Monopoly – An Analysis using Markov Chains

Benjamin Bernard

Columbia University, New York

PRIME event at Pace University, New York

April 19, 2017

Introduction

Applications of Markov chains

Applications of Markov chains:

- Equilibrium distribution in networks: page rank in search engines,
- Speech analysis: voice recognition, word predictions on keyboard,
- Agriculture: population growth and harvesting,
- Research in related fields: modelling of random systems in physics and finance, sampling from arbitrary distributions
- Markov decision processes in artificial intelligence: sequential decision problems under uncertainty, reinforcement learning,
- Games: compute intricate scenarios in a fairly simple way.

Applications of Markov chains

Applications of Markov chains:

- Equilibrium distribution in networks: page rank in search engines,
- Speech analysis: voice recognition, word predictions on keyboard,
- Agriculture: population growth and harvesting,
- Research in related fields: modelling of random systems in physics and finance, sampling from arbitrary distributions
- Markov decision processes in artificial intelligence: sequential decision problems under uncertainty, reinforcement learning,
- Games: compute intricate scenarios in a fairly simple way.

Skills required to use Markov chains:

- Knowledge of the system we are modelling,
- Some basic programming skills.

Which squares are the best?

Markov chains allow us to:

- Calculate the probability distribution of pieces on the board that arises after sufficiently long play.
- Compute expected income/return of properties and their developments.
- Evaluate the quality of a trade by its impact on the ruin probability.

Distribution after the first turn

Rules

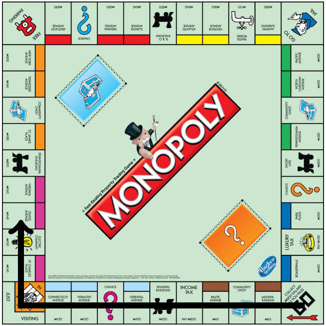

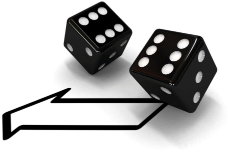

- Roll dice to advance on board.

Rules

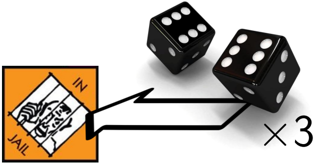

- Roll dice to advance on board. If you roll 3 doubles in a row $\rightarrow$ jail.

Rules

- Roll dice to advance on board. If you roll 3 doubles in a row $\rightarrow$ jail.
- Properties can be bought or traded. Visitors pay “rent”.

Rules

- Roll dice to advance on board. If you roll 3 doubles in a row $\rightarrow$ jail.
- Properties can be bought or traded. Visitors pay “rent”.
- If all properties of one color are owned, they can be developed for a substantial increase of rent.

Rules

- Roll dice to advance on board. If you roll 3 doubles in a row $\rightarrow$ jail.
- Properties can be bought or traded. Visitors pay "rent".
- If all properties of one color are owned, they can be developed for a substantial increase of rent.
- Players who cannot afford rent are eliminated. Last remaining player wins.

Probability distribution of the sum of two dice

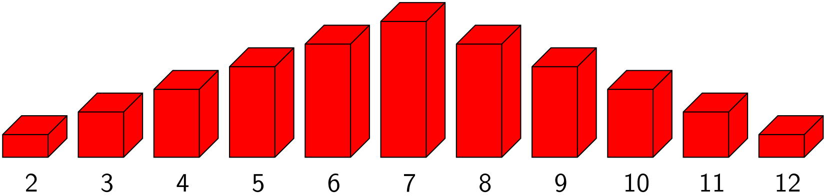

Each combination of dice has a probability of $\frac{1}{6} \cdot \frac{1}{6} = \frac{1}{36}$.

- 2 is attained by ⚀⚀
- 3 is attained by ⚀⚁ and ⚁⚀
- 4 is attained by ⚀⚂, ⚁⚁ and ⚂⚀
- 5 is attained by ⚀⚃, ⚁⚂, ⚂⚁ and ⚃⚀

Advancing on the monopoly board

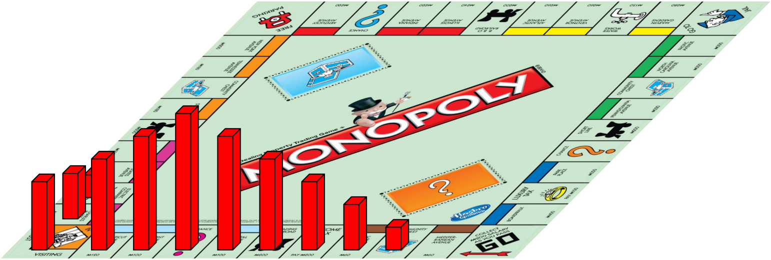

Three possible ways of moving your piece:

