

0.1 Learning Objectives

By the end of this lecture, you will be able to:

- Define stochastic processes and state spaces
- State the Markov property and explain its significance
- Construct transition matrices from verbal descriptions
- Compute n -step transition probabilities via matrix powers
- Find the steady-state distribution of a Markov chain

1. Stochastic Processes

1.1 What Is a Stochastic Process?

Definition: Stochastic Process

A **stochastic process** is a collection of random variables $\{X_1, X_2, X_3, \dots\}$ indexed by time (or space).

1.1 What Is a Stochastic Process?

Definition: Stochastic Process

A **stochastic process** is a collection of random variables $\{X_1, X_2, X_3, \dots\}$ indexed by time (or space).

- Can be **discrete** or **continuous** in time
- Can be **discrete** or **continuous** in state space
- Today: discrete time, discrete state space

1.2 States and State Spaces

The values that the random variables X_1, X_2, X_3, \dots can take are called **states**.

1.2 States and State Spaces

The values that the random variables X_1, X_2, X_3, \dots can take are called **states**.

The **state space** S is the set of all possible states.

1.2 States and State Spaces

The values that the random variables X_1, X_2, X_3, \dots can take are called **states**.

The **state space** S is the set of all possible states.

Example: Coin Flip Process

Flip a fair coin 100 times. Let X_t = outcome of flip t .

- State space: $S = \{H, T\}$
- Each X_t is a random variable taking values in S

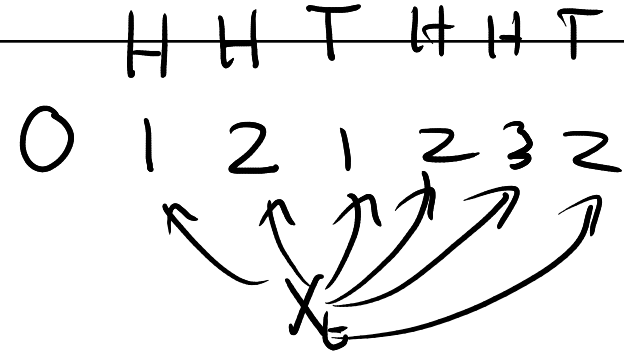
1.3 Example: Random Walk

Example: Random Walk

Flip a fair coin 100 times. Keep a running score:

- Heads: +1
- Tails: -1

Let X_t = score after t flips.



1.3 Example: Random Walk

Example: Random Walk

Flip a fair coin 100 times. Keep a running score:

- Heads: +1
- Tails: -1

Let X_t = score after t flips.

What's the state space? $[-100, 100]$

1.3 Example: Random Walk

Example: Random Walk

Flip a fair coin 100 times. Keep a running score:

- Heads: +1
- Tails: -1

Let X_t = score after t flips.

What's the state space?

$$S = \{-100, -99, \dots, 99, 100\}$$

1.3 Example: Random Walk

Example: Random Walk

Flip a fair coin 100 times. Keep a running score:

- Heads: +1
- Tails: -1

Let X_t = score after t flips.


What's the state space?

$$S = \{-100, -99, \dots, 99, 100\}$$

This is a **random walk**


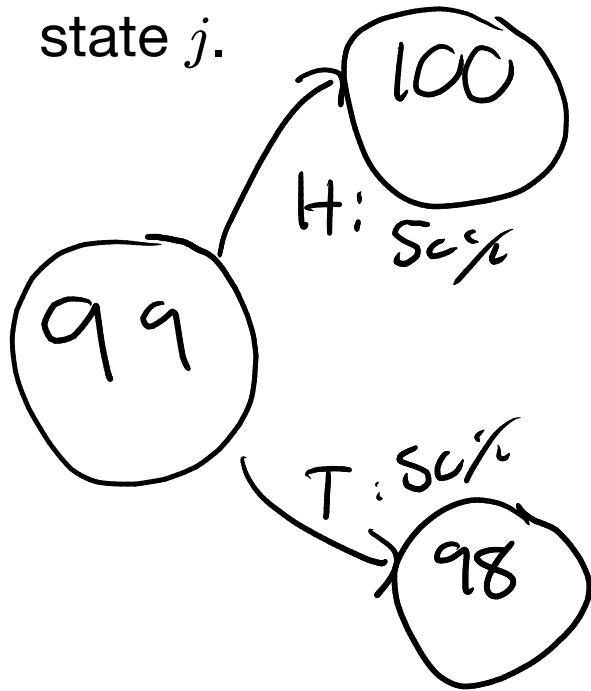
1.3 Example: Random Walk

one of the simplest stochastic processes.



