Stationary Distributions, Ergodicity, and Detailed Balance

- We want to **simulate samples** from a target distribution π.
- Markov chain theory gives a powerful recipe:
  1. Construct a Markov chain with transition matrix P.
  2. Ensure the chain has **stationary distribution** π (so $\pi P = \pi$).
  3. Ensure the chain is **ergodic** (so it *converges* to π from any start).
- Then, after running long enough, the chain's states behave like draws from π.

Markov chain setup and notation

- Finite state space $\mathcal{S} = \{1, 2, \ldots, m\}$.
- Transition matrix $P = (p_{ij})$ where

$$p_{ij} = \mathbb{P}(Y_{n+1} = j \mid Y_n = i), \qquad \sum_{j=1}^{m} p_{ij} = 1, \quad p_{ij} \geq 0.$$

- A probability distribution over states is a **probability vector**

$$q = (q_1, \ldots, q_m), \qquad q_i \geq 0, \; \sum_{i=1}^{m} q_i = 1.$$

- If $q^{(n)}$ denotes the distribution of Y_n, then

$$q^{(n+1)} = q^{(n)} P \quad \Rightarrow \quad q^{(n)} = q^{(0)} P^n.$$

Definition 5.7: Stationary distribution

Definition 5.7. Suppose there exists a probability vector π such that

$$\pi P = \pi.$$

Then π is called a **stationary distribution** of the Markov chain.

- $\pi P = \pi$ means π is a **left eigenvector** of P with eigenvalue 1.
- Intuition: applying one step of the chain *does not change* the distribution.

Why the name "stationary"? (invariance over time)

If the chain is started at π, i.e. $q^{(0)} = \pi$, then

$$q^{(n)} = q^{(0)} P^n = \pi P^n.$$

But since $\pi P = \pi$, we get

$$\pi P^n = (\pi P) P^{n-1} = \pi P^{n-1} = \cdots = \pi.$$

Interpretation

Starting the chain in π makes the marginal distribution of Y_n **constant in time**:

$$\mathbb{P}(Y_n = i) = \pi_i \quad \text{for all } n.$$

Theorem 5.1: Irreducible $\Rightarrow$ unique stationary distribution

Theorem 5.1. If a finite Markov chain is **irreducible**, then a **unique** stationary distribution exists.

- **Irreducible** (recall): every state can be reached from every other state with positive probability in some number of steps.
- What this theorem does *not* say:
  - It does not guarantee that $q^{(0)} P^n \to \pi$ as $n \to \infty$.
  - It only guarantees existence and uniqueness of π.

Takeaway

Uniqueness of π is a **structural** property of the graph of transitions. Convergence needs an extra condition.

Definition 5.8: Ergodicity (convergence to stationarity)

Definition 5.8. A finite Markov chain is called **ergodic** if for all $i, j \in \mathcal{S}$,

$$\lim_{n\to\infty} \mathbb{P}(Y_n = i \mid Y_0 = j) = \pi_i.$$

- This means: **no matter where you start** $(Y_0 = j)$, the distribution at time n converges to π.

Equivalent matrix form

For any initial distribution q,

$$\lim_{n\to\infty} qP^n = \pi.$$

Irreducible is not enough: a 2-state counterexample

Consider the 2-state Markov chain with transition matrix

$$P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

- The chain deterministically alternates: $1 \to 2 \to 1 \to 2 \to \cdots$
- **Irreducible:** from 1 you can reach 2 in one step; from 2 you can reach 1 in one step.

Stationary distribution exists and is unique

Check that $\pi = \left(\frac{1}{2}, \frac{1}{2}\right)$ satisfies $\pi P = \pi$:

$$\left(\tfrac{1}{2}, \tfrac{1}{2}\right) \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \left(\tfrac{1}{2}, \tfrac{1}{2}\right).$$

But it does not converge (periodicity problem)

Start at state 1 with probability 1:

$$q^{(0)} = (1, 0).$$

Then

$$q^{(1)} = q^{(0)}P = (0, 1), \quad q^{(2)} = q^{(1)}P = (1, 0),$$

and in general

$$q^{(2n)} = (1, 0), \qquad q^{(2n+1)} = (0, 1).$$

No convergence

For example,

$$\mathbb{P}(Y_{2n} = 1) = 1 \neq \tfrac{1}{2} = \pi_1.$$

So $q^{(0)}P^n$ does **not** approach π.

- The chain keeps **oscillating** rather than mixing.

Aperiodicity: what went wrong?

The issue is **periodicity**. In the 2-state alternating chain:

- You return to state 1 only at even times: $n = 2, 4, 6, \ldots$
- You return to state 2 only at even times as well.

Period (informal)

The period of a state is the gcd of all return times with positive probability. Here both states have period 2:

$$d_1 = d_2 = 2 > 1.$$

- A chain is **aperiodic** if all states have period 1.
- Aperiodicity prevents this kind of deterministic "cycling".

Theorem 5.2: Characterisation of ergodicity

Theorem 5.2. A finite Markov chain is **ergodic** if and only if it is

- **irreducible**, and
- **aperiodic**.

Meaning

If the chain is irreducible + aperiodic, then for any start q,

$$qP^n \longrightarrow \pi,$$

where π is the unique stationary distribution.

