in Section 5.2 this is equivalent to requiring (5). The expression for V_0 is extracted from the variance of an associated Ornstein-Uhlenbeck process as in p. 332 of [19]. □

4. Variations on the Algorithm with Improved Rate of Convergence

In this section we study what occurs to the rate of convergence when the sufficient conditions for the CLT are not met. We shall study a simple example consisting of two states.

4.1. Counterexample to Square Root Convergence

Consider the Markov chain with states $\{0, 1, 2\}$, the state 0 is absorbing and the matrix Q satisfies

$$Q = \begin{pmatrix} (1-\varepsilon)/2 & (1-\varepsilon)/2 \\ (1-\varepsilon)/2 & (1-\varepsilon)/2 \end{pmatrix}. \tag{6}$$

By symmetry the recursion that we analyze, namely (2), can be tracked by a simple process, $\{\bar{X}_m : m \geq 0\}$, which we describe now. Assume that the distribution of $\bar{X}_0$ is given. At step m, the value of $\bar{X}_m$ is decided according to a Bernoulli trial which call the *type*. The type is Bernoulli with success parameter equal to $1 - \varepsilon$.

If the type is a success, we sample a second Bernoulli trial with probability 1/2 of success, if the second trial is successful we let $\bar{X}_m = 1$, if it is a failure we let $\bar{X}_m = 2$.

If the type is a failure (which occurs with probability ε), the we sample state 1 or 2 according to the empirical measure of $\{\bar{X}_k : 0 \leq k \leq m-1\}$.

Let T_n be time at which the n-th failure type occurs. Then we can have that recursion (2) is equivalent to studying

$$\mu_n(x) = \frac{1}{T_n} \sum_{k=0}^{T_n} I\left(\bar{X}_k = x\right).$$

The process $\left\{\bar{X}_m : m \geq 0\right\}$ is known as a self-interacting Markov chain, see [11]. Inequality (5) in Theorem 3.1 applied to this case corresponds to requiring $\varepsilon < 1/2$. Of course, $\mu = (1/2, 1/2)$. Reference [11] is applicable to this example and shows that if $f \neq 0$ there exists $\delta > 0$ such that for $n \geq 1$

$$\delta \frac{1}{n^{2(1-\varepsilon)}} \leq E\left(\left(\mu_n' f - \mu' f\right)^2\right) \leq \delta^{-1} \frac{1}{n^{2(1-\varepsilon)}}.$$

This results indicates that the rate of convergence is not $O(n^{-1/2})$ but rather $O\left(n^{-(1-\varepsilon)}\right)$.

4.2. Doeblinization

The previous example shows that the rate of convergence can deteriorate substantially if (5) does not hold. We now argue that in fact that there are natural algorithmic "tricks" that one might attempt to use and which are likely to induce a violation of (5).

Note that the expected time to absorption starting from the quasi-stationary distribution satisfies $E(\tau \mid X_0 \sim \mu_*) = (1-\lambda_*)^{-1}$. If this expected time is large the iterations of the algorithm will tend to be long. In order to shorten the length of the iterations one might "Doeblinize" the chain by multiplying Q by a constant $\alpha \in (0,1)$. (The term Doeblinization is adopted from the name of the French-German probabilist Wolfgang Döeblin, who studied Markov chains satisfying certain minorization condition which is satisfied obtained when multiplying Q by α as indicated earlier.) This operation, of course, does not change the principal eigenvector of Q, but it does change all the eigenvalues by the same proportion. Because of the non-linearity of the $1/(1-\lambda_*)$ as a function of λ_* and its presence in (5), we conclude choosing $\alpha > 0$ too small might result in a significant deterioration in rate of convergence (despite the gain in speed at each iteration). This observation further motivates the need for a technique that allows to obtain a CLT with a square-root convergence which can be guaranteed regardless of the eigenvalues of Q.

Doeblinization can also be done for continuous time Markov chain by subtracting αI from the substochastic rate matrix associated with the transient states $\{1, ..., n\}$.

4.2.1. *Projection and Averaging* Now that the method is under the stochastic approximation umbrella, we can modify the algorithm in order to change the step-size by using its projection variant, given by the recursion

$$\bar{\mu}_{n+1} = \Pi_H \left(\bar{\mu}_n + \varepsilon_n \left(\sum_{k=0}^{\tau^{(n+1)}-1} \left(I\left(X_k^{(n+1)} = \cdot \mid X_0^{(n)} \sim \mu_n\right) - \bar{\mu}_n(\cdot)\right)\right)\right), \tag{7}$$

where Π_H denotes the L_2-projection into the probability simplex H. We still require $\sum \varepsilon_n = \infty$ and $\sum \varepsilon_n^2 < \infty$. In practice we only need to perform a small number of projections. The vector inside the projection operator always has components which add up to one. So, projection is only needed when any of the components of the vector inside $\Pi_H(\cdot)$ is negative. By conveniently splitting the recursion defining the iterates in (7) we can gain insight into when one or more components of the updated vector $\bar{\mu}_n$ can become negative. In particular, we can write

$$\bar{\mu}_n + \varepsilon_n \left(\sum_{k=0}^{\tau^{(n+1)}-1} \left(I\left(X_k^{(n+1)} = \cdot \mid X_0^{(n)} \sim \mu_n \right) - \bar{\mu}_n(\cdot) \right) \right)$$
$$= \bar{\mu}_n \left(1 - \varepsilon_n \tau^{(n+1)}\right) + \varepsilon_n \sum_{k=0}^{\tau^{(n+1)}-1} I\left(X_k^{(n+1)} = \cdot \mid X_0^{(n)} \sim \bar{\mu}_n \right).$$

So, for a component to become negative, it is necessary that $\tau^{(n+1)} > \varepsilon_n^{-1}$. It is not difficult to argue, as we shall do in Lemma 5.1 part iii), that there exists $\delta > 0$ such that $\tau^{(n+1)} > \delta \log(n)$ for only finitely many values of $n \geq 1$ with probability one, thus $\tau^{(n+1)} > \varepsilon_n^{-1}$ occurs only finitely many times if $\varepsilon_n = O(n^{-\alpha})$ for $\alpha > 0$. Moreover, it is quite easy to perform the L_2 projection into a probability simplex. In particular,

$$\bar{\mu}_{n+1} = \left(\bar{\mu}_n \left(1 - \varepsilon_n \tau^{(n+1)}\right) + \varepsilon_n \sum_{k=0}^{\tau^{(n+1)}-1} I\left(X_k^{(n+1)} = \cdot \mid X_0^{(n)} \sim \bar{\mu}_n \right) - u_{n+1}\mathbf{e} \right)_+ ,$$

where $u_{n+1} > 0$ is the unique constant such that $\bar{\mu}_{n+1}'\mathbf{e} = 1$ (see [5]).

The advantage of the projection version is that we are free to choose slower step sizes so that we can weaken the condition for the required CLT to hold. In particular, when $\varepsilon_n = n^{-\alpha}$ and $\alpha \in (1/2, 1)$ we always obtain a $\varepsilon_n^{-1/2}$-CLT. We summarize this observation in the following result proved in Section 6.

Proposition 4.1. *If $\varepsilon_n = n^{-\alpha}$ for $\alpha \in (1/2, 1)$ we have that*

$$\varepsilon_n^{-1/2} (\bar{\mu}_n - \mu_*) \Rightarrow N(0, V_1),$$

where V_1 can be characterized via (21).