1.4 Evolution of a Stochastic Process

If $X_1 = i$ and $X_2 = j$, we say the process made a transition from state i to state j .



1.4 Evolution of a Stochastic Process

If $X_1 = i$ and $X_2 = j$, we say the process made a **transition** from state i to state j .

Typically, we care about **long-run behavior**: after many transitions, what are the probabilities of being in each state?

1.4 Evolution of a Stochastic Process

If $X_1 = i$ and $X_2 = j$, we say the process made a **transition** from state i to state j .

Typically, we care about **long-run behavior**: after many transitions, what are the probabilities of being in each state?

Note: In general, the future of a stochastic process can depend on its entire history. This makes analysis very hard. Can we simplify?

1.5 Your Day as a Stochastic Process

Your daily routine is a stochastic process!

1.5 Your Day as a Stochastic Process

Your daily routine is a stochastic process!

- **States:** Wake up, Breakfast, Coffee, Lecture, Studying, Nap, Break, ...
- **Transitions:** probabilistic moves between activities

1.5 Your Day as a Stochastic Process

Your daily routine is a stochastic process!

- **States:** Wake up, Breakfast, Coffee, Lecture, Studying, Nap, Break, ...
- **Transitions:** probabilistic moves between activities

Some states are **transient** (startup behavior), some are **recurrent** (steady-state), and some are **absorbing** (done for today).

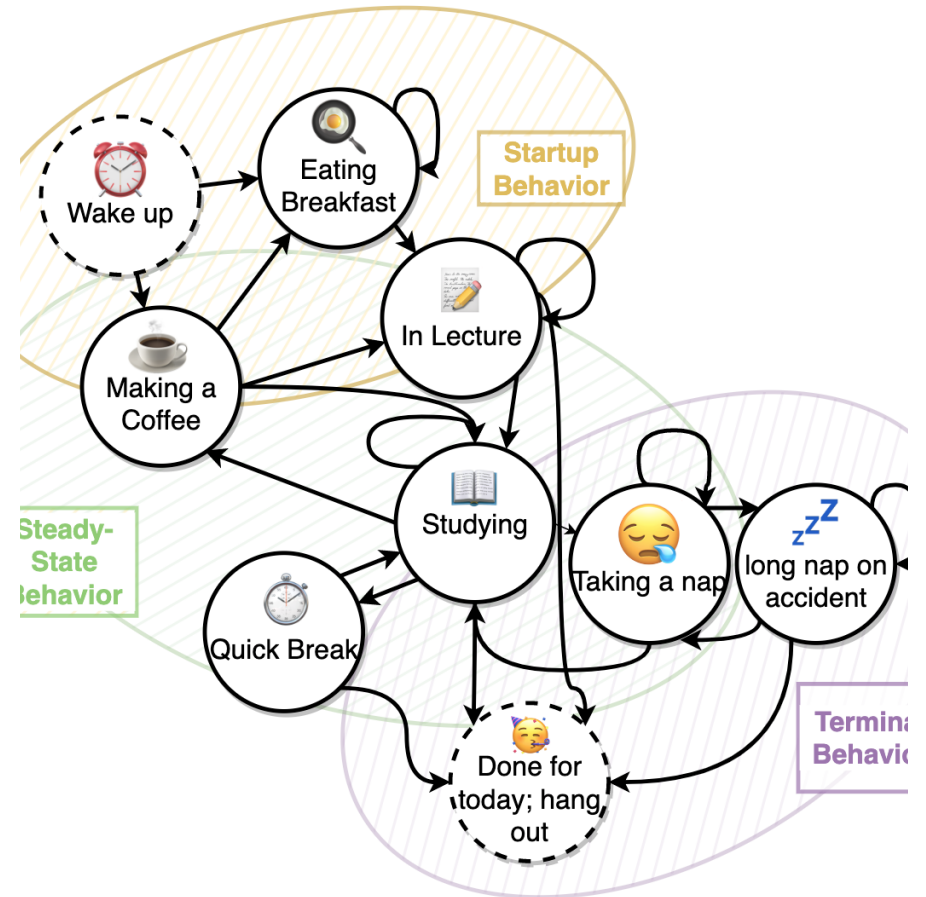
1.5 Your Day as a Stochastic Process

Your daily routine is a stochastic process!

