

Note 9

Discrete Markov chains and ergodicity. Doob's h -transform.

1 Markov chains on discrete spaces

We will restrict now our considerations to the case where the state space E of the homogeneous Markov chain is discrete (maybe infinite). In this case the transition kernel $P: E \rightarrow \Pi(E, \mathcal{E})$ is equivalent to the transition matrix $P: E \times E \rightarrow \mathbb{R}$ given by $P(x, y) = P(x, \{y\})$ for all $x, y \in E$. This matrix is *stochastic*, i.e. $\sum_{y \in E} P(x, y) = 1$ for all $x \in E$ and $P(x, y) \in [0, 1]$ for all $x, y \in E$. We will denote also $T_x = T_{\{x\}}$ the stopping time to reach $x \in E$.

Moreover we will assume that the Markov chain is *irreducible*, that is there exists positive probability to go from any state to any other, i.e. $\mathbb{P}_x(T_y < \infty) > 0$ for all $x, y \in E$. Said differently: for all $x, y \in E$ there exists $n = n(x, y)$ such that $P^n(x, y) > 0$ where P^n is the n -fold matrix product of the transition matrix P (equivalent to the n -fold composition of the transition kernels).

In this context we say that the chain is

- *transient* if $\mathbb{P}_x(T_x < \infty) < 1$ for all $x \in E$;
- *recurrent* if $\mathbb{P}_x(T_x < \infty) = 1$ for all $x \in E$;
- *positive recurrent* if $\mathbb{E}_x[T_x] < \infty$ for all $x \in E$.

Remark 1. Similar notions can be attached to the single states of the chain. By irreducibility is not difficult to show that they are properties which are common to all the states of the chain. Moreover they can be extended to non-irreducible chains introducing the notions of communication classes and noting that a general chain can be decomposed in irreducible components and a set of transient states. We will not concern ourselves here with the general theory which does not present any additional substantial difficulty.

Theorem 2. For all $x \in E$:

$$\mathbb{P}_x(X_n = x \text{ infinitely often}) = 0 \Leftrightarrow \mathbb{E}_x \left[\sum_{n \geq 1} \mathbb{1}_{X_n = x} \right] < +\infty \Leftrightarrow \mathbb{P}_x(T_x < \infty) < 1,$$

$$\mathbb{P}_x(X_n = x \text{ infinitely often}) = 1 \Leftrightarrow \mathbb{E}_x \left[\sum_{n \geq 1} \mathbb{1}_{X_n = x} \right] = +\infty \Leftrightarrow \mathbb{P}_x(T_x < \infty) = 1.$$

Moreover if the chain is irreducible then for any $x, y \in E$

$$\mathbb{E}_x \left[\sum_{n \geq 1} \mathbb{1}_{X_n = x} \right] < +\infty \Leftrightarrow \mathbb{E}_y \left[\sum_{n \geq 1} \mathbb{1}_{X_n = y} \right] < +\infty.$$

Proof. Let $\lambda > 0$. By the strong Markov property we have

$$\begin{aligned} \mathbb{E}_x \left[\sum_{n>0} e^{-\lambda n} \mathbb{1}_{X_n=x} \right] &= \mathbb{E}_x \left[e^{-\lambda T_x} + e^{-\lambda T_x} \mathbb{E} \left[\sum_{n-T_x>0} e^{-\lambda(n-T_x)} \mathbb{1}_{X_n=x} \middle| \mathcal{F}_{T_x} \right] \right] \\ &= \mathbb{E}_x [e^{-\lambda T_x}] + \mathbb{E}_x \left[e^{-\lambda T_x} \mathbb{E}_{X_{T_x}} \left[\sum_{n>0} e^{-\lambda n} \mathbb{1}_{X_n=x} \right] \right] = \mathbb{E}_x [e^{-\lambda T_x}] \left(1 + \mathbb{E}_x \left[\sum_{n>0} e^{-\lambda n} \mathbb{1}_{X_n=x} \right] \right) \end{aligned}$$

and therefore

$$\mathbb{E}_x \left[\sum_{n>0} e^{-\lambda n} \mathbb{1}_{X_n=x} \right] = \frac{\mathbb{E}_x [e^{-\lambda T_x}]}{1 - \mathbb{E}_x [e^{-\lambda T_x}]}.$$