- Rolling the dice,

Advancing on the monopoly board

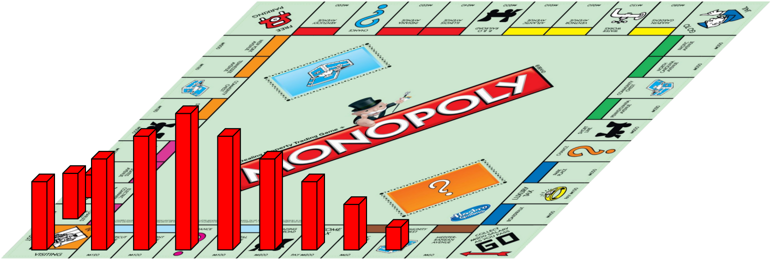

Three possible ways of moving your piece:

- Rolling the dice,
- Being sent to jail,

Advancing on the monopoly board

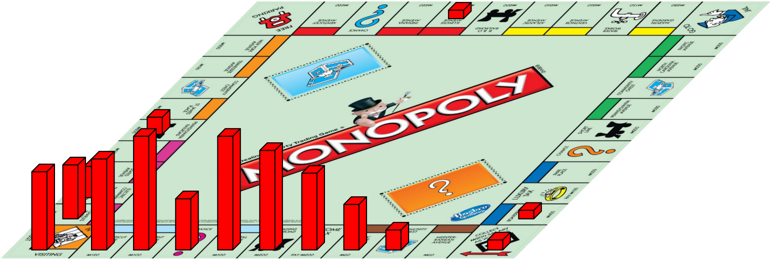

Three possible ways of moving your piece:

- Rolling the dice,
- Being sent to jail,
- Community chest and chance cards.

Steady state distribution

Finite state Markov chain

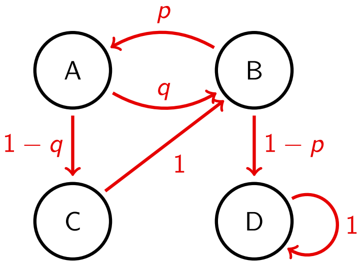

A Markov chain transitions randomly between several possible states:

- Sum of transition probabilities is 1.

Finite state Markov chain

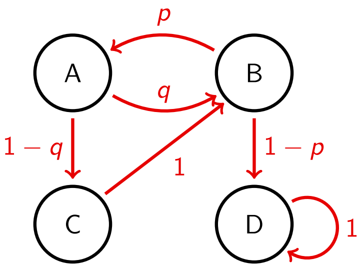

$$\begin{pmatrix} 0 & p & 0 & 0 \\ q & 0 & 1 & 0 \\ 1-q & 0 & 0 & 0 \\ 0 & 1-p & 0 & 1 \end{pmatrix}$$

A Markov chain transitions randomly between several possible states:

- Sum of transition probabilities is 1.
- Representation by a matrix $P = (p_{ij})$: Element p_{ij} in row i and column j is the probability to transition from state j to state i.

Finite state Markov chain

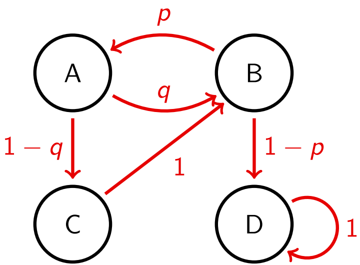

$$\begin{pmatrix} 0 & p & 0 & 0 \\ q & 0 & 1 & 0 \\ 1-q & 0 & 0 & 0 \\ 0 & 1-p & 0 & 1 \end{pmatrix}$$

A Markov chain transitions randomly between several possible states:

- Sum of transition probabilities is 1.
- Representation by a matrix $P = (p_{ij})$: Element p_{ij} in row i and column j is the probability to transition from state j to state i.
- n-step transition probability given by P^n (matrix multiplication).

