Stationary Distributions for Random Markov Chains

Vasco Portilheiro

March 22, 2021

1 Introduction

The purpose of this note is to have written down some (very preliminary) work on the following problem. Consider a finite state space $[n] = \{1, \ldots, n\}$ equipt with a random Markov transition kernel — that is, a random row-stochastic matrix $P \in [0,1]^{n \times n}$. What can we say about the the stationary distribution π? That a unique π exists almost surely can be easily guaranteed by a natural choice of distribution on P, which I consider here. Namely, we let the rows of P be independently Dirichlet.[1]

The rest of this note is organized as follows. I present the model and some of its easily identifiable properties in Section 2, followed by the conjecture on the distribution of π in Section 3. Next, Section 4 is devoted to explicitly finding the distribution of π when the number of states is $n = 2$. For ease of exposition I relegate to Section 5 some discussion of the fraction of relevant literature I have encountered so far [TODO] .

2 Some basic calculations

I apologize in advance to any reader for the slightly cumbersome notation.

With that caveat, let's fully state the model. Considering fixed $\alpha_{ij} > 0$ for $i, j \in [n]$, the rows of P are independently sampled as

$$P_i \sim \mathrm{Dir}(\alpha_{i1}, \ldots, \alpha_{in}), \forall i \in [n].$$

We'll let ν denote the distribution of P, and ρ the distribution of any of its rows. Given an initial state X_0 (possibly random), the process $(X_t : t \geq 0)$ is given by

$$\mathbb{P}\left(X_t = j | X_{t-1} = i, P = p\right) = p_{ij}.$$

To save on space, I will write X_s^t for the sequence $(X_s, \ldots, X_t)$.

[1] I owe my thanks to Jim Pitman for suggesting this setup.

"The most basic question" might be: how are sequences of states distributed? It is known that in a mixture of Markov chains, the initial state X_0 together with the set of transition counts between pairs of states is a sufficient statistic (Diaconis and Freedman 1980). Letting $c_{ij}(x_s^t) = |\{r \in [s, t-1] : x_{r+1} = i, x_r = j\}|$ be the count of i-to-j transitions in x_s^t, we see this fact in our answer to the most basic question.

$$\begin{aligned} \mathbb{P}\left(X_1^t = x_1^t | X_0 = x_0\right) &= \int \mathbb{P}\left(X_1^t = x_1^t | X_0 = x_0, P = p\right) \nu(dp) \\ &= \prod_{i=1}^n \int \prod_{j=1}^n p_{ij}^{c_{ij}(x_0^t)} \rho(dp_i) \\ &= \prod_{i=1}^n \frac{\mathrm{B}(c_{i1}(x_0^t) + \alpha_{i1}, \ldots, c_{in}(x_0^t) + \alpha_{in})}{\mathrm{B}(\alpha_{i1}, \ldots, \alpha_{in})} \end{aligned} \tag{1}$$

Here, I'm using the following notation for the "multivariate Beta" function

$$\mathrm{B}(\alpha_1, \ldots, \alpha_m) = \frac{\Gamma(\alpha_1) \cdots \Gamma(\alpha_m)}{\Gamma(\alpha_1 + \cdots + \alpha_m)}.$$

From (1) we can derive other expressions of interest, such as the predictive distribution

$$\mathbb{P}\left(X_{t+1} = x_{t+1} | X_0^t = x_0^t\right) = \frac{c_{x_t x_{t+1}}(x_0^t) + \alpha_{x_t x_{t+1}}}{\sum_{j=1}^n c_{x_t j}(x_0^t) + \alpha_{x_t j}} \tag{2}$$

and the posterior density for P with respect to Lebesgue measure λ on the set $\{p \in [0,1]^{n \times n} : \sum_{j=1}^n p_{ij} = 1 \ \forall i \in [n]\}$ given an observed sequence x_0^t,

$$\frac{d\nu(p|x_0^t)}{d\lambda} = \prod_{i=1}^n \frac{\prod_{j=1}^n p_{ij}^{c_{ij}(x_0^t) + \alpha_{ij} - 1}}{\mathrm{B}\left(\sum_{k=1}^n c_{ik}(x_0^t) + \alpha_{ik}\right)}. \tag{3}$$

(Actually, the posterior can of course be derived directly.)

Equation (2) shows that we can view $(X_t : t \geq 0)$ as coming from an urn model. Namely, imagine we have n urns, each containing balls of n different colors: the i-th urn starts out with α_{ij} balls of color j. (We allow fractional ball counts here.) Then (2) describes the sequence of colors drawn when, starting from urn X_0, we draw a ball at random of color X_1 and replace two balls of that color in the urn, then move to the urn numbered X_1 and repeat. Equivalently, we can also view this process as an *edge-reinforced* random walk on a directed graph of n vertices with initial weight α_{ij} on edge (i, j), where upon traversing any edge we increase its weight by one.