The Polyak-Ruppert averaging technique, [25], can be applied jointly with the projection algorithm to ensure "square root convergence", regardless of whether (5) holds or not as the next theorem shows. Its proof takes advantage of the analysis behind Proposition 4.1 and Theorem 3.1 and it is given in 6 based on [25].

Theorem 4.1. *Suppose that μ_0 is any probability vector supported on $\{1, ..., n\}$ and pick $T_0 \geq 1$. Selecting $\varepsilon_n = n^{-\alpha}$ for $\alpha \in (1/2, 1)$, let*

$$v_n = \frac{1}{n} \sum_{k=1}^{n} \bar{\mu}_k.$$

Then,

$$n^{1/2} (v_n - \mu_*) \Rightarrow N\left(0, \bar{V}_1\right),$$

where $\bar{V}_1$ is given in equation (25).

We can apply Theorem 4.1 in the estimation of quasi-stationary expectations of the form $E\left(s\left(X\right)|X\sim\mu_*\right)=\mu_*'s$, using the estimator $\mu_n's$ (note that we are encoding the function $s\left(\cdot\right)$ as a column vector). As a consequence of our CLT, we have the following corollary.

Corollary 4.1. *Under the notations defined in Theorem 3.1, we have that*

$$n^{1/2}\left(v_n's-\mu_*'s\right)\Rightarrow N(0,\sigma_s^2),$$

where $\sigma_f^2=s'\bar{V}_1 s$.

4.2.2. *Estimating the Asymptotic Variance* The next result (proved in Section 7) indicates how to estimate V_1 using the outcomes of the improved algorithm, which we ultimately advocate using.

Proposition 4.2. *Let* v_n *be defined as in Theorem 4.1. For* $\varepsilon_n=n^{-\alpha}$ *with* $\alpha\in(1/2,1)$, *set* $n_k=\left[k^{\beta/\alpha}\right]$ *and* $N_n=\left[n^{\alpha/\beta}\right]$, *where* $\beta\in(\alpha,1)$. *Then,*

$$\frac{1}{N_n}\sum_{k=0}^{N_n}\varepsilon_{n_k}^{-1}\left(\bar{\mu}_{n_k}-v_{n_k}\right)\left(\bar{\mu}_{n_k}-v_{n_k}\right)'\longrightarrow\bar{V}_1,$$

as $n\to\infty$ *in probability.*

In the context of Corollary 4.1, we can use Proposition 4.2 noting that

$$\frac{1}{N_n}\sum_{k=0}^{N_n}\varepsilon_{n_k}^{-1}s'\left(\bar{\mu}_{n_k}-v_{n_k}\right)\left(\bar{\mu}_{n_k}-v_{n_k}\right)'s=\frac{1}{N_n}\sum_{k=0}^{N_n}\varepsilon_{n_k}^{-1}\left(\left(\bar{\mu}_{n_k}-v_{n_k}\right)'s\right)^2.$$

In turn, we can recursively compute $\bar{\mu}_n's-v_n's$ the previous estimator can be cheaply computed because the n-th update of the associated recursions takes $O(E(\tau^{(n+1)}|X_0^{n+1}\sim\mu_n))$ function evaluations, in expectation and we saw that the "Doeblinization" procedure allows us to impose a uniform (in n) upper bound on $E\left(\tau^{(n+1)}|X_0^{n+1}\sim\mu_n\right)$.

5. Proofs of Main Results

5.1. Proof of Theorem 3.1: Convergence

We first restate a series of assumptions and notations that are used in Theorem 5.2.1 from [19]. We adopt the abstract form of the recursion $\theta_{n+1}=\theta_n+\varepsilon_n W_n$. In our setting $\theta_n=\left(\mu_n',T_n\right)$ and $W_n=\left(Y_n\left(\theta_n\right),Z\left(\theta_n\right)\right)$ as defined in (4). Recall that $\mathcal{F}_n$ is the σ-field generated by the iterates of the algorithm, namely, $\mathcal{F}_n=\sigma(\theta_0,\theta_i,W_{i-1}:1\leq i\leq n)$.

For the almost sure convergence of μ_n to μ_* we must verify the following conditions.

1. $\varepsilon_n\to 0$, $\sum\varepsilon_n=\infty$, $\sum\varepsilon_n^2<\infty$. This is immediately satisfied with the choice $\varepsilon_n=1/(n+1)$, as in our case. Moreover, define $t_n=\sum_{j=1}^n\varepsilon_j$ (with $t_0=0$) and let $m\left(s\right)=\max\{n:t_n\leq s\}$.

2. Uniformly bounded variance: $\sup_n E\left\|W_n\right\|_\infty^2 < \infty$. This is shown in Lemma 5.4 below. (We can use any norm, but we choose the norm $\|x\|_\infty = \max_i |x_i|$.)

3. Local averaging condition: Define $g_n(\theta_n) := E(W_n|\mathcal{F}_n)$. The family of functions $\{g_n(\cdot)\}_{n\geq 0}$ must be uniformly equicontinuous, and there must exist a continuous function $g(\cdot)$ such that for each θ, and each $t>0$,

$$\left|\sum_{k=n}^{m(t_n+t)} \varepsilon_k (g_k(\theta) - g(\theta))\right| \longrightarrow 0,$$

almost surely. This local averaging condition is proved in Lemma 5.3 below.

In our setting, we write $\theta = (\mu', T)$ and define $g(\theta) = (f'(\theta), h(\theta))$, where f and h are given via

$$f(\mu', T) = \frac{1}{T} E\left(\sum_{k=0}^{\tau-1} (I(X_k = \cdot) - \mu) \mid X_0 \sim \mu\right) = \frac{1}{T}(\mu' R - (\mu' R\mathbf{e})\mu')', \tag{8}$$

$$h(\mu', T) = E(\tau \mid X_0 \sim \mu) - T = \mu' R\mathbf{e} - T, \tag{9}$$

with $R = (I-Q)^{-1}$. We also define

$$f_n(\mu', T) = E\left(\frac{\sum_{k=0}^{\tau-1}(I(X_k = \cdot) - \mu)}{T + \tau/(n+1)} \mid X_0 \sim \mu\right)$$

and set $g_n(\theta) = (f_n'(\theta), h(\theta))$.

Under conditions 1., 2., and 3., Theorem 5.2.1 from [19] indicates that if the ODE

$$\dot{\theta}(t) = g(\theta(t))$$

has an attractor (asymptotically stable point) in some domain $\mathcal{D}$ and the sequence $\{\theta_n\}$ visits a compact subset within the domain $\mathcal{D}$ infinitely often with probability 1, then θ_n converges to the attractor with probability 1.

In our situation it turns out that the entire probability simplex H is the domain of attraction for an attractor which is precisely the quasi-stationary vector. So we just need to be concerned with $\{T_n\}$. We will compute the functions $\{g_n(\cdot)\}$, verify condition 3., then the uniformly bounded variance condition 2., and finally the asymptotic stability behavior of the ODE. We will show that in fact $\{T_n\}$ stays within a compact set throughout the course of the algorithm and the uniform continuity of the functions $\{g_n(\cdot)\}$ holds for every compact set in $H \times [1, \infty)$. First, however, we obtain a result that allows us to control the sequence $\{\tau^{(n)}\}$.

5.1.1. *Auxiliary Results* Define $\bar{\tau}(x)$ to be a random variable with the distribution of the first passage time to the absorbing state, 0, given that the initial condition of the chain is the transient state $x \in \{1, ..., d\}$. Suppose that the random variables $\{\bar{\tau}(j) : 1 \leq j \leq d\}$ are all independent. Then, let $\bar{\tau} = \max\{\bar{\tau}(j) : 1 \leq j \leq d\}$. We have the following simple but useful result.