Finite state Markov chain

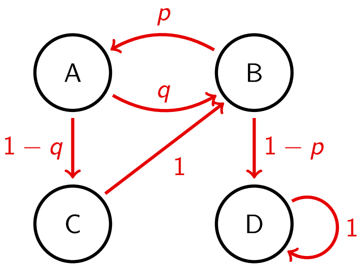

$$\begin{pmatrix} 0 & p & 0 & 0 \\ q & 0 & 1 & 0 \\ 1-q & 0 & 0 & 0 \\ 0 & 1-p & 0 & 1 \end{pmatrix}$$

Example: Initial state is A

- The initial distribution equals $\pi_0 = (1, 0, 0, 0)^\top$.

Finite state Markov chain

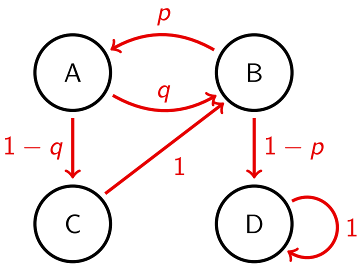

$$\begin{pmatrix} 0 & p & 0 & 0 \\ q & 0 & 1 & 0 \\ 1-q & 0 & 0 & 0 \\ 0 & 1-p & 0 & 1 \end{pmatrix}$$

Example: Initial state is A

- The initial distribution equals $\pi_0 = (1, 0, 0, 0)^\top$.
- After one step, the distribution equals $\pi_1 = (0, q, 1-q, 0)^\top$. This can also be obtained by matrix multiplication: $\pi_1 = P \cdot \pi_0$.

Finite state Markov chain

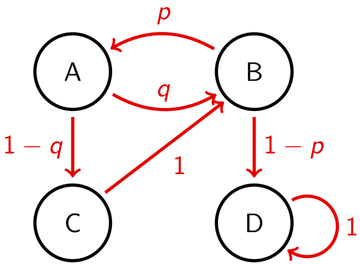

$$\begin{pmatrix} 0 & p & 0 & 0 \\ q & 0 & 1 & 0 \\ 1-q & 0 & 0 & 0 \\ 0 & 1-p & 0 & 1 \end{pmatrix}$$

Example: Initial state is A

- The initial distribution equals $\pi_0 = (1, 0, 0, 0)^\top$.
- After one step, the distribution equals $\pi_1 = (0, q, 1-q, 0)^\top$. This can also be obtained by matrix multiplication: $\pi_1 = P \cdot \pi_0$.
- After two steps: $\pi_2 = P \cdot \pi_1 = P^2 \cdot \pi_0 = \big(qp, 1-q, 0, q(1-p)\big)^\top$.

Convergence of distribution?

Does there exist a limiting distribution $\pi_\infty = \lim_{n \to \infty} P^n \cdot \pi_0$?

Convergence of distribution?

Does there exist a limiting distribution $\pi_\infty = \lim_{n\to\infty} P^n \cdot \pi_0$?

If such a distribution exists, it has to satisfy $\pi_\infty = P \cdot \pi_\infty$. That is, it is a so-called *steady state distribution*.

Theorem

If a steady state distribution exists, it is unique.

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state.

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state. ✓

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state. ✔
- *positive recurrent* if the chain returns to any state in finite time with probability 1.

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state. ✓
- *positive recurrent* if the chain returns to any state in finite time with probability 1. ✓

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state. ✓
- *positive recurrent* if the chain returns to any state in finite time with probability 1. ✓
- *periodic with period k* if returns to the same state is possible only in multiples of k steps.

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state. ✔
- *positive recurrent* if the chain returns to any state in finite time with probability 1. ✔
- *periodic with period k* if returns to the same state is possible only in multiples of k steps. ✗

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state. ✔
- *positive recurrent* if the chain returns to any state in finite time with probability 1. ✔
- *periodic with period k* if returns to the same state is possible only in multiples of k steps. ✗
- *aperiodic* if it is not periodic. ✔

Steady state distribution

A Markov chain is called:

- *irreducible* if any state is reachable from any other state. ✓
- *positive recurrent* if the chain returns to any state in finite time with probability 1. ✓
- *periodic with period k* if returns to the same state is possible only in multiples of k steps. ✗
- *aperiodic* if it is not periodic. ✓

Theorem

An irreducible Markov chain with transition matrix P has a unique steady state distribution π with $\pi = P\pi$ if and only if it is positive recurrent.

Moreover, if the Markov chain is aperiodic, then $\pi = \lim_{n\to\infty} P^n \pi_0$ for any π_0.