- **States:** Wake up, Breakfast, Coffee, Lecture, Studying, Nap, Break, ...
- **Transitions:** probabilistic moves between activities

Some states are **transient** (startup behavior), some are **recurrent** (steady-state), and some are **absorbing** (done for today).

Are the transition probabilities related to long-term academic success?



2. Markov Chains

2.1 The Key Simplification

Question: What if the future only depends on the **present**, not the past?

2.1 The Key Simplification

Question: What if the future only depends on the **present**, not the past?

Definition: Markov Property
A stochastic process $\{X_1, X_2, X_3, \dots\}$ has the **Markov property** if

$$P(\underbrace{X_{n+1} = j}_{\text{next state}} \mid \underbrace{X_n = i_n, X_{n-1} = i_{n-1}, \dots, X_1 = i_1}_{\text{entire history}}) = P(\underbrace{X_{n+1} = j}_{\text{next state}} \mid \underbrace{X_n = i_n}_{\text{present state}})$$

2.1 The Key Simplification

Question: What if the future only depends on the **present**, not the past?

Definition: Markov Property

A stochastic process $\{X_1, X_2, X_3, \dots\}$ has the **Markov property** if

$$P(X_{n+1} = j \mid X_n = i_n, X_{n-1} = i_{n-1}, \dots, X_1 = i_1) = P(X_{n+1} = j \mid X_n = i_n)$$

The next state depends *only* on the current state

2.1 The Key Simplification

not on how we got here.

2.2 Markov Chains

Definition: Markov Chain

A **Markov chain** is a stochastic process with the Markov property and a finite (or countable) state space.

2.2 Markov Chains

Definition: Markov Chain

A **Markov chain** is a stochastic process with the Markov property and a finite (or countable) state space.

It captures **one-step dependence**: the future depends on the present, but not the past.

2.2 Markov Chains

Definition: Markov Chain

A **Markov chain** is a stochastic process with the Markov property and a finite (or countable) state space.

It captures **one-step dependence**: the future depends on the present, but not the past.

We are interested in the **long-run behavior**: once the process no longer “remembers” its starting state.

2.3 Example: Weather


Example: Island Weather

On any day, the weather is sunny (S) or rainy (R).

$$X_n = \begin{cases} 0 & \text{if sunny} \\ 1 & \text{if rainy} \end{cases}$$

Once a pattern begins, it tends to continue:

- If today is sunny, probability 0.97 that tomorrow is sunny
- If today is rainy, probability 0.94 that tomorrow is rainy



0 1 1 0 1 0

2.4 Weather: Transition Probabilities

Today	Tomorrow	Notation	Probability
Sunny	Sunny	$P(X_{n+1} = 0 \mid X_n = 0)$	0.97
Sunny	Rainy	$P(X_{n+1} = 1 \mid X_n = 0)$	0.03
Rainy	Sunny	$P(X_{n+1} = 0 \mid X_n = 1)$	0.06
Rainy	Rainy	$P(X_{n+1} = 1 \mid X_n = 1)$	0.94


2.4 Weather: Transition Probabilities

Today	Tomorrow	Notation	Probability
Sunny	Sunny	$P(X_{n+1} = 0 \mid X_n = 0)$	0.97
Sunny	Rainy	$P(X_{n+1} = 1 \mid X_n = 0)$	0.03
Rainy	Sunny	$P(X_{n+1} = 0 \mid X_n = 1)$	0.06
Rainy	Rainy	$P(X_{n+1} = 1 \mid X_n = 1)$	0.94

Notice: each pair of rows sums to 1 (from any state, we must go *somewhere*).

2.5 Weather: State Diagram

We can also visualize the transitions as a **state diagram**:




Each node is a state, each arrow is a transition with its probability. Self-loops represent staying in the same state.

3. Transition Matrices

3.1 Transition Matrices

Definition: Transition Matrix

The **transition matrix** P of a Markov chain has entry P_{ij} = probability of going from state i to state j in one step.



3.1 Transition Matrices

Definition: Transition Matrix

The **transition matrix** P of a Markov chain has entry P_{ij} = probability of going from state i to state j in one step.