Assume $\mathbb{P}_x(T_x < \infty) < 1$, then as $\lambda \downarrow 0$ we have

$$\mathbb{E}_x \left[\sum_{n>0} \mathbb{1}_{X_n=x} \right] = \lim_{\lambda \downarrow 0} \mathbb{E}_x \left[\sum_{n>0} e^{-\lambda n} \mathbb{1}_{X_n=x} \right] = \lim_{\lambda \downarrow 0} \frac{\mathbb{E}_x [e^{-\lambda T_x}]}{1 - \mathbb{E}_x [e^{-\lambda T_x}]} = \frac{\mathbb{P}_x(T_x < \infty)}{1 - \mathbb{P}_x(T_x < \infty)} < \infty$$

and therefore $\sum_{n>0} \mathbb{1}_{X_n=x} < \infty$ \mathbb{P}_x -a.e. and $\mathbb{P}_x(X_n = x \text{ infinitely often}) = 0$. On the other hand if $\mathbb{P}_x(X_n = x \text{ infinitely often}) > 0$ then also $\mathbb{E}_x[\sum_{n>0} \mathbb{1}_{X_n=x}] = +\infty$ and the same limit necessarily gives $\mathbb{P}_x(T_x < \infty) = 1$. Then it is easy to conclude using the strong Markov property that

$$\mathbb{P}_x(T_x < \infty) = 1 \Rightarrow \mathbb{P}_x(X_n = x \text{ infinitely often}) = 1.$$

From this we can deduce all the remaining implications in the first two claims. For the last, observe that, given two states $x, y \in E$ by irreducibility there exists m_1, m_2 such that $P^{m_1}(x, y) > 0$ and $P^{m_2}(y, x) > 0$, therefore

$$P^{m_1+m_2+k}(y, y) \geq P^{m_2}(y, x) P^k(x, x) P^{m_1}(x, y)$$

and

$$\sum_{n \geq 1} P^n(y, y) \geq P^{m_2}(y, x) P^{m_1}(x, y) \sum_{n \geq 1} P^n(x, x)$$

from which we conclude since $\mathbb{E}_x[\sum_{n \geq 1} \mathbb{1}_{X_n=x}] = \sum_{n \geq 1} P^n(x, x)$. \square

In particular, for any irreducible chain the tail event $\{X_n = x \text{ infinitely often}\}$ is trivial.

We are going now to study the asymptotic behaviour of irreducible chains. A measure μ on (E, \mathcal{E}) (not necessarily a probability) is invariant for the transition kernel P iff it is nontrivial and

$$\mu P = \mu.$$

Similarly a probability measure $\pi \in \Pi(E, \mathcal{E})$ is invariant for a Markov chain with kernel P iff π is invariant wrt. P . In this case, if we choose the initial state of the chain according to π then the law of the chain is invariant under time shift, namely, defining $\mathbb{P}_\pi = \int \mathbb{P}_x \pi(dx)$ we have

$$\mathbb{E}_\pi(F \circ \theta_n) = \mathbb{E}_\pi[F]$$

for all bounded measurable $F: E^{\mathbb{N}} \rightarrow \mathbb{R}$ and all $n \geq 0$.

Let us settle the question of uniqueness of invariant measures for discrete chains.

Theorem 3. *If a discrete chain is irreducible, then any two invariant measures differ by a multiplicative constant and there exists at most one invariant probability.*

Proof. Since the chain is irreducible, if ρ is an invariant measure we need to have $\rho(y) > 0$ for all $y \in E$. Indeed, let $x \in E$ such that $\rho(x) > 0$ and observe that by irreducibility there exists n such that $P^n(x, y) > 0$, but by invariance $\rho(y) = (P\rho)(y) = \dots = (P^n\rho)(y) \geq P^n(x, y)\rho(x) > 0$.

Let ν be an invariant measure for P and define $Q(x, y) := P(x, y)\nu(y) / \nu(x)$ for $x, y \in E$ and observe that $Q(x, y)$ is again a stochastic kernel and that the measure ν^2 is invariant for it, indeed

$$\sum_{x \in E} \nu^2(x)Q(x, y) = \sum_{x \in E} \nu(x)P(x, y)\nu(y) = \nu^2(y), \quad y \in E.$$

Moreover, if ρ is another invariant measure for P we can define the function $f(x) = \rho(x) / \nu(x)$ since $\nu(x) > 0$ for all $x \in E$. This function is positive and in $L^2(\nu^2)$, indeed

$$\sum_{x \in E} f(x)^2 \nu^2(x) = \sum_{x \in E} \rho(x)^2 \leq 1.$$

Now consider the positive quadratic form

$$\mathcal{D}(h) := \sum_{x, y \in E} (h(x) - h(y))^2 \nu^2(x)Q(x, y)$$

which is well defined for all $h \in L^2(\nu^2)$ since

$$\sum_{x, y \in E} (h(x) - h(y))^2 \nu^2(x)Q(x, y) \leq 2 \sum_{x, y \in E} (h(x)^2 + h(y)^2) \nu^2(x)Q(x, y) \leq 4\nu^2(h^2) < \infty.$$