Monopoly as a Markov chain

Monopoly as a Markov chain

There are 119 states:

- There are 40 distinct squares on the board (being in jail and visiting jail are distinct).
- For each non-jail square, we keep track of how many previous doubles have been rolled $\rightarrow$ 3 states for each non-jail square.
- In jail, we keep track of how many turns the player has been in jail already $\rightarrow$ 3 states for jail.
- The jail state "been there two turns already" is equivalent to the state "visiting jail without previous doubles" $\rightarrow$ 1 fewer states in total.

Monopoly as a Markov chain

There are 119 states:

- There are 40 distinct squares on the board (being in jail and visiting jail are distinct).
- For each non-jail square, we keep track of how many previous doubles have been rolled $\rightarrow$ 3 states for each non-jail square.
- In jail, we keep track of how many turns the player has been in jail already $\rightarrow$ 3 states for jail.
- The jail state "been there two turns already" is equivalent to the state "visiting jail without previous doubles" $\rightarrow$ 1 fewer states in total.

A more extensive Markov chain would involve the amount of players' cash.

Decomposition of transition matrix

Transition D due to dice and transition C due to cards is simple. Because we roll dice first, then pick cards, it follows that $P = C \cdot D$.

$$D = \begin{pmatrix} (0,0,0) & (0,0,0) & \cdots \\ (0,0,0) & (0,0,0) & \cdots \\ (0,\frac{1}{36},0) & (0,0,0) & \cdots \\ (\frac{2}{36},0,0) & (0,\frac{1}{36},0) & \cdots \\ (\frac{2}{36},\frac{1}{36},0) & (\frac{2}{36},0,0) & \cdots \\ (\frac{4}{36},0,0) & (\frac{2}{36},\frac{1}{36},0) & \cdots \\ (\frac{4}{36},\frac{1}{36},0) & (\frac{4}{36},0,0) & \cdots \\ (\frac{6}{36},0,0) & (\frac{4}{36},\frac{1}{36},0) & \cdots \\ (\frac{4}{36},\frac{1}{36},0) & (\frac{6}{36},0,0) & \cdots \\ (\frac{4}{36},0,0) & (\frac{2}{36},\frac{1}{36},0) & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}$$

Decomposition of transition matrix

Transition D due to dice and transition C due to cards is simple. Because we roll dice first, then pick cards, it follows that $P = C \cdot D$.

$$D = \begin{pmatrix} (0,0,0) & (0,0,0) & \cdots \\ (0,0,0) & (0,0,0) & \cdots \\ (0,\frac{1}{36},0) & (0,0,0) & \cdots \\ (\frac{2}{36},0,0) & (0,\frac{1}{36},0) & \cdots \\ (\frac{2}{36},\frac{1}{36},0) & (\frac{2}{36},0,0) & \cdots \\ (\frac{4}{36},0,0) & (\frac{2}{36},\frac{1}{36},0) & \cdots \\ (\frac{4}{36},\frac{1}{36},0) & (\frac{4}{36},0,0) & \cdots \\ (\frac{6}{36},0,0) & (\frac{4}{36},\frac{1}{36},0) & \cdots \\ (\frac{4}{36},\frac{1}{36},0) & (\frac{6}{36},0,0) & \cdots \\ (\frac{4}{36},0,0) & (\frac{2}{36},\frac{1}{36},0) & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \qquad \pi_0 = \begin{pmatrix} (1,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ \vdots \end{pmatrix}$$

- Initial distribution π_0: probability 1 on "Go" without previous doubles.

Decomposition of transition matrix

Transition D due to dice and transition C due to cards is simple. Because we roll dice first, then pick cards, it follows that $P = C \cdot D$.

$$D = \begin{pmatrix} (0,0,0) & (0,0,0) & \cdots \\ (0,0,0) & (0,0,0) & \cdots \\ (0,\frac{1}{36},0) & (0,0,0) & \cdots \\ (\frac{2}{36},0,0) & (0,\frac{1}{36},0) & \cdots \\ (\frac{2}{36},\frac{1}{36},0) & (\frac{2}{36},0,0) & \cdots \\ (\frac{4}{36},0,0) & (\frac{2}{36},\frac{1}{36},0) & \cdots \\ (\frac{4}{36},\frac{1}{36},0) & (\frac{4}{36},0,0) & \cdots \\ (\frac{6}{36},0,0) & (\frac{4}{36},\frac{1}{36},0) & \cdots \\ (\frac{4}{36},\frac{1}{36},0) & (\frac{6}{36},0,0) & \cdots \\ (\frac{4}{36},0,0) & (\frac{2}{36},\frac{1}{36},0) & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \qquad \pi_0 = \begin{pmatrix} (1,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ (0,0,0) \\ \vdots \end{pmatrix}$$

- Initial distribution π_0: probability 1 on "Go" without previous doubles.
- Distribution after n rolls: $\pi_n = P^n \cdot \pi_0$.