For the weather example:

$$P = \begin{matrix} & \begin{matrix} 0 & 1 \end{matrix} \\ \begin{matrix} 0 \\ 1 \end{matrix} & \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \end{matrix}$$

row-stochastic

3.1 Transition Matrices

Definition: Transition Matrix

The **transition matrix** P of a Markov chain has entry P_{ij} = probability of going from state i to state j in one step.

For the weather example:

$$P = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix}$$

Properties:

- All entries are between 0 and 1
- Each row sums to 1 (each row is a probability distribution)

3.2 Try It: Build a Transition Matrix

Example: Social Media

A user is on one of three pages: **Home** (H), **Profile** (P), or **Search** (S).

- From Home: 60% stay, 30% go to Profile, 10% go to Search
- From Profile: 40% go to Home, 40% stay, 20% go to Search
- From Search: 30% go to Home, 20% go to Profile, 50% stay

Write the transition matrix.

$$X_t \begin{matrix} & \begin{matrix} H & P & S \end{matrix} \\ \begin{matrix} H \\ P \\ S \end{matrix} & \begin{bmatrix} 0.6 & 0.3 & 0.1 \\ 0.4 & 0.4 & 0.2 \\ 0.3 & 0.2 & 0.5 \end{bmatrix} \end{matrix}$$

3.2 Try It: Build a Transition Matrix

Example: Social Media

A user is on one of three pages: **Home** (H), **Profile** (P), or **Search** (S).

- From Home: 60% stay, 30% go to Profile, 10% go to Search
- From Profile: 40% go to Home, 40% stay, 20% go to Search
- From Search: 30% go to Home, 20% go to Profile, 50% stay

Write the transition matrix.

Try it yourself

Talk to your neighbor and try to solve this problem.

3.3 Solution: Social Media Transition Matrix

$$P = \begin{pmatrix} 0.6 & 0.3 & 0.1 \\ 0.4 & 0.4 & 0.2 \\ 0.3 & 0.2 & 0.5 \end{pmatrix}$$


3.3 Solution: Social Media Transition Matrix

$$P = \begin{pmatrix} 0.6 & 0.3 & 0.1 \\ 0.4 & 0.4 & 0.2 \\ 0.3 & 0.2 & 0.5 \end{pmatrix}$$

Reading row by row: $P_{HH} = 0.6$, $P_{HP} = 0.3$, $P_{HS} = 0.1$, etc.

3.3 Solution: Social Media Transition Matrix

$$P = \begin{matrix} & \begin{matrix} H & P & S \end{matrix} \\ \begin{matrix} H \\ P \\ S \end{matrix} & \begin{pmatrix} 0.6 & 0.3 & 0.1 \\ 0.4 & 0.4 & 0.2 \\ 0.3 & 0.2 & 0.5 \end{pmatrix} \end{matrix}$$



Reading row by row: $P_{HH} = 0.6$, $P_{HP} = 0.3$, $P_{HS} = 0.1$, etc.

Note: Each row sums to 1. We can think of each row as the conditional distribution of the next state, given the current state.

$$P(X_2 = P \mid X_0 = H) = 0.6 \times 0.3 + 0.1 \times 0.2 + 0.3 \times 0.4$$

4. Multi-Step Transitions

4.1 n -Step Transition Probabilities

Definition: ~~n~~ -Step Transition Probability

The probability of going from state i to state j in exactly n steps:

$$P_{ij}^{(n)} = P(X_n = j \mid X_0 = i)$$

$$= P(X_{t+n} = j \mid X_t = i)$$

4.1 n -Step Transition Probabilities

Definition: n -Step Transition Probability

The probability of going from state i to state j in exactly n steps:

$$P_{ij}^{(n)} = P(X_n = j \mid X_0 = i)$$

Key fact: The n -step transition matrix is $P^{(n)}$ (the matrix P raised to the n -th power).

$$P_{ij}^2$$

4.1 n -Step Transition Probabilities

Definition: n -Step Transition Probability

The probability of going from state i to state j in exactly n steps:

$$P_{ij}^{(n)} = P(X_n = j \mid X_0 = i)$$

Key fact: The n -step transition matrix is P^n (the matrix P raised to the n -th power).

Row i of P^n gives the distribution of X_n given that the chain started in state i .

4.2 Try It: Two-Step Weather Transitions

Example: Weather After 2 Days

Using the weather transition matrix $P = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix}$:

If today is sunny, what is the probability it's rainy in 2 days?