Lemma 5.1. *The following claims hold: i) $\tau^{(n+1)}$ is stochastically bounded by $\bar{\tau}$, ii) there exists $\delta > 0$ such that $E\exp(\delta\bar{\tau}) < \infty$, iii) $P(\tau^{(n+1)} > \log(n)$ i.o.$) = 0$, iv) almost surely we have that $1 \leq \overline{\lim} \sum_{k=1}^{n} \tau^{(k)}/n \leq E(\widetilde{\tau}) < \infty$.*

Proof of Lemma 5.1: The proof is almost immediate part i) follows regardless of any assumption, for parts ii) to iv) we need $Q^n \to 0$ as $n \to \infty$ because this ensures that $\bar{\tau}(x)$ has a finite moment generating function in a neighborhood of the origin. □

We also have a useful expression for $f_n(\mu', T)$. Define $v(x, s) := E(\exp(-s\tau) | X_0 = x)$.

Lemma 5.2. *For each $(\mu', T) \in H \times (0, \infty)$, the x-th component of $f_n(\mu', T)$, namely $f_n(\mu', T)(x)$, is equal to*

$$\int_0^\infty e^{-Tu} \left[v\left(x, \frac{u}{n+1}\right) \mu' \left(I - e^{-\frac{u}{n+1}} Q\right)^{-1} e_x - \left(\mu' \left(I - e^{-\frac{u}{n+1}} Q\right)^{-1} v\left(\cdot, \frac{u}{n+1}\right)\right) \mu(x) \right] du,$$

where e_x denotes the vector which has 1 in the x-th coordinate and zeroes elsewhere.

Proof of Lemma 5.2: First note that

$$\frac{1}{T + \tau/(n+1)} = \int_0^\infty \exp(-(T + \tau/(n+1))u)\, du.$$

Then we have (applying Fubini's theorem since $\int_0^\infty E\left(e^{-Tu}\tau | X_0 \sim \mu\right) du < \infty$),

$$f_n(\mu', T)(x) = \int_0^\infty E\left(e^{-\left(T + \frac{\tau}{n+1}\right)u} \sum_{k=0}^{\tau-1} (I(X_k = x) - \mu(x)) \mid X_0 \sim \mu\right) du.$$

Again, another application of Fubini's theorem (also valid because $E(\tau | X_0 \sim \mu) < \infty$) yields that the previous expression equals

$$\begin{aligned}
&\int_0^\infty E\left(e^{-\left(T + \frac{\tau}{n+1}\right)u} \sum_{k=0}^{\infty} I(\tau > k)(I(X_k = x) - \mu(x)) \mid X_0 \sim \mu\right) du \\
&= \int_0^\infty e^{-Tu} \sum_{k=0}^{\infty} e^{-\frac{ku}{n+1}} E\left(e^{-u\frac{(\tau-k)}{n+1}} I(\tau > k)(I(X_k = x) - \mu(x)) \mid X_0 \sim \mu\right) du \\
&= \int_0^\infty e^{-Tu} \sum_{k=0}^{\infty} e^{-\frac{ku}{n+1}} E\left(v\left(X_k, \frac{u}{n+1}\right)(I(X_k = x) - \mu(x)) \mid X_0 \sim \mu\right) du \\
&= \int_0^\infty e^{-Tu} \sum_{k=0}^{\infty} e^{-\frac{ku}{n+1}} \left[v\left(x, \frac{u}{n+1}\right) \mu' Q^k e_x - \left(\mu' Q^k v\left(\cdot, \frac{u}{n+1}\right)\right) \mu(x)\right] du \\
&= \int_0^\infty e^{-Tu} \left[v\left(x, \frac{u}{n+1}\right) \mu' \left(I - e^{-\frac{u}{n+1}} Q\right)^{-1} e_x - \left(\mu' \left(I - e^{-\frac{u}{n+1}} Q\right)^{-1} v\left(\cdot, \frac{u}{n+1}\right)\right) \mu(x)\right] du.
\end{aligned}$$

□

5.1.2. *Local Averaging and Uniformly Bounded Variance* We first verify the uniformly bounded variance condition

Lemma 5.3. *We have that $\{g_n(\cdot)\}$ is a sequence of uniformly equicontinuous functions on $H \times [1,\infty)$ and we have that for each $t > 0$,*

$$\lim_{n\to\infty}\left|\sum_{k=n}^{m(t_n+t)} \varepsilon_k (g_k(\mu', T) - g(\mu', T))\right| = 0.$$

Proof of Lemma 5.3: Clearly, we have that

$$\begin{aligned}&|f_k(\mu', T)(x) - f(\mu', T)(x)| \\ &\leq E\left(\sum_{k=0}^{\tau-1} \frac{|(I(X_k = x) - \mu(x))|}{T} \frac{\tau}{(T(k+1)+\tau)} \mid X_0 \sim \mu\right) \leq E\left(\frac{\bar{\tau}^2}{k+1}\right).\end{aligned}$$

Therefore,

$$\sum_{k=n}^{m(t_n+t)} \varepsilon_k |(f_k(\mu', T) - f(\mu', T))| \leq E\left(\frac{\bar{\tau}^2}{n+1}\right) \sum_{k=n}^{m(t_n+t)} \varepsilon_k \leq tE\left(\frac{\bar{\tau}^2}{n+1}\right) \to 0$$

as $n \to \infty$. Finally, we need to argue that $g_n(\cdot)$ is uniformly equicontinuous on compact sets in $H \times [1,\infty)$. This follows easily by noting from the expression obtained in Lemma 5.2 that the Jacobian $(Df_n)(\mu', T)$ is uniformly bounded for a neighborhood around any point $(\mu', T) \in H \times [1,\infty)$.

The coordinate of g_n corresponding to h does not depend on n and thus the result follows immediately in this case. □

Now we turn our attention to condition 2., namely, uniformly bounded variance.

Lemma 5.4. *We have that $\sup_n E\|W_n\|_\infty^2 < \infty$.*

Proof of Lemma 5.4: Clearly, we have that $\|Y_n(\mu_n, T_n)\|_\infty \leq \tau^{(n+1)}$ and therefore Lemma 5.1, parts i) and ii) yield that

$$E\|Y_n(\mu_n, T_n)\|_\infty^2 \leq E\bar{\tau}^2 < \infty. \tag{10}$$

We also have that $|Z(\mu_n, T_n)| \leq \tau^{(n+1)}$ and therefore a similar bound to (10) applies, thus concluding the proof. □

5.1.3. *Stability of the Dynamical System and Final Convergence Argument* The dynamical system of interest, namely $\dot{\theta}(t) = g(\theta(t))$ takes the form

$$\begin{aligned}\dot{\mu}(t)' &= f(\mu(t), T(t)) = \frac{1}{T(t)}\left(\mu(t)' R - (\mu(t)' R\mathbf{e})\mu'(t)\right) \qquad (11) \\ \dot{T}(t) &= h(\mu(t), T(t)) = \mu(t)' R\mathbf{e} - T(t).\end{aligned}$$

In the proof of Theorem 5.2.1 in [19] it is shown that any converging subsequence of the suitably normalized iterates converges to a function $\theta(\cdot,\omega)$ which is a solution to the ODE $\acute{\theta}(t) = g(\theta(t))$. We only need to show that these solutions converge as $t \to \infty$ to $(\mu_*, 1/(1-\lambda_*))$. We will actually show that all solutions of a suitably reduced ODE (starting from points in H) converge to the quasi-stationary distribution and then we will show how to map the solutions $\theta(\cdot,\omega)$ for the full system (11) into solutions to the reduced ODE. We can invoke then Theorem 5.2.1 in [19] after arguing that $\{(\mu_n, T_n)\}$ visits a compact set of $H \times [1,\infty)$ infinitely often almost sure; actually we will show that $\{(\mu_n, T_n)\}$ will stay inside a compact set eventually.