Distribution after several moves

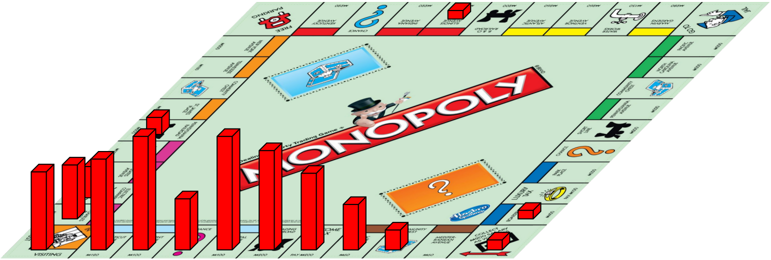

Move count: 1

Distribution after several moves

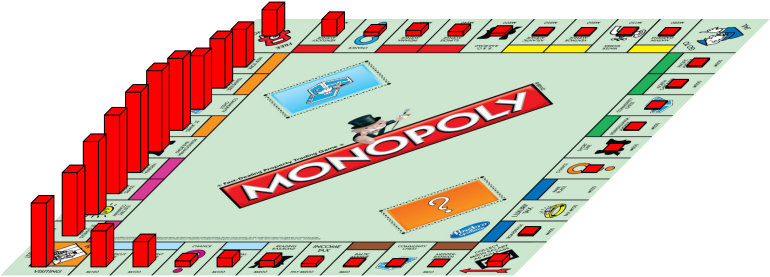

Move count: 2

Distribution after several moves

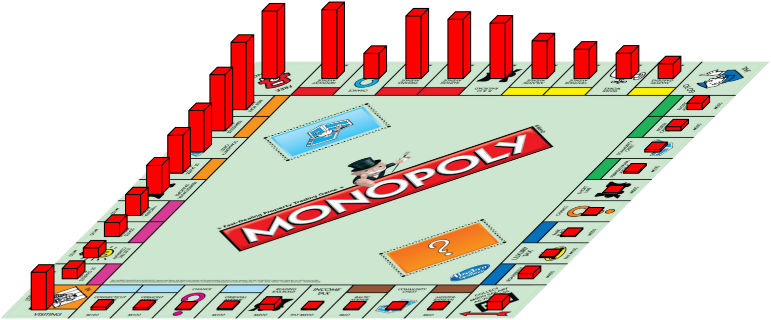

Move count: 3

Distribution after several moves

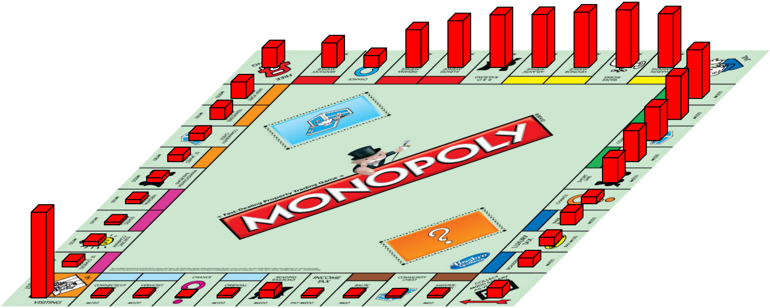

Move count: 4

Distribution after several moves

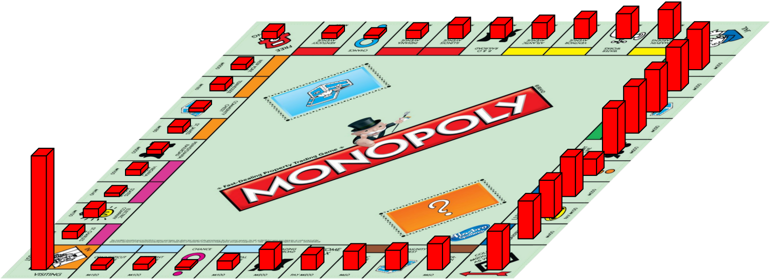

Move count: 5

Distribution after several moves

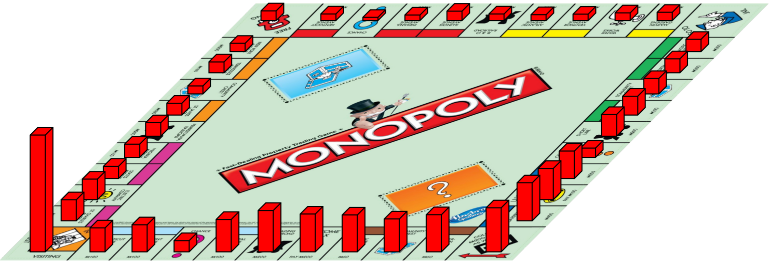