- This is the key convergence guarantee used in MCMC.

MCMC interpretation: why ergodicity matters

- In MCMC we want to sample from a **target distribution** π.
- We build a Markov chain whose stationary distribution is π.
- If the chain is ergodic, then:
  - regardless of initialisation, Y_n eventually behaves like a draw from π;
  - empirical averages along the chain approximate expectations under π.

Practical wording

It may take a while to "reach stationarity" (burn-in / mixing), but in principle an ergodic chain will get there.

Checklist for convergence to the stationary distribution

To ensure a finite Markov chain converges to its stationary distribution, check:

1. **Irreducible** (can reach everywhere).
2. **Aperiodic** (does not get trapped in cycles).

If you also want to sample from a *prescribed* target π, you additionally need:

3. π **is stationary:** $\pi P = \pi$.

Two separate design goals

(1)–(2) guarantee convergence *to something*; (3) ensures the "something" is the distribution you want.

Example 5.2: King Markov and prescribed stationarity

Recall Example 5.1 (the King and the ring of islands):

- States are islands $i = 1, \dots, 10$.
- Target long-run behaviour: spend time proportional to island size:

$$\mathbb{P}(Y = i) \propto i.$$

- The King's proposal-and-accept rule was designed so that:
  - the chain is **aperiodic** (not stuck in a regular alternating cycle);
  - the chain is **irreducible** (can reach all islands);
  - the resulting stationary distribution matches $\pi(i) \propto i$.

Designing irreducible + aperiodic chains is often straightforward. Designing a chain with **a prescribed stationary distribution** is the subtle part.

How to enforce a desired stationary distribution?

Suppose we want a chain with stationary distribution π.

- Directly solving $\pi P = \pi$ can be awkward, because P must also:
  - have nonnegative entries,
  - have rows summing to 1,
  - satisfy irreducible + aperiodic structure.
- A very useful sufficient condition is **detailed balance**.

Key idea

Instead of enforcing the global condition $\pi P = \pi$, we enforce a *pairwise symmetry of flow* between every pair of states.

Definition 5.9: Detailed balance

Definition 5.9. A Markov chain with transition matrix P satisfies the **detailed balance equation** with respect to π if for all states i, j,

$$\pi_i p_{ij} = \pi_j p_{ji}.$$

- $\pi_i p_{ij}$ is the probability mass *flowing* from i to j in one step if the chain is distributed as π.
- Detailed balance says: for each pair (i, j) the flow $i \to j$ equals the flow $j \to i$.

Interpretation

No net flow anywhere $\Rightarrow$ distribution stays unchanged $\Rightarrow$ stationarity.

Theorem 5.3: Detailed balance $\Rightarrow$ stationarity

Theorem 5.3 (Detailed Balance). Let P satisfy detailed balance with respect to π. Then

$$\pi P = \pi,$$

so π is stationary.

- This theorem provides a convenient way to **construct** Markov chains with a known stationary distribution.
- Many MCMC algorithms (e.g. Metropolis–Hastings) are designed by enforcing detailed balance.

Proof of Theorem 5.3 (line by line)

We prove that the jth entry of πP equals π_j.

The jth entry is:

$$(\pi P)_j = \sum_i \pi_i p_{ij}.$$

Using detailed balance $\pi_i p_{ij} = \pi_j p_{ji}$:

$$\sum_i \pi_i p_{ij} = \sum_i \pi_j p_{ji} = \pi_j \sum_i p_{ji}.$$

Now use that probabilities out of state j sum to 1:

$$\sum_i p_{ji} = 1 \quad \Rightarrow \quad (\pi P)_j = \pi_j.$$

Conclusion

Since this holds for every j, we have $\pi P = \pi$.

Putting it together: the MCMC design recipe

To sample from a target distribution π using a Markov chain:

1. Choose a state space $\mathcal{S}$ and a transition mechanism.
2. Ensure **irreducibility** (connect the space).
3. Ensure **aperiodicity** (avoid deterministic cycles).
4. Ensure **stationarity** of π:
    - often via **detailed balance**: $\pi_i p_{ij} = \pi_j p_{ji}$.

Guarantee

Then the chain is ergodic and

$$qP^n \to \pi \quad \text{for any initial } q,$$

so long-run samples behave like draws from π.

A quick comment: why do this for discrete π?

So far, everything was for **finite** state Markov chains.

- If π is a discrete distribution on a finite set, we can already sample from it directly (e.g. inverse transform / alias method).
- So using a Markov chain can feel like **overkill** in the finite-discrete case.

Why we still study it

The same concepts generalise naturally to:

- countably infinite state spaces,
- continuous state spaces (e.g. $\mathbb{R}^d$),
- complicated posteriors in Bayesian inference.

That is where MCMC becomes essential.

Transition to general state spaces (preview)

- In continuous or high-dimensional problems we often cannot:
  - write down the normalising constant of π,
  - sample from π directly.
- We therefore construct a Markov chain whose stationary distribution is π and rely on ergodic convergence.

Next

We will extend these ideas from finite transition matrices to **Markov kernels** and algorithms like **Metropolis–Hastings**.