Let's turn our attention now to the main goal: the random stationary distribution π. Since P is almost surely irreducible, the ergodic theorem guarantees that as $t \to \infty$

$$\frac{\sum_{j=1}^n c_{ij}(X_0^t)}{t} \stackrel{\text{a.s.}}{\to} \pi(i).$$

This motivates writing down one more expression: the probability of a given number of occurrences of a state k up to but not including time t, or equivalently, the number of transitions out of k up to time t. Let

$$\mathcal{P}_k(s,t) = \left\{ x_0^t \in [n]^{t+1} : \sum_{j=1}^{n} c_{ij}(X_0^t) = s \right\}$$

be the set of paths x_0^t for which out of the t total transitions, s were from state k. Then

$$\mathbb{P}\left(\sum_{j=1}^{n} c_{kj}(X_0^t) = s \right) = \sum_{x_0^t \in \mathcal{P}_k(s,t)} \mathbb{P}\left(X_0 = x_0\right) \prod_{i=1}^{n} \frac{\mathrm{B}(c_{i1}(x_0^t)+\alpha_{i1}, \ldots, c_{in}(x_0^t)+\alpha_{in})}{\mathrm{B}(\alpha_{i1}, \ldots, \alpha_{in})} \tag{4}$$

follows from (1). We thus have a first characterization of the distribution of π

$$\mathbb{P}\left(\pi(k) \le r\right) = \lim_{t\to\infty} \sum_{s=0}^{\lfloor rt \rfloor} \sum_{x_0^t \in \mathcal{P}_k(s,t)} \mathbb{P}\left(X_0 = x_0\right) \prod_{i=1}^{n} \frac{\mathrm{B}(c_{i1}(x_0^t)+\alpha_{i1}, \ldots, c_{in}(x_0^t)+\alpha_{in})}{\mathrm{B}(\alpha_{i1}, \ldots, \alpha_{in})}. \tag{5}$$

3 A conjecture

I think it's uncontroversial to call (5) unsatisfying. Fortunately, it's rather easy to investigate π directly through simulation. We don't even need to run the Markov chain; we just sample P and find its left eigenvector π of eigenvalue 1, and plot the histogram of values of $\pi(1)$, say. Below are these exact plots, where the Dirichlet distributions are identically symmetric, that is $\alpha_{ij} = \alpha$ for all i, j.

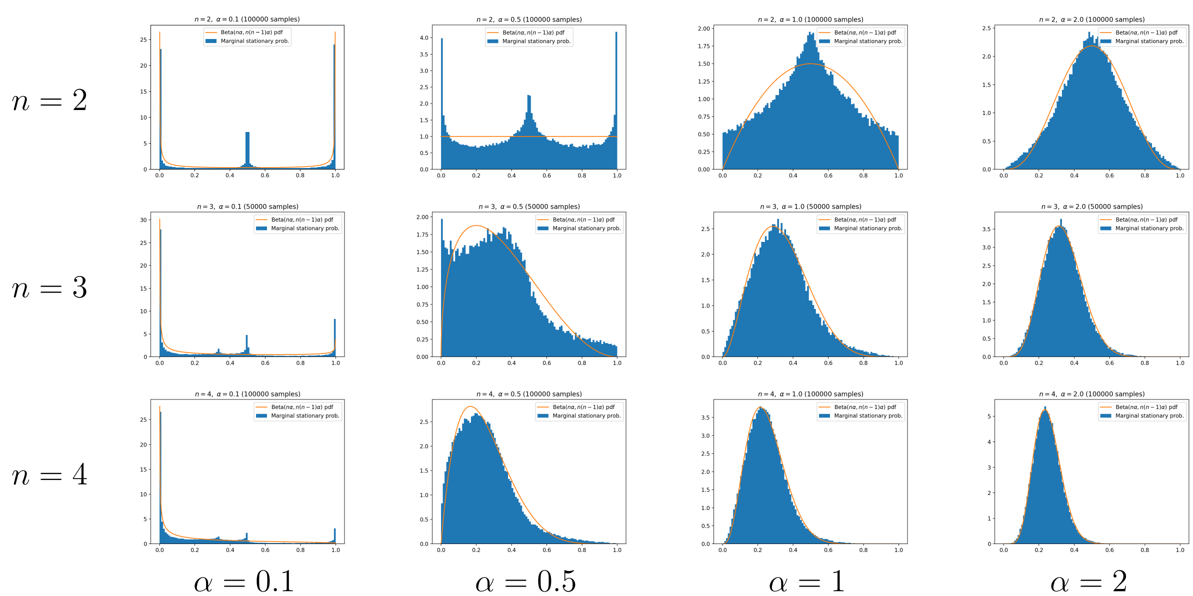

The orange lines are density functions of $\mathrm{Beta}(n\alpha, n(n-1)\alpha)$. These would be the marginals $\pi(1)$ if $\pi \sim \mathrm{Dir}(n\alpha\vec{1})$. The empirical densities seem to match these orange lines both as n increases and as α increases, giving us the following conjecture.