We now define the reduced ODE as follows

$$\begin{aligned} \dot{v}(t) &= \left(v(t)' R - (v(t)' R\mathbf{e}) v'(t)\right)' \\ v(0) &= \mu_0 \in H. \end{aligned} \tag{12}$$

Note that the gradient of the vector field in (12) is continuously differentiable in H, therefore $v(\cdot)$ has a unique solution given $v(0) \in H$.

Suppose that $T_0 \geq 1$ and let $(\mu(\cdot), T(\cdot))$ be a solution to (11) obtained by the subsequence procedure in Theorem 5.2.1 from [19], then define $\Gamma(t) = \int_0^t (1/T(s))\, ds$. It follows by formal differentiation that $v(t) = \mu(\Gamma^{-1}(t))$ solves (12). The following result ensures regularity properties of $\Gamma(\cdot)$.

Lemma 5.5. *$\Gamma(t) > 0$, $\Gamma(\cdot)$ is strictly increasing, and $\Gamma(t) \to \infty$ as $t \to \infty$.*

Proof of Lemma 5.5: Clearly, we have that $T(s) \geq 1$ so $\Gamma(\cdot)$ is strictly increasing and non-negative. Now, suppose that $\Gamma(\infty) < \infty$, this implies that

$$\begin{aligned} T(t) &= T_0 \exp\left(\int_0^t \frac{E(\tau \mid X_0 \sim \mu(s))}{T(s)} ds - t\right) \\ &\leq T_0 \exp\left(E(\bar{\tau}) \int_0^t \frac{ds}{T(s)} - t\right) \leq T_0 \exp\left(E(\bar{\tau})\, \Gamma(\infty)\right) \in (0,\infty). \end{aligned}$$

This bound would imply that there exists $\delta > 0$ so that $1/T(t) > \delta$ and obtaining a contradiction to the assumption that $\Gamma(\infty) < \infty$. Thus we must have that $\Gamma(\infty) = \infty$. □

Lemma 5.6. *Any solution to the reduced ODE in (12) (regardless of $v(0) \in H$) converges to the quasi-stationary distribution μ as $t \to \infty$.*

Proof of Lemma 5.6: By applying inner product with the vector $\mathbf{e}$ we can see that the entire trajectory stays in H. We let $R = (I - Q)^{-1} = \sum_{k=0}^{\infty} Q^k$ which has only non-negative entries (actually, the entries are strictly positive if Q is irreducible). By Duhamel's principle all the solutions $v(\cdot)$ can be represented by

$$v(t)' = v(0)' \exp\left(Rt - \int_0^t \left(v(s)' R\mathbf{e}\right) ds\right). \tag{13}$$

Because R has non-negative entries we have that $\exp(Rt)$ has non-negative entries and thus $v(t) \geq 0$ (in fact the entries are strictly positive if Q is irreducible). Rearranging equation (13) we arrive at

$$v(t)' \exp\left(\int_0^t \left(v(s)' R\mathbf{e}\right) ds - \lambda_R t\right) = v(0)' \exp\left(Rt - \lambda_R t\right), \tag{14}$$

where $\lambda_R > 0$ is the principal eigenvalue of R. The matrix $\exp(R) / \exp(\lambda_R)$ is a matrix with strictly positive entries and it has spectral radius equal to one. By the Perron-Frobenius theorem, see [16], we have that there exists a strictly positive vector w such that

$$v(0)' \exp\left(Rn - \lambda_R n\right) \to w'$$

as $n \to \infty$, where w is a principal eigenvector of $R - \lambda_R I$. The convergence holds also along real numbers t (not only natural numbers n) by virtue of a continuity argument noting that $(R - \lambda_R I)/m$ has the same eigenvectors regardless of the value of $m > 0$.

Now, take the inner product with $\mathbf{e}$ in both sides of equation (14) to obtain (because $v(t) \in H$) that

$$\exp\left(\int_0^t \left(v(s)' R\mathbf{e}\right) ds - \lambda_R t\right) \to \gamma := w'\mathbf{e} \in (0, \infty).$$

Finally, rewrite (13) as

$$v(t)' = v(0)' \exp\left(Rt - \lambda_R t\right) \exp\left(-\left(\int_0^t \left(v(s)' R\mathbf{e}\right) ds - \lambda_R t\right)\right) \to \frac{w'}{\gamma}.$$

The fact that finishes the proof of the lemma is that the Perron-Frobenius eigenvectors of R and Q are identical, so we see that $v(t)$ converges to the quasi-stationary distribution as $t \to \infty$. □

Now we are ready to conclude the consistency portion of Theorem 3.1 by invoking Theorem 5.2.1 in [19] together with the following proposition.

Proposition 5.1. *Any subsequence solution (obtained as in Theorem 5.2.1 in [19]) of the system (11) satisfies that $\mu_0 \in H$ and $T_0 \geq 1$ and it is such that $\mu(t) \to \mu_*$ and*

$$T(t) \to 1/(1 - \lambda_*) = E\left(\tau \mid X_0 \sim \mu_*\right)$$

as $t \to \infty$. The sequence $\{(\mu_n', T_n)\}$ stays in a compact set of the attractor domain $H \times [1, \infty)$ eventually. Therefore, $\mu_n \to \mu_*$ and $T_n \to 1/(1 - \lambda_*)$ with probability one.

Proof of Proposition 5.1: We have that $\mu_n \in H$ and $T_n \geq 1$ because T_n is the average of the $\tau^{(n)}$'s which are greater than one, so the subsequence procedure in Theorem 5.2.1 in [19] produces trajectories that lie in $H \times [1, \infty)$ and which are solutions to (11). Now, we have noted that $\Gamma(\cdot)$ is non-negative and strictly increasing, according to Lemma 5.5 and thus Lemma 5.6 implies that $v(t) = \mu\left(\Gamma^{-1}(t)\right) \to \mu_*$. Moreover, Lemma 5.5 indicates that $\Gamma^{-1}(t) \to \infty$, therefore we have that $\mu(t) \to \mu_*$ as $t \to \infty$.

Now, observe that

$$T(t) = \frac{\int_0^t E(\tau \mid X_0 \sim \mu(s)) \exp(s)\, ds + T_0}{\exp(t)}.$$

Because $E(\tau \mid X_0 \sim \mu(t)) \to E(\tau \mid X_0 \sim \mu_*)$, we can use L'Hopital's rule and conclude that

$$\lim_{t\to\infty} T(t) = \lim_{t\to\infty} \frac{E(\tau \mid X_0 \sim \mu(t)) \exp(t)}{\exp(t)} = E(\tau \mid X_0 \sim \mu_*) = 1/(1-\lambda_*).$$

The fact that $\{(\mu'_n, T_n)\}$ stays in a compact set of $H \times [1, \infty)$ follows from Lemma 5.1 part iv). □

5.2. Proof of Theorem 1: CLT

In order to prove the CLT portion of Theorem 3.1, we shall invoke Theorem 10.2.1 in [19]; this requires verifying the following conditions.

1. The sequence $\{W_n I(\|\theta_n - \theta_*\| \leq \delta)\}$ is uniformly integrable. (This follows immediately because due to Lemma 5.4 we have that W_n is L_2 bounded.)