Move count: 6

Distribution after several moves

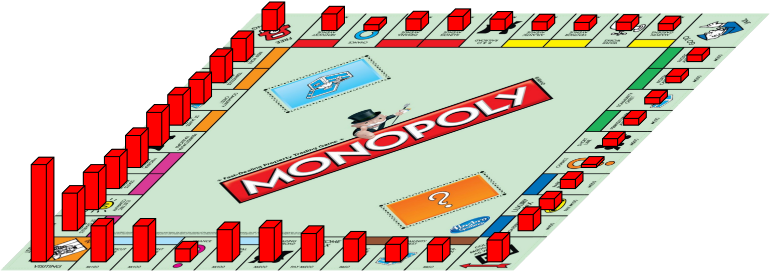

Move count: 7

Distribution after several moves

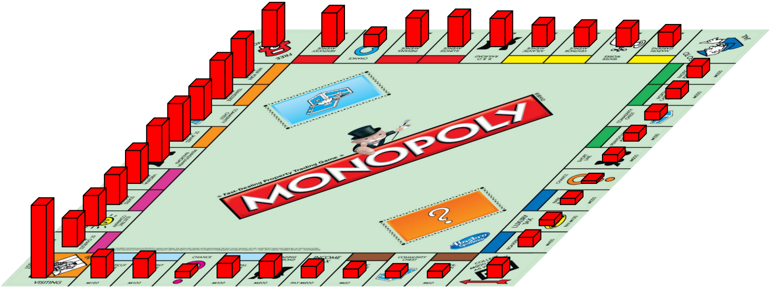

Move count: 8

Distribution after several moves

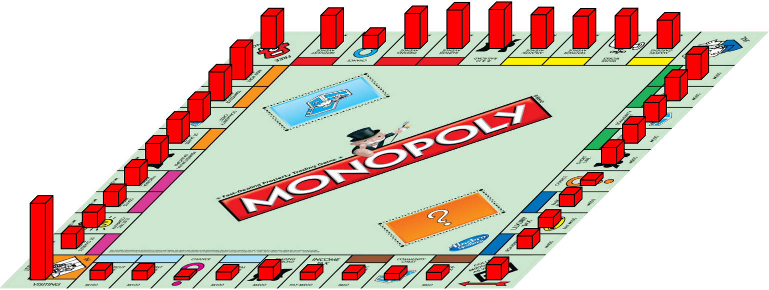

Move count: 9

Distribution after several moves

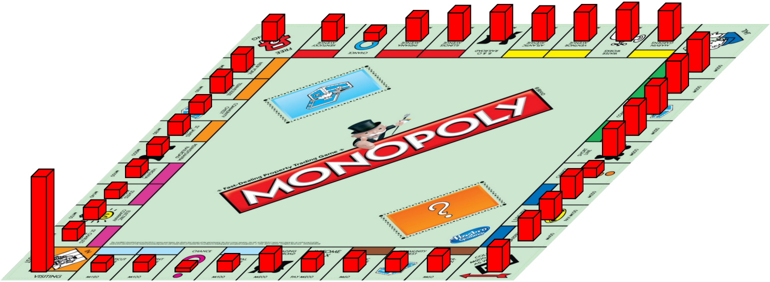

Move count: 10

Steady state distribution

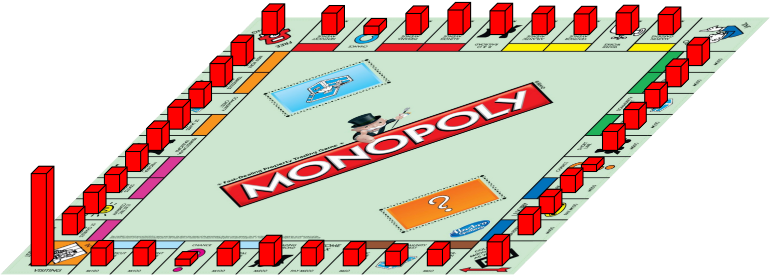

Move count: ∞

Implications for strategy

Expected income per turn on empty property[1]