Conjecture 1. Let $P \in \mathbb{R}^{n \times n}$ with i.i.d. rows $P_i \sim \text{Dir}(\alpha\vec{1})$ for some $\alpha > 0$. Then for π the unique probability distribution such that $\pi = \pi P$, we have

$$d(\mathcal{L}(\pi), \text{Dir}(n\alpha\vec{1})) \to 0$$

both as $n \to \infty$ and as $\alpha \to \infty$, for some distance between distributions d. Furthermore, this convergence happens "quickly."

A couple notes on this conjecture. First, I believe it's natural to consider a fixed symmetric Dirichlet as the distribution of i.i.d. rows of P, as a way to encode the idea that no state is "special." Secondly, a restricted version of this conjecture, with $\alpha = 1$ fixed, has been made by Bordenave, Caputo, and Chafai (2008), but as far as I know no work has been published towards its proof (or disproof).

4 The case of two states

In the case that $n = 2$, it's straightforward to find the distribution of $\pi(1)$ explicitly. We can write

$$P = \begin{bmatrix} 1-\theta_1 & \theta_1 \\ \theta_2 & 1-\theta_2 \end{bmatrix}$$

where $\theta_1 \sim \text{Beta}(\alpha_{11}, \alpha_{12})$ and $\theta_2 \sim \text{Beta}(\alpha_{21}, \alpha_{22})$ are independent. Some linear algebra shows that

$$\pi(1) = \frac{\theta_2}{\theta_1 + \theta_2} \qquad \pi(2) = \frac{\theta_1}{\theta_1 + \theta_2}.$$

It can be shown (see Pham-Gia (2000)) that $\pi(1)$ has density at t of

$$\frac{t^{\alpha_{21}-1}(1-t)^{\alpha_{21}+1}\text{B}(\alpha_{11}+\alpha_{21}, \alpha_{12})_2F_1\left(\alpha_{11}+\alpha_{21}, 1-\alpha_{22}; \alpha_{11}+\alpha_{21}+\alpha_{12}; \frac{t}{1-t}\right)}{\text{B}(\alpha_{11}, \alpha_{12})\text{B}(\alpha_{21}, \alpha_{22})}$$

for $t \in (0, 1/2]$ and

$$\frac{t^{-(\alpha_{11}+1)}(1-t)^{\alpha_{11}-1}\text{B}(\alpha_{11}+\alpha_{21}, \alpha_{22})_2F_1\left(\alpha_{11}+\alpha_{21}, 1-\alpha_{12}; \alpha_{11}+\alpha_{21}+\alpha_{22}; \frac{t}{1-t}\right)}{\text{B}(\alpha_{11}, \alpha_{12})\text{B}(\alpha_{21}, \alpha_{22})}$$

for $t \in [1/2, 1]$, where

$${}_2F_1(a, b; c; x) = \sum_{m=0}^{\infty} \frac{a^{(m)} b^{(m)} x^m}{c^{(m)} m!} = \int_0^1 \frac{u^{a-1}(1-u)^{c-a-1}(1-xu)^{-b}}{\text{B}(a, c-a)} du$$

is the ordinary hypergeometric function. (As is common, if not standard, $a^{(m)}$ refers to the m-th rising factorial of a.)

At the moment, it is unclear to me whether the expression above is helpful in proving Conjecture 1 for fixed $n = 2$ — probably in large part due to my unfamiliarity with the hypergeometric function. Furthermore, I don't see how the above might generalize to $n > 2$.

5 Some relevant work and ideas

TODO

References

Bordenave, Charles, Pietro Caputo, and Djalil Chafai (2008). "Circular Law Theorem for Random Markov Matrices". In: May 2010. DOI: 10.1007/s00440-010-0336-1. arXiv: 0808.1502. URL: http://arxiv.org/abs/0808.1502%7B%5C%%7D0Ahttp://dx.doi.org/10.1007/s00440-010-0336-1.

Diaconis, P. and D. Freedman (1980). "De Finetti's Theorem For Markov Chains". In: *The Annnals of Probability* 8.1, pp. 115–130. ISSN: 0091-1798. URL: http://projecteuclid.org/euclid.aop/1176996548.

Pham-Gia, T. (2000). "Distributions of the ratios of independent beta variables and applications". In: *Communications in Statistics - Theory and Methods* 29.12, pp. 2693–2715. ISSN: 03610926. DOI: 10.1080/03610920008832632.