4.2 Try It: Two-Step Weather Transitions

Example: Weather After 2 Days

Using the weather transition matrix $P = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix}$:

If today is sunny, what is the probability it's rainy in 2 days?

$$P^2_{SR}$$

Try it yourself

Talk to your neighbor and try to solve this problem.

4.3 Solution: Two-Step Weather

$$P^2 = P \cdot P = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \cdot \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix}$$

4.3 Solution: Two-Step Weather

$$\begin{aligned} P^2 &= P \cdot P = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \cdot \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \\ &= \begin{pmatrix} 0.97 \cdot 0.97 + 0.03 \cdot 0.06 & 0.97 \cdot 0.03 + 0.03 \cdot 0.94 \\ 0.06 \cdot 0.97 + 0.94 \cdot 0.06 & 0.06 \cdot 0.03 + 0.94 \cdot 0.94 \end{pmatrix} \end{aligned}$$

4.3 Solution: Two-Step Weather

$$\begin{aligned} P^2 &= P \cdot P = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \cdot \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \\ &= \begin{pmatrix} 0.97 \cdot 0.97 + 0.03 \cdot 0.06 & 0.97 \cdot 0.03 + 0.03 \cdot 0.94 \\ 0.06 \cdot 0.97 + 0.94 \cdot 0.06 & 0.06 \cdot 0.03 + 0.94 \cdot 0.94 \end{pmatrix} \\ &= \begin{pmatrix} 0.9427 & 0.0573 \\ 0.1146 & 0.8854 \end{pmatrix} \end{aligned}$$

4.3 Solution: Two-Step Weather

$$\begin{aligned} P^2 &= P \cdot P = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \cdot \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \\ &= \begin{pmatrix} 0.97 \cdot 0.97 + 0.03 \cdot 0.06 & 0.97 \cdot 0.03 + 0.03 \cdot 0.94 \\ 0.06 \cdot 0.97 + 0.94 \cdot 0.06 & 0.06 \cdot 0.03 + 0.94 \cdot 0.94 \end{pmatrix} \\ &= \begin{matrix} & \begin{matrix} S & R \end{matrix} \\ \begin{matrix} S \\ R \end{matrix} & \begin{pmatrix} 0.9427 & 0.0573 \\ 0.1146 & 0.8854 \end{pmatrix} \end{matrix} \end{aligned}$$

$$P_{0,1}^{(2)} = 0.0573$$

4.3 Solution: Two-Step Weather

about a 5.7% chance of rain in 2 days given sun today.

5. Steady-State Distribution

5.1 Long-Run Behavior

What happens as $n \rightarrow \infty$? Do the rows of P^n converge?

5.1 Long-Run Behavior

What happens as $n \rightarrow \infty$? Do the rows of P^n converge?

Example: Weather: Powers of P

$$P^1 = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \quad P^{10} \approx \begin{pmatrix} 0.76 & 0.24 \\ 0.48 & 0.52 \end{pmatrix}$$
$$P^{50} \approx \begin{pmatrix} 0.667 & 0.333 \\ 0.667 & 0.333 \end{pmatrix} \quad P^{100} \approx \begin{pmatrix} 0.667 & 0.333 \\ 0.667 & 0.333 \end{pmatrix}$$

5.1 Long-Run Behavior

What happens as $n \rightarrow \infty$? Do the rows of P^n converge?

Example: Weather: Powers of P

$$P^1 = \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} \quad P^{10} \approx \begin{pmatrix} 0.76 & 0.24 \\ 0.48 & 0.52 \end{pmatrix}$$
$$P^{50} \approx \begin{pmatrix} 0.667 & 0.333 \\ 0.667 & 0.333 \end{pmatrix} \quad P^{100} \approx \begin{pmatrix} 0.667 & 0.333 \\ 0.667 & 0.333 \end{pmatrix}$$

The rows converge to the **same** distribution! The starting state no longer matters.

5.2 Steady-State Distribution

Definition: Steady-State (Stationary) Distribution

A probability distribution π is a **steady-state distribution** if it satisfies:

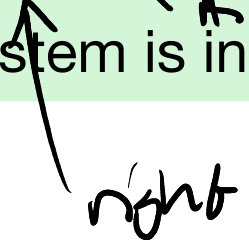
$$\pi P = \pi$$

That is, π is unchanged by the transition — the system is in equilibrium.

