Markov and Mr. Monopoly Make Millions
SP.268 Spring 2011

Probability

Probability is key to many fields, such as econometrics, quantum mechanics, signal processing, and theoretical computer science. We will go through a gentle introduction to the basics of probability, then discuss how probability can be used to analyze Monopoly. We will focus on discrete probability here, though we could easily convert to the continuous analogs.

Sets

A **set** is a collection of items. An example of a set can be all the Course XIV classes offered at MIT: $\{14.01, 14.02, 14.04, 14.05, 14.32, 14.33, 14.36 \ldots\}$. For the following definitions and examples, let A and S be arbitrary sets.

An **element** of a set is something belonging to that set. We write $a \in A$ if a is a member of the set A, and $a \notin A$ if a is not a member of the set A.

A **subset** is a set contained within another set, in other words, if all members of a set belongs to another set. A is a subset of S if all members of A belong to S, and we write $A \subseteq S$. Note that:

- $A = S$ if and only if $A \subseteq S$ and $S \subseteq A$.

- The **empty set** $\emptyset$, or a set with no elements, is a proper subset of every set.
- A **proper subset** is a set that is strictly contained in another set. That is, A is a proper subset of S if and only if there is at least one element contained in S that is not contained in A, and all elements of A are contained in S.

The **cardinality** of a set, denoted $|A|$ here, is the number of elements in that set. If $A \subseteq S$, then $|A| \leq |S|$. If $A \subset S$, then $|A| < |S|$.

Probability and Sets

Now that we have defined sets generally, let's look at how sets are used when applied to probability. The 'things' or 'items' that we're concerned with are **outcomes**–outcomes from flipping coins, dealing hands or cards, etc. The **sample space** is the set of all possible outcomes, denoted $\mathbf{\Omega}$. A subset of a sample space consists of the outcomes that we're interested in, called **events**.

Suppose we have the events A, B, and C. The **interesction** of two events is the event that they both occur. $C = A \cap B$ if the event C represents both A and B occuring. If $A \cap S = \phi$, then the two sets are called **disjoint** or **mutually exclusive**. The **union** of two events is the event that either one or the other occurs, denoted $C = A \cup B$.

Laws of set operations:

- Commutative:

$$A \cup B = B \cup A$$
$$A \cap B = B \cap A$$

- Associative:

$$(A \cup B) \cup C = A \cup (B \cup C)$$
$$(A \cap B) \cap C = A \cap (B \cap C)$$

- Distributive:

$$(A \cup B) \cap C = (A \cap C) \cup (B \cap C)$$
$$(A \cap B) \cup C = (A \cup C) \cap (B \cup C)$$

The probability of an event is a mapping from the set of events to the interval [0,1]. When we talk about the probability of some event A in Ω, it will always follow these axioms:

1. The probability of the sample space, Ω, is $P(\Omega) = 1$;
2. $P(A) \geq 0$ for all $A \in \Omega$.
3. If A_1 and A_2 are disjoint, then

$$P(A_1 \cup A_2) = P(A_1) + P(A_2).$$

More generally, if A_i for $i = 1, 2, 3, \ldots$ are disjoint, then

$$P\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} P(A_i).$$

The **inclusion-exclusion** principle is very useful in calculating probabilities. It states that for two events, A_1 and A_2, not necessarily disjoint as in the third axiom above,

$$|A_1 \cup A_2| = |A_1| + |A_2| - |A_1 \cap A_2|.$$

The third term in in the equation above subtracts the overlap in A_1 and A_3, which was counted twice. The probability version of the inclusion-exclusion principle is

$$P(A_1 \cup A_2) = P(A_1) + P(A_2) - P(A_1 \cap A_2).$$

The **complement** of a set A, commonly denoted A', A^c, or $\overline{A}$, is all elements in the sample space that don't belong to that set.

$$\begin{aligned} A^c &= \Omega - A, \\ |A^c| &= |\Omega| - |A|, \\ P(A^c) &= P(\Omega) - P(A) \\ &= 1 - P(A). \end{aligned}$$

For most problems, the goal will be to find the likelihood that an event E happens, or $P(E)$, out of the set of possible outcomes S. When all the outcomes are equally likely,

$$P(E) = \frac{|E|}{|S|}.$$

We're adding up all the elements in E and all the elements in S, then dividing them. This leads us to the topic of counting, which is used when dealing with discrete, finite sample spaces.

Counting

We make the assumption that all the outcomes are equally likely, also known as the assumption of uniform probability. All that needs to be done then is add up the number of outcomes that we care about and divide that by the number of all possible outcomes. The trickiest part is defining the event and sample space and making sure that we count everything the right number of times.