Computing $\mathcal{D}(f)$ we have

$$\begin{aligned} \mathcal{D}(f) &= \sum_{x, y \in E} (f(x)^2 + f(y)^2 - 2f(x)f(y)) \nu^2(x)Q(x, y) \\ &= 2 \sum_{x \in E} f^2(x) \nu^2(x) - 2 \sum_{x, y \in E} f(y) \rho(x) P(x, y) \nu(y) = 2 \sum_{x \in E} f^2(x) \nu^2(x) - 2 \sum_{x \in E} f^2(x) \nu^2(x) = 0 \end{aligned}$$

where we used $\rho P = \rho$.

Therefore it must be that $(f(x) - f(y))^2 \nu^2(x)Q(x, y) = 0$ for all $x, y \in E$. This means that $f(x) = f(y)$ for all x, y for which $P(x, y) > 0$ since ν is strictly positive everywhere. However irreducibility also implies that for any two states $x, y \in E$ there exists an $n > 0$ and states x_1, \dots, x_{n-1} such that $P(x_{i-1}, x_i) > 0$ for all $i = 1, \dots, n$ with $x_0 = x$ and $x_n = y$. As a consequence $f(x) = f(x_1) = \dots = f(x_{n-1}) = f(y)$ and therefore the function f must be constant. From this we deduce that f must be constant and that the two measures are proportional. Then it is also clear that if it exists an invariant probability then it is unique. \square

Note that it could exist an invariant measure but not an invariant probability, think to the case of the simple random walk on \mathbb{Z} where any constant measure is invariant. Note that it cannot exist an invariant probability since by the above argument it must be proportional to the constant measure which is non-normalizable.

If the chain is recurrent, for any $x \in E$ we can define the measure ν^x as

$$\nu^x(y) = \mathbb{E}_x \left[\sum_{n=1}^{T_x} \mathbb{1}_{X_n=y} \right], \quad y \in E.$$

Theorem 4. *If the chain is irreducible and recurrent then ν^x is an invariant measure.*

Proof. Let $\nu^x(y) = \mathbb{E}_x[\sum_{n=1}^{T_x} \mathbb{1}_{X_n=y}]$. We show that ν^x is an invariant measure, namely that $\nu^x P = \nu^x$:

$$\begin{aligned} \nu^x(y) &= \mathbb{E}_x \left[\sum_{n=1}^{T_x} \mathbb{1}_{X_n=y} \right] = \sum_{z \in E} \mathbb{E}_x \left[\sum_{n=1}^{T_x} \mathbb{1}_{X_n=y, X_{n-1}=z} \right] \\ &= \sum_{z \in E} \mathbb{E}_x \left[\sum_{n \geq 1} \mathbb{1}_{T_x \geq n} \mathbb{E}_x[\mathbb{1}_{X_n=y, X_{n-1}=z} | \mathcal{F}_{n-1}] \right] \\ &= \sum_{z \in E} \mathbb{E}_x \left[\sum_{n \geq 1} \mathbb{1}_{T_x \geq n} \mathbb{1}_{X_{n-1}=z} \right] P(z, y) = \sum_{z \in E} \mathbb{E}_x \left[\sum_{n \geq 1} \mathbb{1}_{T_x \geq n} \mathbb{1}_{X_n=z} \right] P(z, y) \\ &= \sum_{z \in E} \nu^x(z) P(z, y) = (\nu^x P)(y) \end{aligned}$$

where we used the Markov property and the fact that, by recurrence, $X_0 = X_{T_x} = x$ to rewrite the summation over n . As a consequence μ^x is an invariant probability for the chain. Note that if the chain is not assumed recurrent then the only thing we can deduce is that $\nu^x(y) \geq (\nu^x P)(y)$. By the previous theorem \square