5.2 Steady-State Distribution

Definition: Steady-State (Stationary) Distribution

A probability distribution π is a **steady-state distribution** if it satisfies:

$$\underline{\pi}P = \underline{\pi} \cdot 1 \quad \underline{A} \underline{V} = \lambda \cdot \underline{V}$$


That is, π is unchanged by the transition — the system is in equilibrium.

π is a **left eigenvector** of P with eigenvalue 1.

5.2 Steady-State Distribution

Definition: Steady-State (Stationary) Distribution

A probability distribution π is a **steady-state distribution** if it satisfies:

$$\pi P = \pi$$

That is, π is unchanged by the transition — the system is in equilibrium.

π is a **left eigenvector** of P with eigenvalue 1.

Recall: An eigenvector of a matrix A satisfies $Av = \lambda v$. Here we need $\pi P = 1 \cdot \pi$.

5.3 Finding the Steady-State Distribution

Three approaches:

5.3 Finding the Steady-State Distribution

Three approaches:

1. **Simulate**: run the chain for many steps, record the fraction of time in each state

5.3 Finding the Steady-State Distribution

Three approaches:

1. **Simulate**: run the chain for many steps, record the fraction of time in each state
2. **Matrix powers**: compute P^n for large n and read off any row

5.3 Finding the Steady-State Distribution

Three approaches:

1. **Simulate**: run the chain for many steps, record the fraction of time in each state
2. **Matrix powers**: compute P^n for large n and read off any row
3. **Solve directly**: solve $\pi P = \pi$ subject to $\sum_i \pi_i = 1$

5.4 Try It: Weather Steady-State

Example: Weather Steady-State

Find $\pi = (\pi_0, \pi_1)$ such that $\pi P = \pi$ and $\pi_0 + \pi_1 = 1$.

$P(S)$ \nearrow
 $P(R)$ \nearrow

$$(\pi_0, \pi_1) \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} = (\pi_0, \pi_1)$$

5.4 Try It: Weather Steady-State

Example: Weather Steady-State

Find $\pi = (\pi_0, \pi_1)$ such that $\pi P = \pi$ and $\pi_0 + \pi_1 = 1$.

$$(\pi_0, \pi_1) \begin{pmatrix} 0.97 & 0.03 \\ 0.06 & 0.94 \end{pmatrix} = (\pi_0, \pi_1)$$

Try it yourself

Talk to your neighbor and try to solve this problem.

5.5 Solution: Weather Steady-State

From $\pi P = \pi$:

$$0.97\pi_0 + 0.06\pi_1 = \pi_0$$

5.5 Solution: Weather Steady-State

From $\pi P = \pi$:

$$0.97\pi_0 + 0.06\pi_1 = \pi_0$$

$$\rightarrow -0.03\pi_0 + 0.06\pi_1 = 0 \rightarrow \pi_1 = \frac{\pi_0}{2}$$

5.5 Solution: Weather Steady-State

From $\pi P = \pi$:

$$0.97\pi_0 + 0.06\pi_1 = \pi_0$$

$$\rightarrow -0.03\pi_0 + 0.06\pi_1 = 0 \rightarrow \pi_1 = \frac{\pi_0}{2}$$

Using $\pi_0 + \pi_1 = 1$: $\pi_0 + \frac{\pi_0}{2} = 1 \rightarrow \pi_0 = 2/3$

5.5 Solution: Weather Steady-State

From $\pi P = \pi$:

$$0.97\pi_0 + 0.06\pi_1 = \pi_0$$

$$\rightarrow -0.03\pi_0 + 0.06\pi_1 = 0 \rightarrow \pi_1 = \frac{\pi_0}{2}$$

Using $\pi_0 + \pi_1 = 1$: $\pi_0 + \frac{\pi_0}{2} = 1 \rightarrow \pi_0 = 2/3$

$\sum \pi_i = 1$
normalization

$$\pi = (2/3, 1/3) = (0.667, 0.333)$$

5.5 Solution: Weather Steady-State

From $\pi P = \pi$:

$$0.97\pi_0 + 0.06\pi_1 = \pi_0$$

$$\rightarrow -0.03\pi_0 + 0.06\pi_1 = 0 \rightarrow \pi_1 = \frac{\pi_0}{2}$$

Using $\pi_0 + \pi_1 = 1$: $\pi_0 + \frac{\pi_0}{2} = 1 \rightarrow \pi_0 = 2/3$

$$\pi = (2/3, 1/3) = (0.667, 0.333)$$

In the long run, about 67% of days are sunny and 33% are rainy

5.5 Solution: Weather Steady-State

regardless of today's weather.