Counting Rules

We've seen the **Sum Rule** already, just not labeled with the name. If $A_1, A_2, \ldots, A_n$ are disjoint sets, then

$$|A_1 \cup A_2 \cup \ldots \cup A_n| = |A_1| + |A_2| + \ldots + |A_n|.$$

What's the probability version of the Sum Rule?

Example 1

Excuse this somewhat lame example, but its purpose is to show the sum rule at work. In a group of 150 students, 15 use Internet Explorer as their web browser of choice, 80 use Firefox, 15 use Safari, and 40 use Chrome. If being "cool" means you use Firefox or Chrome as your main web browser, what is the probability that we pick one student who is "cool?"

Let the set C be the set of "cool" students; there are $80 + 40$ students in C, by the sum rule. Let S be the set of all students; there are 150 students in total, as stated in the problem. Therefore, the probability of picking a cool student is:

$$P(\text{picking a cool student}) = \frac{|C|}{|S|} = \frac{120}{150} = \frac{4}{5}.$$

The **multiplication rule** states that for a length-k sequence, where the first term is chosen out of set S_1, the 2nd term is chosen out of set S_2 ... the last term is chosen out of S_k, then

$$\begin{aligned} |\text{Total \# of sequences}| &= |S_1 \times S_2 \times \cdots \times S_k| \\ &= |S_1| \cdot |S_2| \cdots |S_k|. \end{aligned}$$

Example 2

The Athena combination lock just got changed again. You're far from any Quickstation and there's no one else nearby. Suppose you wanted to try your luck at guessing the combo (and you don't have SIPB's hint board). How many possible combinations could you try?

The athena door locks have 5 buttons, 3 on the top row and 1 on the bottom row (the bottom-right button is a reset, so it doesn't count). The athena passcode is 5 digits. Let D_i, for $i = 1, 2, 3, 4, 5$, represent the set of buttons possible for each digit.

$$
\begin{aligned}
|\text{Total \# of combinations}| &= |D_1 \times D_2 \times D_3 \times D_4 \times D_5| \\
&= |D_1| \cdot |D_2| \cdot |D_3| \cdot |D_4| \cdot |D_5| \\
&= 5 \cdot 5 \cdot 5 \cdot 5 \cdot 5 \\
&= 5^5
\end{aligned}
$$

Permutations

The set of **permutations** on a collection of objects is an example of the algebraic structure called a 'group,' which was covered in the Rubik's Cube lecture. Here we'll use 'a set of ordered objects' as a working definition of permutations. For a collection of n objects, there are $n(n-1)(n-2)\cdots(1) = n!$ different orderings of the objects.

Example 3: The Birthday Problem

You're in a room with a bunch of people, say $n \leq 365$ people.

a) What is the probability that two people in the room have the same birthday? Ignore complications with leap years and assume there are 365 days in a year. We also assume that birthdays are random (not exactly true).

   This problem is best approached the other way around, with the probability that no two people have the same birthday.

   Let A be the event that two people have the same birthday. Then A^c is the event that no two people have the same birthday. Note that $P(A) = 1-P(A^c)$. We start with person 1; this person can have any 1 of 365 days out of the year. A second person can only have a birthday on the 364 days out of the year that hasn't been 'taken.' By assumption of random birthdays, and of uniform probability, the chance that this person has any of the 364 birthdays is $\frac{364}{365}$. A third person can only have a birthday out of the 353 days not 'taken,' and the corresponding probability of such an event is $\frac{363}{365}$. This continues until we've covered all n people.

$$\begin{aligned} P(A^c) &= \frac{365 \cdot 364 \cdot 363 \cdots (365 - n + 1)}{365^n} \\ P(A) &= 1 - P(A^c) \\ &= 1 - \frac{365!}{365^n \cdot n!} \end{aligned}$$

b) What is the probability that someone shares your birthday?

Each person can have your birthday with probability $\frac{1}{365}$. There are $n-1$ people besides you, so the probability that someone shares your birthday is $\frac{n-1}{365}$.

The answers to part a) and part b) are quite different, but the way the questions were phrased were only slightly different. Half the work in probability questions is usually figuring out what the question wants from you...

What happens if $n > 365$? You can answer part a) without doing any math, by the **Pigeonhole Principle**. The Pigeonhole Principle states that in a mapping from set X to set Y, if $|X| > |Y|$, then more than one element of X map to some element in Y.

Combinations

We will also want to deal with collections that are unordered. How many ways are there to take r objects out of a set of n objects?