2. θ_* is an isolated stable point of the ODE. (This follows from Perron-Frobenius theorem and the analysis in Proposition 5.1.)

3. The following expansion is valid

$$g_n(\theta) = g_n(\theta_*) + (Dg_n)(\theta_*)(\theta - \theta_*) + o(\|\theta - \theta_*\|),$$

   where the error term is uniform in n. (This estimate will be elaborated in Lemma 5.7.)

4. We must have that

$$\lim_{n,m\to\infty} \frac{1}{m^{1/2}} \sum_{k=n}^{n+mt-1} (Dg_k)(\theta_*) = 0$$

   uniformly over t in compact sets. (See Lemma 5.9.)

5. There exists a matrix A such that

$$\lim_{n,m} \sum_{k=n}^{n+m-1} ((Dg_k)(\theta_*) - A) = 0.$$

   (Let $A = (Dg)(\theta_*)$, then this condition will hold true by Lemma 5.10 which shows that $(Dg_n)(\theta_*) \to (Dg)(\theta_*)$.)

6. The matrix A must also be such that $A + I/2$ is Hurwitz (i.e. all its eigenvalues have negative real part). (This corresponds precisely to condition (5) and it will be established in Proposition 5.2.)
7. The sequence $\{(\theta_n - \theta_*)/\varepsilon_n^{1/2}\}$ is tight. (See Lemma 5.12.)
8. Define $\delta M_n = W_n - g_n(\theta_n)$, then there exists $p > 0$ such that
$$\sup_n E \|\delta M_n\|_\infty^{2+p} < \infty,$$
and a non-negative matrix Σ such that
$$E\left(\delta M_n (\delta M_n)' |\mathcal{F}_n\right) \to \Sigma$$
in probability as $n \to \infty$. (This is established in Lemma 5.13.)

5.2.1. *Analysis of the Jacobian: Conditions 3, 4, and 5*

Lemma 5.7. *We have that*
$$\lim_{\|\theta-\theta_*\|_\infty \to 0} \sup_{n\geq 1} \left| \frac{g_n(\theta) - g_n(\theta_*) - (Dg_n)(\theta_*)(\theta - \theta_*)}{\|\theta - \theta_*\|_\infty} \right| = 0.$$

Proof of Lemma 5.7: We consider the analysis only for $f_n(\mu', T)$ because h is a simpler quantity and does not depend on n. The analysis follows as an application of the representation derived in Lemma 5.2. It is easy to justify the interchange of differentiation and integration in the representation given in Lemma 5.2 because the integrand consists of products of a second degree polynomial in μ, the exponential factor $\exp(-uT)$ on the region of interest which is $T \geq 1$, and the term including $v(x, s) \in (0, 1]$. Thus the second derivatives of f_n will be bounded uniformly in n around a neighborhood of the stationary point θ_*. □

Next we turn to condition 4., but first we have an auxiliary result.

Lemma 5.8. *It follows that $g_n(\theta_*) = O(1/n)$.*

Proof of Lemma 5.8: Note that $h(\theta_*) = 0$, so we only focus on $f_n(\theta_*)$. On the other hand, we observed in the proof of Lemma 5.3 that $|f_n(\mu', T)(x) - f(\mu', T)(x)| \leq E(\bar{\tau}^2)/(n + 1)$, but (with $R = (I - Q)^{-1}$),
$$f(\mu'_*, T_*) = \frac{1}{T_*}(\mu'_* R - (\mu'_* R\mathbf{e})\mu'_*)' = \frac{1}{T^*}\left(\frac{1}{1-\lambda}\mu'_* - \frac{1}{1-\lambda}\mu'_*\right)' = 0.$$
Hence the Lemma 5.8 follows. □

Lemma 5.9. *We have that*
$$\lim_{n,m\to\infty} \frac{1}{m^{1/2}} \sum_{k=n}^{n+mt-1} (Dg_k)(\theta_*) = 0,$$
uniformly over compact sets in t.

Proof of Lemma 5.9: Using Lemma 5.8 we have that there exists a constant c (independent of n) such that

$$\frac{1}{m^{1/2}}\sum_{k=n}^{n+mt-1}\left\|\left(Dg_k\right)\left(\theta_*\right)\right\|_\infty \leq \frac{c}{m^{1/2}}\log\left(1+\frac{mt}{n}\right).$$

Changing the variables via the transformation $u = mt/n$ we have that

$$\frac{1}{m^{1/2}}\sum_{k=n}^{n+mt-1}\left\|\left(Dg_k\right)\left(\theta_*\right)\right\|_\infty \leq \frac{ct^{1/2}}{n^{1/2}}\times\sup_{u\geq 0}\frac{\log\left(1+u\right)}{u^{1/2}}.$$

We have that $\sup_{u\geq 0}\log\left(1+u\right)/u^{1/2} < \infty$ and therefore, we can send $n\to\infty$ in the right hand side to conclude the statement of Lemma 5.9. □

Lemma 5.10.

$$\left(Dg_n\right)\left(\theta_*\right)\to\left(Dg\right)\left(\theta_*\right)$$

as $n\to\infty$.

Proof of Lemma 5.10: Once again, it suffices to concentrate on f_n. Because of Lemma 5.7, we know that

$$f_n\left(\theta\right) = f_n\left(\theta_*\right)+\left(Df_n\right)\left(\theta_*\right)\left(\theta-\theta_*\right)+o\left(\theta-\theta_*\right).$$

Taking the limit as $n\to\infty$, we arrive at

$$f\left(\theta\right) = \lim_{n\to\infty}\left(Df_n\right)\left(\theta_*\right)\left(\theta-\theta_*\right)+o\left(\theta-\theta_*\right).$$

Expanding the left hand side, we have that

$$f\left(\theta\right) = \left(Df\right)\left(\theta_*\right)\left(\theta-\theta_*\right)+o\left(\theta-\theta_*\right).$$

Matching these terms and noting that $\theta-\theta_*$ can have any direction, we conclude the result of the lemma. □

5.2.2. *The Hurwitz Property: Condition 6*

Proposition 5.2. *Let $A = \left(Dg\right)\left(\theta_*\right)$, then $A+I/2$ is Hurwitz assuming that the eigenvalues of Q satisfy condition (5).*

Proof of Proposition 5.2: Recall that $g\left(\mu',T\right) = \left(f'\left(\mu',T\right),h\left(\mu',T\right)\right)$ and expressions (8) and (9). Letting $B = R' = \left(I-Q'\right)^{-1}$ we have that the Jacobians are given by (using D_μ and D_T to denote the derivatives with respect to μ and T respectively),

$$\begin{aligned}
D_\mu f\left(\mu',T\right) &= \frac{1}{T}\left[B-\left(\mu' B\mathbf{e}\right)I-\mu\mathbf{e}' B\right],\\
D_T f\left(\mu',T\right) &= -\frac{1}{T^2}\left[B\mu-\left(\mu' B\mathbf{e}\right)\mu\right],\\
D_\mu h\left(\mu',T\right) &= \mathbf{e}' B,\ \text{and}\ \ D_T h\left(\mu',T\right) = -1.
\end{aligned}$$

We consider the stationary point and note that $\mu_*' B\mathbf{e} = T_*$. Then, define

$$J := D_\mu f\left(\mu_*', T_*\right) = \frac{1}{T_*}\left[B - T_* I - \mu_* \mathbf{e}' B\right]. \tag{15}$$

Also, note that $D_T f\left(\mu_*', T_*\right) = 0$. We now establish a one-to-one correspondence between the eigenvectors of J and the eigenvectors of B. The overall Jacobian in block form would take the form

$$A = \left[\begin{array}{cc} J & 0 \\ \mathbf{e}' B & -1 \end{array}\right].$$

This matrix has the same eigenvalues as J, with the addition of the eigenvalue -1 which has no effect on the Hurwitz property. Hence we only need to ensure that $J + I/2$ is Hurwitz.