6. Properties of Markov Chains

6.1 When Does Convergence Happen?

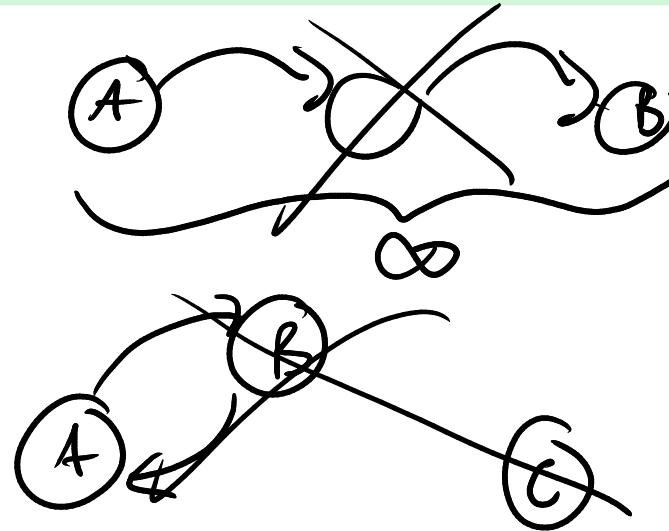
Not all Markov chains converge to a steady state. Two conditions guarantee it:

6.1 When Does Convergence Happen?

Not all Markov chains converge to a steady state. Two conditions guarantee it:

Definition: Irreducibility

A Markov chain is **irreducible** if it is possible to reach any state from any other state in a finite number of steps.



6.1 When Does Convergence Happen?


Not all Markov chains converge to a steady state. Two conditions guarantee it:

Definition: Irreducibility

A Markov chain is **irreducible** if it is possible to reach any state from any other state in a finite number of steps.

Definition: Aperiodicity

A state i has **period** d if returns to i must occur in multiples of d steps. A chain is **aperiodic** if all states have period 1.



6.2 Convergence Theorem

If a Markov chain is **irreducible** and **aperiodic**, then:

$$P_{ij}^n \rightarrow \pi_j \text{ as } n \rightarrow \infty$$

for all states i and j .

6.2 Convergence Theorem

If a Markov chain is **irreducible** and **aperiodic**, then:

$$P_{ij}^n \rightarrow \pi_j \text{ as } n \rightarrow \infty$$

for all states i and j .

Furthermore, π is **unique** and satisfies:

- $\pi P = \pi$ (stationarity)
- $\sum_i \pi_i = 1$ (valid probability distribution)
- $\pi_i =$ expected long-run fraction of time spent in state i

6.3 Properties of the Stationary Distribution

Uniqueness: π is the unique distribution satisfying $\pi P = \pi$.

6.3 Properties of the Stationary Distribution

Uniqueness: π is the unique distribution satisfying $\pi P = \pi$.

Steady-state: If X_n has distribution π , then X_m also has distribution π for all $m \geq n$.

6.3 Properties of the Stationary Distribution

Uniqueness: π is the unique distribution satisfying $\pi P = \pi$.

Steady-state: If X_n has distribution π , then X_m also has distribution π for all $m \geq n$.

Long-run proportion: π_i is the expected fraction of time the chain spends in state i as $n \rightarrow \infty$.

6.3 Properties of the Stationary Distribution

Uniqueness: π is the unique distribution satisfying $\pi P = \pi$.

Steady-state: If X_n has distribution π , then X_m also has distribution π for all $m \geq n$.

Long-run proportion: π_i is the expected fraction of time the chain spends in state i as $n \rightarrow \infty$.

Note: These properties make Markov chains a powerful modeling tool. Many real-world processes — web browsing, weather, queues, genetics — are well-approximated by Markov chains.

Prague

6.4 Recap

Today we covered:

- A **stochastic process** is a collection of random variables indexed by time
- The **Markov property**: the future depends only on the present, not the past
- **Transition matrices** encode one-step probabilities; P^n gives n -step probabilities
- The **steady-state distribution** π satisfies $\pi P = \pi$ — it's where the chain “settles”
- Irreducible + aperiodic chains converge to a unique steady state