For the first object, we have n to choose from. For the 2nd object, we have $n-1$ to choose from. For the rth object, we have $n-r+1$ to choose from. But note that once we've selected r objects this way, they are in some kind of order, and the answer $n(n-1)(n-2)\cdots(n-r+1) = \frac{n!}{(n-r)!}$ is not correct. We must divide by $r!$, which is the number of ways you can order (permute) r objects.

The number of ways that we can take r objects out of a set of n objects is therefore

$$\frac{n!}{r!(n-r)!} = \binom{n}{r}.$$

Example 4: Hands of cards

In this example we're using a standard 52-card deck.

a) How many ways are there to deal a 5-card hand?

$$\binom{52}{5} = \frac{52!}{5!(47)!} = 2598960$$

b) How many ways are there to deal a flush, a 5-card hand with all cards the same suit?

There are $\binom{4}{1}$ ways to choose the suit, and $\binom{13}{5}$ ways to choose the 5 cards out of that suit.

$$\#\text{ of ways to deal a flush} = 4\binom{13}{5} = 5148$$

c) How many ways are there to deal a 5-card hand with 1 pair?

There are $\binom{13}{1}$ ways to choose the card value of the pair, and $\binom{4}{2}$ ways to choose the suits of the pair; then there are $\binom{50}{3}$ ways to choose the remaining 3 cards of the 5-card hand.

$$\#\text{ of ways to deal a hand with 1 pair} = 13\binom{4}{2}\binom{50}{3} = 1528800$$

d) How many ways are there to deal a 5-card hand with *only* 1 pair?

As before, there are $\binom{13}{1}$ ways to choose the card value of the pair, and $\binom{4}{2}$ ways to choose the suits of the pair. But the problem specifies only 1 pair. The remaining 3 cards in the hand cannot contain a pair. So there are are $\binom{12}{3}$ ways to choose 3 different values besides the value that's already a pair, and they can be from any suit.

$$\text{\# of ways to deal a hand with only 1 pair} = 13\binom{4}{2}\binom{12}{3}4^3 = 1098240$$

e) How many ways are there to deal a 3-of-a-kind?

There are $\binom{13}{1}$ ways to choose the card value of the 3-of-a-kind, and $\binom{4}{3}$ ways to choose the suits of the pair; then there are $\binom{48}{2}$ ways to choose the remaining 2 cards of the 5-card hand, making sure that the value of the 3-of-a-kind doesn't get chosen (otherwise we'd get a 4-of-a-kind).

$$\text{\# of ways to deal a 3-of-a-kind} = 13\binom{4}{3}\binom{48}{2} = 58656$$

f) How many ways are there to deal a full house, a 5-card hand with 3 of one kind and 2 of another?

As before, there are $\binom{13}{1}$ ways to choose the card value of the 3-of-a-kind, and $\binom{4}{3}$ ways to choose the suits of the pair. Then there are $\binom{12}{1}$ ways to choose the value of the 2-of-a-kind and $\binom{4}{2}$ ways to choose the suit.

$$\text{\# of ways to deal a full house} = 13\binom{4}{3}\binom{12}{1}\binom{4}{2} = 3744$$

Conditional Probability

If we're interested in the probability that some event A occurs *given* that some event B has already occurred, the sample space becomes B. The probability of A conditioned on B becomes a probability on the space B.

The **Multiplication Law** states that $P(A \cap B) = P(A|B)P(B)$.

So we get that

$$P(A|B) = \frac{P(A \cap B)}{P(B)}, \text{ for } P(B) \neq 0.$$

With some rearranging, we get **Baye's Rule**, which is commonly seen in many different forms:

$$P(B|A)P(A) = P(A|B)P(B).$$

The **Law of Total Probability** gives us the ability to isolate the probability of one event on a partitioned probability space. Given a space Ω that is partitioned by $B_n : n = 1, 2, \ldots$, and an event A,

$$P(A) = \sum_{i=1}^{n} P(A \cap B_i)$$

Example 5

Melissa and I are going to assign your P/F grades for this seminar by picking them out of a hat. We take 100 slips of paper and mark P on half of them, F on the other half. Then we put the slips of paper in two hats, and pick a slip of paper from one of the hats. Whatever we pick will be your grade. *But*, being as merciful and fair as we so obviously are, and curious how much you got out of this class, we'll leave it up to you to place the slips of paper into the two hats any way you want. How will you do it?