By irreducibility we must have $\nu^x = C_{x,y} \nu^y$ for some constant $C_{x,y}$. However note that

$$\nu^x(E) = \mathbb{E}_x[T_x], \quad \nu^x(x) = 1, \quad x \in E,$$

since $\sum_{n=1}^{T_x} \mathbb{1}_{X_n=x} = 1$, \mathbb{P}_x -a.s. Therefore if the chain is positive recurrent then $C_{x,y} = \mathbb{E}_x[T_x] / \mathbb{E}_y[T_y]$ and $\mu^x = \nu^x / \mathbb{E}_x[T_x]$ is a probability measure for which

$$\mu^x(y) = \mu^y(y) = \frac{1}{\mathbb{E}_y[T_y]},$$

so it does not depend on x and calling it simply μ we have $\mu(y) = (\mathbb{E}_y[T_y])^{-1}$. We record this result in the following

Corollary 5. *If the chain is positive recurrent and irreducible then the probability measure*

$$\mu(x) = 1 / \mathbb{E}_x(T_x), \quad x \in E,$$

is the only invariant measure of the chain.

Theorem 6. *If μ is an invariant measure for an irreducible chain, then the chain is recurrent. Moreover if μ is finite then the chain is positive recurrent and therefore $\pi = \mu / \mu(E)$ is the only invariant probability.*

Proof. Assume the chain is transient and irreducible then for any x, y there exists n such that $P^n(y, x) > 0$ and at the same time $0 = \mathbb{P}_y(X_n = y \text{ i.o.}) \geq P^n(y, x) \mathbb{P}_x(X_n = y \text{ i.o.})$. Therefore for all $x, y \in E$ we have $\mathbb{P}_x(X_n = y \text{ i.o.}) = 0$. But then we must also have (by dominated convergence) $\lim_{n \rightarrow \infty} \mathbb{P}_x(X_n = y) = 0$ and also

$$\lim_{n \rightarrow \infty} \mathbb{P}_\mu(X_n = y) = \lim_{n \rightarrow \infty} \sum_{x \in E} \mu(x) \mathbb{P}_x(X_n = y) = 0.$$

But this is not possible since by invariance $\mathbb{P}_\mu(X_n = y) = \mu(y)$ for all n and this implies $\mu \equiv 0$. Therefore the chain is recurrent and we can consider the measure ν^x . By irreducibility we deduce that $\nu^x = C\mu$. Summing over all E we obtain $\mathbb{E}_x[T_x] = \nu^x(E) = C\mu(E) < \infty$ whenever the measure μ is finite and the chain is positive recurrent. We conclude that π is the unique invariant probability. \square

2 Doob's h -transform

We work here on the canonical space of a Markov process $(X_n)_{n \geq 0}$ with the probability \mathbb{P}_{x_0} . Let h be a positive harmonic function such that $h(x_0) = 1$ for some $x_0 \in E$. Then $(h(X_n))_{n \geq 0}$ is a martingale with average 1 under \mathbb{P}_{x_0} . On \mathcal{F}_n we can define the probability \mathbb{Q}_n by

$$\mathbb{E}_{\mathbb{Q}_n}[F] = \mathbb{E}_{x_0}[h(X_n)F].$$

The family $(\mathbb{Q}_n)_{n \geq 0}$ is a consistent family of probabilities defined on the increasing family of σ -algebras $(\mathcal{F}_n)_{n \geq 0}$, indeed by the martingale property of $(h(X_n))_{n \geq 0}$, for all $F \in \mathcal{F}_n$ and for all $n \leq m$ we have

$$\mathbb{E}_{\mathbb{Q}_n}[F] = \mathbb{E}_{x_0}[h(X_n)F] = \mathbb{E}_{x_0}[h(X_m)F] = \mathbb{E}_{\mathbb{Q}_m}[F].$$

Therefore it defines a unique probability measure \mathbb{Q} on $\mathcal{F} = \bigvee_{n \geq 0} \mathcal{F}_n$ by Caratheodory extension theorem. The measure \mathbb{Q} is called Doob's h -transform of \mathbb{P} . We have

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_n} = h(X_n), \quad n \geq 0.$$

Note that the \mathbb{Q} probability that $(X_n)_{n \geq 0}$ visits the set $Z = \{x: h(x) = 0\}$ is zero. Indeed

$$\mathbb{Q}(T_Z < \infty) = \lim_{n \rightarrow \infty} \mathbb{E}[h(X_n) \mathbb{1}_{T_Z \leq n}] = \lim_{n \rightarrow \infty} \mathbb{E} \left[\underbrace{\mathbb{E}[h(X_n) | \mathcal{F}_{T_Z}]}_{=0} \mathbb{1}_{T_Z \leq n} \right] = 0.$$

Moreover since $(h(X_n))_{n \geq 0}$ is a positive martingale, if we have $h(X_n) = 0$ for some n then $h(X_m) = 0$ for all $m \geq 0$ and therefore $h(X_n) = h(X_{n \wedge T_Z})$.