Let y be any vector such that $Jy = \lambda_J y$ and such that y is linearly independent of μ_*. Note that $Jy = \lambda_J y$ is equivalent to the equation

$$By = T_* \lambda_J y + T_* y + \left(\mathbf{e}' By\right) \mu_*,$$

and therefore if we let $x = y + r\mu_*$, for some r to be characterized momentarily, we have that

$$Bx = By + rT_*\mu_* = T_*(\lambda_J + 1)y + \left(\mathbf{e}' By + rT_*\right) \mu_*.$$

So, the value of r that would make x an eigenvector of B is such that $rT_*\lambda_J = \mathbf{e}' By$. Since $T_* > 0$ the existence of r is guaranteed if $\lambda_J \neq 0$ and the corresponding eigenvalue for B would be $\lambda_B = T_*\left(1 + \lambda_J\right)$. On the other hand, if y is a multiple of μ_*, its eigenvalue for the matrix J equals -1 (the eigenvalue for B is, of course, T_*). In Lemma 5.11 below we will argue that λ_J cannot be zero, so the argument just given shows that every eigenvector of J is an eigenvector of B.

Conversely, given any vector z, such that $Bz = \lambda_B z$, we can define $u = z + r\mu_*$. If choose $r = \left(T_* + \lambda_B \mathbf{e}' z\right) / \left(T_* - \lambda_B\right)$ then we conclude that

$$Ju = \left(\lambda_B / T_* - 1\right) u.$$

This selection of r is valid if z is not the principal right eigenvector of B (i.e. in case z is not μ_*). In case we select $z = \mu_*$, then trivially $Jz = -z$.

Consequently, we conclude that there is a one-to-one correspondence between the eigenvectors (and eigenvalues) of J and B, and the relationship between the eigenvalues is as follows

$$\lambda_J = \frac{\lambda_B}{T_*} - 1, \text{ for } \lambda_B \neq T_* \text{ or } \lambda_J \neq 0,$$
$$\lambda_J = -1 \text{ if } \lambda_B = T_*.$$

Therefore, in order to ensure that $J + I/2$ is Hurwitz, we must have that

$$\mathrm{Re}\,(\lambda_B) < T_*/2$$

for all $\lambda_B \neq T_*$, which is precisely condition (5). □

We finish the analysis of the Hurwitz condition, with the following result invoked in the previous proof.

Lemma 5.11. *We have that $\lambda_J \neq 0$.*

Proof of Lemma 5.11: Assume that y is such that $Jy = 0$. This implies that

$$By = T_* y + \mu_* \left(\mathbf{e}' By\right) = T_* y + \left(\mu_* \mathbf{e}'\right) By.$$

Therefore

$$\left(I - \left(\mu_* \mathbf{e}'\right)\right) By = T_* y. \tag{16}$$

We recognize that $\bar{P} = (I - (\mu_* \mathbf{e}'))$ is a (non-orthogonal) projection in the sense that $\bar{P}^2 = \bar{P}$. Also we have that $\bar{P}\mu_* = 0$ and $\mathbf{e}'\bar{P} = 0$. This means that T_* is an eigenvalue of $\bar{P}B$, which in turn implies that there would exist a left eigenvector x such that

$$x'\bar{P}B = T_* x',$$

or

$$x'\bar{P} = T_* x' B^{-1}. \tag{17}$$

Now let $x'\bar{P} = z'$ and consider all possible solutions w such that $w'\bar{P} = z'$ which must be written as the sum of an element of the null space and a particular solution. Observe, because $\bar{P}^2 = \bar{P}$, that $z'\bar{P} = z'$ is a particular solution and therefore any solution x (i.e. any eigenvector corresponding to T_* for $\bar{P}B$) must take the form $x = c\mathbf{e} + z$ for some constant c. Observe from (17), multiplying by μ_* from the right, that

$$0 = T_* x' B^{-1} \mu_* = T_* x' \left(I - Q'\right) \mu_* = x' \mu_*.$$

Therefore, we have that $c = 0$ because $x'\mu_* = 0$ and

$$0 = x'\mu_* = c + z'\mu_* = c + 0.$$

Consequently, $x' = z' = T_* x B^{-1}$, which implies that $x'B = T_* x'$, therefore concluding that x is the principal left-eigenvector of B. Consequently, x must have strictly positive entries and, in turn, we must have that $x'\mu_* > 0$ thus arriving at a contradiction. So, there is no eigenvalue T_* for the matrix $\bar{P}B$ and thus $\lambda_J = 0$ is not possible. □

5.2.3. *Tightness: Condition 7*

Lemma 5.12. *The sequence* $\{(\theta_n - \theta_*)/\varepsilon_n^{1/2}\}$ *is tight.*

Proof of Lemma 5.12: We use the local techniques discussed in Section 10.5.2 of [19] and apply them as in the proof of Theorem 10.4.1 in [19], albeit with some modifications. We shall use the Lyapunov function $V(\theta) = \|\theta - \theta_*\|_2^2$. However, we now have to deal with the gradient of g_n as opposed to the gradient of g as in the proof of Theorem 10.4.1 in [19]. We shall control the changes in g_n by expanding around the stationary point. We have that

$$\begin{aligned}
&E\left(V\left(\theta_{n+1}\right)|\mathcal{F}_n\right) - V\left(\theta_n\right) \\
&= 2\varepsilon_n\left(\theta_n - \theta_*\right)' g_n\left(\theta_n\right) + O(\varepsilon_n^2) \\
&= 2\varepsilon_n\left(\theta_n - \theta_*\right)' g_n\left(\theta_*\right) + 2\varepsilon_n\left(\theta_n - \theta_*\right)' A_n\left(\theta_n - \theta_*\right)' + \varepsilon_n o(\|\theta_n - \theta_*\|_2^2) + O(\varepsilon_n^2),
\end{aligned}$$

where the first equality uses an idea similar to that of the proof of Lemma 5.4 to arrive at the error term $O(\varepsilon_n^2)$, and the second inequality is just an expansion of g_n around θ_* followed by an application of Lemma 5.7. Since $2A_n$ has eigenvalues with negative real part less than -1 (i.e. $A_n + I/2$ is Hurwitz) for n large enough, we conclude that there exists $\delta > 0$ such that for all n sufficiently large

$$\left(\theta_n - \theta_*\right)' A_n\left(\theta_n - \theta_*\right)' < -\left(1 + 2\delta\right)\|\theta_n - \theta_*\|_2^2.$$

Moreover, because $g_n\left(\theta_*\right) = O\left(1/n\right)$ due to Lemma 5.8, we have that

$$2\varepsilon_n\left(\theta_n - \theta_*\right)' g_n\left(\theta_*\right) \leq O\left(\varepsilon_n^2\left(\theta_n - \theta_*\right)\right)$$

and

$$o(\|\theta_n - \theta_*\|_2^2) \leq \delta V\left(\theta_n\right).$$

We then conclude that

$$E\left(V\left(\theta_{n+1}\right)|\mathcal{F}_n\right) - V\left(\theta_n\right) \leq -\varepsilon_n\left(1 + \delta\right) V\left(\theta_n\right) + O\left(\varepsilon_n^2\right).$$