Monopoly

The game of MONOPOLY® came about during the Great Depression, originating from Charles Darrow of Germantown, Pennsylvania. It started out as handmade sets sold in a shop in Philadelphia, and as people grew to love the game, Darrow approached Parker Brothers to enlarge the production scale (he'd actually been rebuffed by the Parker Brothers the first time in 1933 due to 52 'fundamental playing flaws'). Today, MONOPOLY® is the best-selling board game in the world, distributed in 111 countries and 43 languages. Some fun facts from the MONOPOLY® website:

- The longest MONOPOLY® game ever played was 1,680 hours long.
- The MONOPOLY® man isn't a Parker Brother. His name is Mr. Monopoly.
- Parker Brothers once sent an armored car with a million MONOPOLY® dollars to Pittsburgh because a marathon game there had run out of money.
- MONOPOLY® comes in a Braille version.
- The four most-landed-on squares are Jail, Illinois Avenue, "Go", and the B&O Railroad.

The last in the list of fun facts above is more than meets the eye. What makes certain game squares more likely to be landed-on than others? Illinois Avenue doesn't seem to be special compared to other properties... It turns out that we can model the MONOPOLY® game board to calculate the probability of landing on a certain square.

Rules

The objective of the game is to bankrupt all opponents, though most games played with family and friends end when it is apparent that someone will win. A typical game of MONOPOLY® uses the following items:

- 1 game board
- 2 dice
- token for players (11 official MONOPOLY® ones)

- 32 houses
- 12 hotels
- 16 Chance cards
- 16 Community Chest cards
- property deeds for each of the 22 MONOPOLY® properties
- $15140 in MONOPOLY® money

The Chance and Community Chest cards are placed face down on the game board, and a player must pick one of the cards when he lands on the Chance or the Community Chest game squares. Each player is given $1500 to begin the game. All remaining money, game piece, houses, hotels, and deeds of unsold property go to the Bank. The Bank collects all taxes, fines, loans, and interest. The Bank never goes 'broke.' If the Bank runs out of MONOPOLY® money, then more can be issued (see fun fact above).

Players begin on the Go square, roll two dice, and advance as many steps as dots displayed on the the two dice. A player can buy any property, utility, or railroad that isn't already owned by another player, or must to draw Chance/Community Chest cards, pay rent, fines, or go to Jail as dictated by the square he lands on. If a player throws a double, then he moves his token the number of steps, is subject to whatever privileges or penalties of the square he lands on, and then tosses the dice again. If a player tosses three doubles in a single turn, he must go to Jail.

Landing on the Jail square is just 'visiting Jail', while landing on the 'Go to Jail' square, drawing a 'Go to Jail' card, and tossing doubles 3 times during a turn are actual Jail sentences. Any Jail term lasts 3 turns. A player tosses dice at each turn, and if he tosses a double, then he is free to get out of jail and advances the number of steps as his double shows. That player does not take another turn. A player gets out of Jail if he has a 'Get out of Jail Free' card, or if another player is willing to sell him a 'Get out of Jail Free' card at a negotiated price, or if the player pays a $50 fine.

If a player lands on a property owned by another player, then the owner collects rent based on the information on the property deed. Rents are much higher for properties with houses or hotels. When a player owns all the MONOPOLY® properties in a color group, then he has the option to build houses on those properties. If he buys one house, he can put them on any one of those properties. The next house he buys must be erected on one of the unimproved properties of that or any other complete color group, and so on. Thus, players must build evenly across all his properties in a color group.

More details of the rules of the game will unfurl as we analyze the game. First, a bit of linear algebra.

Matrices, Eigenvalues, and Eigenvectors

The linear equation $A\mathbf{x} = \mathbf{b}$ is at the heart of most introductory linear algebra courses. A is a **matrix**, and $\mathbf{x}$ and $\mathbf{b}$ are **vectors**; the matrix A 'operates' on $\mathbf{x}$ to give $\mathbf{b}$; $\mathbf{x}$ and $\mathbf{b}$ lie on the same vector space but are in different directions unless A is the identity matrix I.

Eigenvectors are special vectors associated with every operating matrix. These vectors don't change directions when multiplied by the matrix, and we get the equation $A\mathbf{x} = \lambda\mathbf{x}$. Each eigenvector has its own *eigenvalue* λ. Most 2×2 matrices have two eigenvectors and their two corresponding eigenvalues.

What happens when A operates on $\mathbf{x}$ more than once? As in, what's $A^2\mathbf{x}$? $A^3\mathbf{x}$? $A^{100}\mathbf{x}$?

The number λ is an eigenvalue of A if and only if $A - \lambda I$ (which is a matrix) is singular, or $\det(A - \lambda I) = 0$. A **singular** matrix is a square matrix that has no inverse, or $\det A = 0$. Then, for each eigenvalue, we solve $(A - \lambda I)\mathbf{x} = 0$, or $A\mathbf{x} = \lambda\mathbf{x}$ to find the eigenvector $\mathbf{x}$.