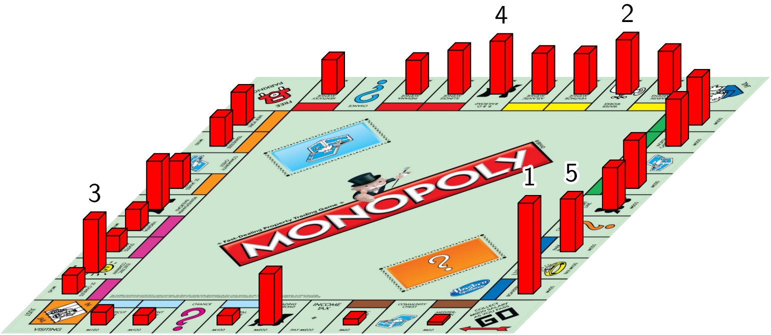

$$\text{Expected income} = \text{rent} \times \text{probability}$$

[1] Railroads and utilities: empty means only a single property of that kind is owned.

Expected income per turn for fully developed property[2]

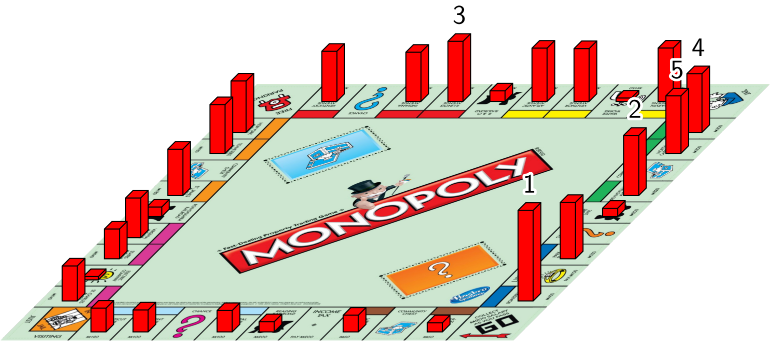

Expected income = rent × probability

[2] Railroads and utilities: fully developed means all properties of that kind are owned.

Expected return per turn for fully developed property[3]

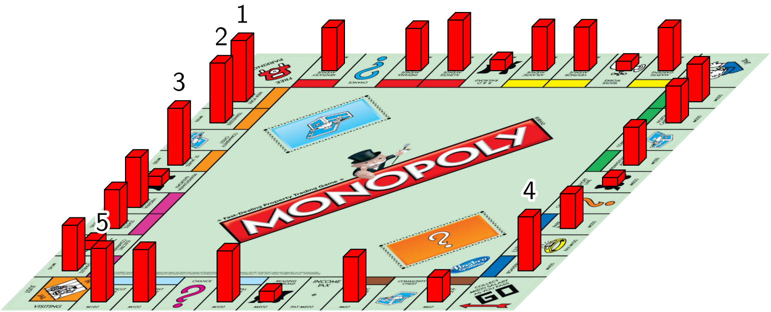

$$\text{Expected return} = \frac{\text{rent} \times \text{probability}}{\text{cost of purchase and development}}$$

[3] Railroads and utilities: fully developed = all properties of that kind are owned.