Under the measure \mathbb{Q} the process $(X_n)_{n \geq 0}$ is a Markov process with generator

$$\mathcal{L}^h f = h^{-1}(\mathcal{L}(hf))$$

We have that for all $f = h^{-1}g$ bounded and measurable

$$\hat{M}_n^f = (h^{-1}g)(X_n) - (h^{-1}g)(X_0) - \sum_{k=0}^{n-1} (\mathcal{L}^h f)(X_k)$$

is a martingale under \mathbb{Q} : for all $n \geq 0$ and all $A \in \mathcal{F}_n$

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}[\Delta \hat{M}_{n+1}^f \mathbb{1}_A] &= \mathbb{E}_x[h(X_{n+1}) \Delta M_{n+1}^f \mathbb{1}_A] \\ &= \mathbb{E}_x[h(X_{n+1})((h^{-1}g)(X_{n+1}) - (h^{-1}g)(X_n) - \mathcal{L}^h f(X_n)) \mathbb{1}_A] \\ &= \mathbb{E}_x[(g(X_{n+1}) - g(X_n) - h(X_n) \mathcal{L} f(X_n)) \mathbb{1}_A] = \mathbb{E}_x[(Pg(X_n) - g(X_n) - (h \mathcal{L}^h f)(X_n)) \mathbb{1}_A] = 0. \end{aligned}$$

The h -transform is a useful tool to describe Markov processes conditioned upon certain events.

For example consider the event that the Markov process never touch a given set $A \in \mathcal{E}$ and assume that $\mathbb{P}_{x_0}(T_A = +\infty) > 0$. Then let

$$h(x) = \frac{\mathbb{P}_x(T_A = +\infty)}{\mathbb{P}_{x_0}(T_A = +\infty)}, \quad x \in E.$$

By construction the function h is a positive harmonic function in A^c with $h(x_0) = 1$ and $h(x) = 0$ on A . This is not quite what we had before but indeed all we used is the availability of a positive martingale with average 1. In this case we can construct the h -transformed measure \mathbb{Q} using the positive martingale $(h(X_{n \wedge T_A}))_n$.

For all event $B \in \mathcal{F}_n$

$$\mathbb{P}_{x_0}(B|T_A = +\infty) = \frac{\mathbb{P}_{x_0}(B, T_A = +\infty)}{\mathbb{P}_{x_0}(T_A = +\infty)} = \frac{\mathbb{E}_{x_0}[\mathbb{1}_B \mathbb{E}[\mathbb{1}_{T_A = +\infty} | \mathcal{F}_n]]}{\mathbb{P}_{x_0}(T_A = +\infty)}$$

But now $\mathbb{E}[\mathbb{1}_{T_A = +\infty} | \mathcal{F}_n] = \mathbb{E}[\mathbb{1}_{T_A = +\infty} \circ \theta_n | \mathcal{F}_n] \mathbb{1}_{T_A \geq n} = h(X_n) \mathbb{1}_{T_A \geq n} = h(X_{n \wedge T_A})$ so

$$\mathbb{P}_{x_0}(B|T_A = +\infty) = \frac{\mathbb{P}_{x_0}(B, T_A = +\infty)}{\mathbb{P}_{x_0}(T_A = +\infty)} = \mathbb{E}_{x_0}[\mathbb{1}_B h(X_{n \wedge T_A})] = \mathbb{E}_{\mathbb{Q}}[\mathbb{1}_B]$$

which allows to identify the conditional measure $\mathbb{P}(\cdot | T_A = +\infty)$ with the h -transformed measure \mathbb{Q} .

In order to extend this result to more general probabilities we need the notion of space-time harmonic functions, namely functions $f: \mathbb{N} \times E$ such that $(f(n, X_n))_n$ is a martingale. Given a positive space-time harmonic function h we can again construct its h transform as above and in general we will obtain an inhomogeneous Markov chain, even starting from an homogeneous one.

Example 7. Consider the measure $\mathbb{P}(\cdot | X_N \in A)$ for some fixed $N > 0$ and $A \in \mathcal{E}$. Let

$$h(n, x) = \mathbb{P}_x(X_{N-n} \in A).$$

Then h is a space-time harmonic function and the corresponding h transformed measure \mathbb{Q} coincides with $\mathbb{P}(\cdot | X_N \in A)$. It describes the Markov chain $(X_n)_n$ conditioned to reach A at time N .