The rest of the proof now can be concluded as in Theorem 10.4.1 from [19]. □

5.2.4. *Quadratic Variation of the Martingales: Condition 8*

Lemma 5.13. *Let* $\delta M_n = W_n - g_n\left(\theta_n\right)$*, then*

$$\sup_{n\geq 0} E\,||\delta M_n||_2^4 < \infty. \tag{18}$$

Moreover,

$$E\left(\delta M_n \delta M_n' \mid \mathcal{F}_n\right) \to \Sigma \tag{19}$$

for some matrix Σ *in probability.*

Proof of Lemma 5.13: We use the notation $E_n(\cdot)$ for $E(\cdot \mid \mathcal{F}_n)$,

$$\begin{aligned} ||\delta M_n||_2^4 &\leq 2(\|Y_n(\theta_n) - E_n Y_n(\theta_n)\|_2^4 + |Z_n(\theta_n) - E_n Z_n(\theta_n)|^4) \\ &\leq 16\left(\|Y_n(\theta_n)\|_2^4 + \|EY_n(\theta_n)\|_2^4 + |Z_n(\theta_n)|^4 + |E_n Z_n(\theta_n)|^4\right). \end{aligned}$$

An argument similar to Lemma 5.3 yields that $\|Y_n(\theta_n)\|_2^4 \leq \bar{\tau}^4$ and $|Z_n(\theta_n)|^4 \leq \bar{\tau}^4$ (stochastically) and therefore, by Lemma 5.1 we conclude bound (18).

To establish (19) let us write $\delta M_n \delta M_n'$ in block matrix form, we obtain

$$\delta M_n \delta M_n' = \left[\begin{array}{cc} (Y_n(\theta_n) - f_n(\theta_n))(Y_n(\theta_n) - f_n(\theta_n))' & (Y_n(\theta_n) - f_n(\theta_n))(Z(\theta_n) - h(\theta_n)) \\ (Z(\theta_n) - h(\theta_n))(Y_n(\theta_n) - f_n(\theta_n))' & (Z(\theta_n) - h(\theta_n))^2 \end{array}\right].$$

By Lemma 5.8 we have that $f_n(\theta_n) \to 0$, and $h(\theta_n) \to 0$. Note that the distribution of $Z(\theta)$ and $Y(\theta)$ can be written in a way that is continuous in θ (as a mixture of the initial distribution), therefore $Z(\theta_n) \Rightarrow Z(\theta_*)$ and $Y(\theta_n) \Rightarrow Y(\theta_*)$; consequently we also have that $Y_n(\theta_n) \Rightarrow Y(\theta_*)$. We observe that each entry of the matrix is stochastically dominated by $2\bar{\tau}$ and thus we can apply Lemma 5.1 to conclude uniform integrability (since $E(\bar{\tau}^2) < \infty$), thereby obtaining that

$$E_n(\delta M_n \delta M_n') \to \Sigma := E\left(\left[\begin{array}{c} Y(\theta_*) \\ Z(\theta_*) \end{array}\right]\left[\begin{array}{cc} Y(\theta_*)' & Z(\theta_*) \end{array}\right]\right).$$

In turn, following the expression in p. 332 from [19], we have that the asymptotic covariance equals

$$V_0 = \int_0^\infty \exp((J + I/2)t)\,\Sigma \exp((J' + I/2)t)\,dt. \tag{20}$$

□

6. Proof of Proposition 4.1

Proof of Proposition 4.1: The proof is almost identical to that of Theorem 3.1. In fact, the analysis is somewhat simpler because there is no denominator and so we just need to analyze the reduced ODE in (12). The proof of tightness also follows similar steps as the argument give in Theorem 10.4.1 in [19], which distinguishes the cases $\varepsilon_n = 1/n$ and the case $\varepsilon_n = n^{-\alpha}$ for $\alpha \in (1/2, 1)$ as we do here.

Now, recall that J be the Jacobian of the vector field obtained in (15), evaluated at the unique stability point μ_*. We need to ensure that J is Hurwitz (i.e. all the eigenvalues have strictly negative real part). This is in contrast to requiring that $J + I/2$ is Hurwitz – which is a stronger condition because then one needs that all the eigenvalues have real part less than $-1/2$, which leads to (5). Instead, requiring that J be Hurwitz is equivalent to the condition that $\mathrm{Re}(\bar{\lambda}) < \lambda_*$ for all non-principal eigenvalue $\bar{\lambda}$ of the matrix Q, which is automatic by Perron-Frobenius Theorem, [16]. Hence we can conclude the result by invoking Theorem 10.2.1 from [19]. The asymptotic covariance matrix in this case takes the form

$$V_1 = \int_0^\infty \exp(Jt)\,\Sigma_0 \exp(J't)\,dt, \tag{21}$$

where

$$\Sigma_0 = E\left(\bar{Y}(\mu'_*)\,\bar{Y}(\mu'_*)'\right),$$

with $\bar{Y}(\mu') = \sum_{k=0}^{\tau-1} \left(I\left(X_k = \cdot \mid X_0 \sim \mu\right) - \mu(\cdot)\right)$. □

We now are ready to discuss the proof of Theorem 4.1

Proof of Theorem 4.1: We shall verify the conditions in Theorem 2 of [25]. First, define $\bar{f}(\mu) = f(\mu', 1)$, recall that $f(\cdot)$ was introduced in (8) and it coincides with $f(\mu', 1) = E\bar{Y}(\mu')$. We must first verify that,

- There exists a function L a (globally) Lipschitz continuous function $V(\cdot)$ such that $L(\mu_*) = 0$, and for some positive definite matrix G,

$$DL(\mu)\,G\bar{f}(\mu') < 0, \tag{22}$$

for $\mu \neq \mu_*$, there exists $\varepsilon, \delta > 0$ such that

$$DL(\mu)\,G\bar{f}(\mu') \leq -\delta V(\mu) \tag{23}$$

if $\|\mu - \mu_*\| \leq \varepsilon$, and $L(\mu - \mu_*) \geq \delta \|\mu - \mu_*\|_2^2$ for some $\delta > 0$.

This condition is satisfied if we construct $L(\cdot)$ by noting that μ_* is the unique root of $\bar{f}(\mu_*) = 0$ and we have established in Lemma 5.2 that $\bar{J} := \left(D\bar{f}\right)(\mu_*) = JT_*$ is Hurwitz (in fact $J + I/2 = \bar{J}/T_* + I/2$ is Hurwitz), therefore

$$\bar{f}(\mu) = \bar{J}(\mu - \mu_*) + O(\|\mu - \mu_*\|_2^2). \tag{24}$$

Now, it is standard in stability of dynamical systems that given a Hurwitz matrix $\bar{J}$ and a given symmetric positive definite matrix G (which we might take as $G = I$, the identity, here) there is a unique symmetric positive definite matrix K such that

$$\bar{J}'K + K\bar{J} = -G = -I$$

and there is a positive constant $\gamma > 0$ such that $I - \delta K$ is symmetric and positive definite. We first set $\bar{L}(\mu) = (\mu - \mu_*)'\,K(\mu - \mu_*)$ so that

$$\begin{aligned}
D\bar{L}(\mu)\,G\bar{f}(\mu') &= 2(\mu - \mu_*)'\,K\bar{J}(\mu - \mu_*) + O\left(\|\mu - \mu_*\|_2^3\right) \\
&= (\mu - \mu_*)'\left(\bar{J}'K + K\bar{J}\right)(\mu - \mu_*) + O\left(\|\mu - \mu_*\|_2^3\right) \\
&= -\|\mu - \mu_*\|_2^2 + O\left(\|\mu - \mu_*\|_2^3\right).
\end{aligned}$$

From this bound we obtain the existence of L satisfying (22) and (23) by modifying $\bar{L}$ outside a neighborhood of μ_* inside the compact set H.