Expected income/return per turn and color

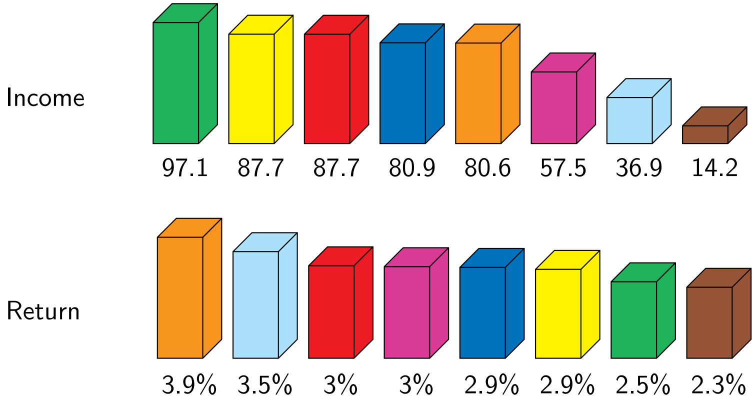

Expected income/return per turn and color

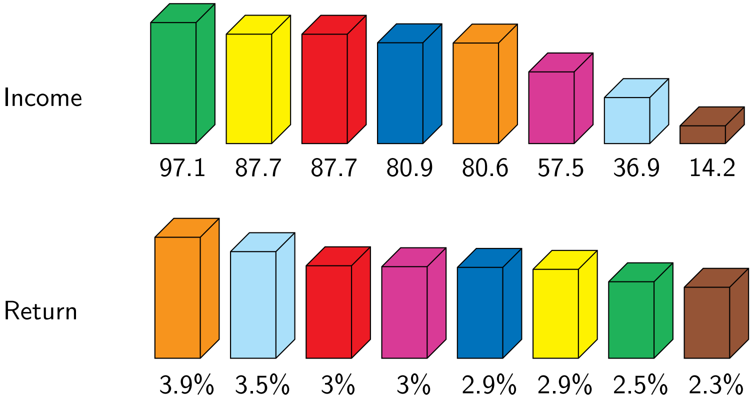

After fully developing the orange and green properties, it takes an expected 6.41 and 10 rounds, respectively, to earn back the costs in a 5-player game.

Expected return per color and stage of development

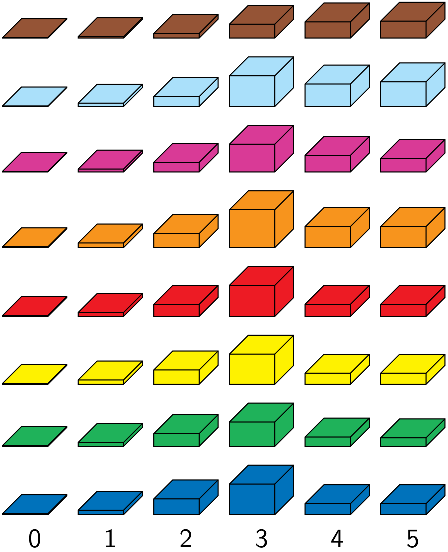

Quality of trades

Extend Markov chain by adding players' capital to the state space:

- If player A lands on property of player B, rent is transferred.
- 0 capital is an absorbing state $\rightarrow$ player is eliminated.
- Summing up probabilities of all states where player A's capital is 0 gives player A's ruin probability.
- A trade is "good" if it reduces one's ruin probability.

Quality of trades

Extend Markov chain by adding players' capital to the state space:

- If player A lands on property of player B, rent is transferred.
- 0 capital is an absorbing state $\rightarrow$ player is eliminated.
- Summing up probabilities of all states where player A's capital is 0 gives player A's ruin probability.
- A trade is "good" if it reduces one's ruin probability.

Difficulty: Depends on players' future actions $\rightarrow$ Markov decision process.

Quality of trades

Extend Markov chain by adding players' capital to the state space:

- If player A lands on property of player B, rent is transferred.
- 0 capital is an absorbing state $\rightarrow$ player is eliminated.
- Summing up probabilities of all states where player A's capital is 0 gives player A's ruin probability.
- A trade is "good" if it reduces one's ruin probability.

Difficulty: Depends on players' future actions $\rightarrow$ Markov decision process.

Assume optimal play in the future:

- Trades are executed only if it is beneficial for all parties involved in the trade $\rightarrow$ no trades if there are only two players left.
- Players develop the properties that minimize their ruin probability.

Conclusions

Monopoly:

- Aim to obtain all orange or light blue properties in the beginning of the game. Next best choices are red and yellow.
- Develop these properties until the third house before investing in another property.
- Aim to have 2 colors of orange, red, yellow, green, or dark blue for end game.

Conclusions

Monopoly:

- Aim to obtain all orange or light blue properties in the beginning of the game. Next best choices are red and yellow.
- Develop these properties until the third house before investing in another property.
- Aim to have 2 colors of orange, red, yellow, green, or dark blue for end game.

Markov chains:

- Math is awesome.
- This is not a simulation: calculated probabilities are exact.
- You don't have to be a mathematician to use Markov chains.
- Make sure you check the conditions of any theorems you use.