The second condition in Theorem 2 of [25] follows directly from (24) and the fact that $\bar{J}$ is Hurwitz. There are two more conditions to check for the application of Theorem 2 of [25], the fourth condition is trivially satisfied for $\varepsilon_n = n^{-\alpha}$ with $\alpha \in (1/2, 1)$ and the third condition involves the martingale difference process and its quadratic variation, the verification is completely parallel to that of Proposition 5.2.□

7. Proof of Proposition 4.2

In this section we provide the proof of consistency for the estimator of

$$\bar{V}_1 = J^{-1}\Sigma_0 J^{-1}. \tag{25}$$

Expression (25) follows from Theorem 11.1.1 in [19].

Proof of Proposition 4.2: We write

$$\begin{aligned}
&\varepsilon_n^{-1}\left(\bar{\mu}_n - v_n\right)\left(\bar{\mu}_n - v_n\right)' && (26)\\
&= \varepsilon_n^{-1}\left(\bar{\mu}_n - \mu_*\right)\left(\bar{\mu}_n - \mu_*\right)' + \varepsilon_n^{-1}\left(v_n - \mu_*\right)\left(v_n - \mu_*\right)' && (27)\\
&- \varepsilon_n^{-1}\left(\bar{\mu}_n - \mu_*\right)\left(v_n - \mu_*\right)' - \varepsilon_n^{-1}\left(\bar{\mu}_n - \mu_*\right)\left(v_n - \mu_*\right)'.
\end{aligned}$$

Note that

$$\varepsilon_n^{-1/2}\left(\bar{\mu}_n - \mu_*\right)\left(\varepsilon_n n\right)^{-1/2} n^{1/2}\left(v_n - \mu_*\right)' \to 0$$

in probability as $n \to \infty$ because of Proposition 4.1 and Theorem 4.1, since $\varepsilon_n n \to \infty$ as $n \to \infty$. A similar argument applies to all the terms in (26) involving $(v_n - \mu_*)$; so it suffices to study the limit of

$$\frac{1}{N_n}\sum_{n_k=1}^{N_n} \varepsilon_{n_k}^{-1}\left(\bar{\mu}_{n_k} - \mu_*\right)\left(\bar{\mu}_{n_k} - \mu_*\right)'$$

as $n \to \infty$. The rest of the calculation is similar to analysis of the asymptotic covariance in Theorem 11.3.1 of [19]. The idea is that the sequence $\{(\bar{\mu}_{n_k} - \mu_*)\,\varepsilon_{n_k}^{-1/2}\}$ is weakly dependent and each term is asymptotically normal (as $k \to \infty$) with variance $\bar{V}_1$. □.

Acknowledgement

We acknowledge support from the NSF through the grants CMMI-1069064 and DMS-1320550.

References

[1] Aldous, D., Flannery, B., and Palacios, J. L. (1988). Two applications of urn processes to the fringe analysis of search trees and the simulation of quasi-stationary distributions of Markov chains. *Probability in the Engineering and Informational Sciences,* **2,** pp. 293-307.

[2] Athreya, K. and Karlin, S. (1968). Embedding of urn schemes into continuous time Markov branching processes and related limit theorems. *The Annals of Mathematical Statistics,* **6,** pp. 1801—1817.

[3] Benaïm, M. and Cloez, B. (2014). A stochastic approximation approach to quasi-stationary distributions on finite spaces. *http://arxiv.org/abs/1405.0448*

[4] Blanchet, J., Glynn, P., and Zheng, S. (2014). Empirical Analysis of a Stochastic Approximation Approach for Computing Quasi-stationary Distributions. *EVOLVE 2012,* pp. 19–37.

[5] Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization.* Cambridge University Press, Cambridge UK.

[6] Burzdy, K., Holyst, R., and March, P. (2000). A Fleming-Viot particle representation of the Dirichlet Laplacian. *Communications of Mathematical Physics,* **214**, pp. 679—703.

[7] Derroch, J. and Seneta, E. (1965). On quasi-stationary distributions in absorbing discrete-time finite Markov chains. *Journal of Applied Probability,* **2**, pp. 88—100.

[8] Derroch, J. and Seneta, E. (1967). On quasi-stationary distributions in absorbing continuous-time finite Markov chains. *Journal of Applied Probability,* **4**, pp. 192—196.

[9] De Oliveira, M. and Dickman, R. (2005). How to simulate the quasi-stationary state. *Physical Review E*, **71**.

[10] De Oliveira, M. and Dickman, R. (2006). Quasi-stationary simulation: the subcritical contact process. *Brazilian Journal of Physics*, **71**.

[11] Del Moral, P. and Miclo, L. (2006). Self-interacting Markov chains. *Stochastic Analysis and Applications,* **24**, *pp. 615—660.*

[12] Dickman, R. and Vidigal, R. (2002). Quasi-stationary distributions for stochastic processes with an absorbing state. *Journal of Physics A,* **35**, pp. 1147—1166.

[13] Dimov, I. T., Karaivanova, A. N., and Yordanova, P. I. (1998). Monte Carlo algorithms for calculating eigenvalues. In *Monte Carlo and Quasi-Monte Carlo Methods 1996: proceedings of a conference at the University of Salzburg, Austria, July 9-12, 1996,* **127**, p. 205.

[14] Ferrari, P. and Maric, N. (2007). Quasi-stationary distributions and Fleming-Viot processes in countable spaces. *Electronic Journal of Probability,* **12**, pp. 684—702.

[15] Golub, G. and van Loan, C. (1996). *Matrix Computations.* Johns Hopkins University Press; 3rd edition.

[16] Karlin, S., and Taylor, H. (1975). *A First Course in Stochastic Processes.* Academic Press.

[17] Krasulina, T. The method of stochastic approximation for the determination of the least eigenvalue of a symmetrical matrix. *USSR Computational Mathematics and Mathematical Physics*, **9,** pp 189—195.

[18] Krasulina, T. Method of stochastic approximation in the determination of the largest eigenvalue of the mathematical expectation of random matrices. *Automatation and Remote Control*, pp 50—56.

[19] Kushner, H. and Yin, G. (2003). *Stochastic Approximation and Recursive Algorithms and Applications.* Springer-Verlag, New York.

[20] Liggett, T. (1985). *Interacting Particle Systems.* Springer, New York.

[21] Marro, J. and Dickman, R. (1999). *Non-equilibrium Phase Transitions in Lattice Models,* Cambridge University Press.

[22] Méléard, S., and Villemonais, D. (2011). Quasi-stationary distributions and population processes. ArXiV preprint arXiv:1112.4732.

[23] Oja, E. and Karhunen, J. (1985). On stochastic approximation of the eigenvectors and eigenvalues of the expectation of a random matrix. *Journal of Mathematical Analysis and Applications*, **106**, pp. 69—84.

[24] Permentle, R. (2007). A survey of random processes with reinforcement. *Probability Surveys,* **4**, pp. 1—79.

[25] Polyak, B. T., and Juditsky, A. B. (1992). Acceleration of stochastic approximation by averaging. *SIAM Journal of Control and Optimization,* **30**, pp. 838—855.

[26] Robbins, H., and Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics,* **22**, pp. 400—407.

[27] Zheng, S. (2014). *Stochastic Approximation Algorithms in the Estimation of Quasi-stationary Distributions of Finite and General State Space Markov Chains.* Ph.D. Thesis, Columbia University.