Example 6

Find the eigenvalues and eigenvectors of $A = \left[\begin{smallmatrix} 1 & 2 \\ 2 & 4 \end{smallmatrix}\right]$.

$$A - \lambda I = \left[\begin{smallmatrix} 1-\lambda & 2 \\ 2 & 4-\lambda \end{smallmatrix}\right]$$

Take the determinant of this matrix.

$$det\left[\begin{smallmatrix} 1-\lambda & 2 \\ 2 & 4-\lambda \end{smallmatrix}\right] = (1-\lambda)(4-\lambda) - (2)(2) = \lambda^2 - 5\lambda.$$

Set the determinant to 0, and solve for λ.

$$\lambda^2 - 5\lambda = 0 \text{ gives } \lambda_1 = 0 \text{ and } \lambda_2 = 5.$$

Solve $(A - \lambda I)\mathbf{x} = \mathbf{0}$ separately for $\lambda_1 = 0$ and $\lambda_2 = 5$.

$$(A - 0I) = \left[\begin{smallmatrix} 1 & 2 \\ 2 & 4 \end{smallmatrix}\right]\left[\begin{smallmatrix} x_1 \\ x_2 \end{smallmatrix}\right] = \left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right] \text{ gives eigenvector } \left[\begin{smallmatrix} 2 \\ -1 \end{smallmatrix}\right]$$
$$(A - 5I) = \left[\begin{smallmatrix} 1 & 2 \\ 2 & 4 \end{smallmatrix}\right]\left[\begin{smallmatrix} x_1 \\ x_2 \end{smallmatrix}\right] = \left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right] \text{ gives eigenvector } \left[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}\right]$$

As a side note, because the vectors that make up A are constant multiples of each other, we know that A itself is a singular matrix. The determinant of a matrix can be found by taking the product of all its eigenvalues, so if the determinant is zero, then we know one of the eigenvalues must be zero.

Markov Chains

Markov chains are the probabilistic versions of deterministic finite automata. For our analysis of MONOPOLY®, we'll consider each of the 40 game squares to be a **state** H. At each time step n, the probabilistic state distribution X_n will be a 40×1 vector, with each element representing the probability that a player ends his turn in that state. Our state, H, belongs to the set S of the **state space** of size 40. The Markov chain is described in terms of its **transition probabilities** p_{ij}, which is the probability that we'll go from state i to state j at a time step. The transition probabilities sum to 1.

$$p_{ij} = P(H_{n+1} = j | H_n = i), \quad i, j \in S$$

$$\sum_j p_{ij} = 1$$

The probability that we're in a certain state at time step n depends only on our state at time step $n-1$, and is independent of all states besides the previous state:

$$P(H_{n+1} = j | H_n = i, H_{n-1} = i-1, \ldots, H_o = i_o) = P(H_{n+1} = j | H_n = i) = p_{ij}$$

The **transition matrix** captures all the transition probabilities and operates on our state distribution vector. Such a matrix is called a Markov matrix, and it is also a square matrix.

$$\begin{bmatrix} p_{00} & p_{01} & \cdots & p_{0m} \\ p_{10} & p_{11} & \cdots & p_{1m} \\ \vdots & \vdots & \ddots & \vdots \\ p_{m0} & p_{m1} & \cdots & p_{mm} \end{bmatrix}^T$$

Special matrices will have special eigenvalues and eigenvectors, and for Markov matrices, all entries are positive and every column sums to 1. Can you see why they must add to 1? The largest eigenvalue is 1, and the corresponding eigenvector is the state that comes out at the end. The eigenvectors of other eigenvalues fall to 0 over time.

Example 7

The ESG elevator has two states: SLOW and BROKEN. If it is SLOW today, then the probability that it becomes BROKEN tomorrow is 0.6, and the probability that it stays SLOW the next day is 0.4. If it is BROKEN, then the probability that MIT Facilities comes to try to fix it (making it SLOW) is 0.2, but most likely, with probability 0.8, it'll just stay BROKEN. If we start the school year in the fall with a SLOW elevator, what kind of elevator will we have at the end of the school year?

Our initial state distribution is $\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$, with x_1 representing the probability of having a BROKEN elevator, and x_2 representing the probability of having a SLOW elevator.

Our transition matrix is $A = \begin{bmatrix} .6 & .8 \\ .4 & .2 \end{bmatrix}$. It has eigenvalues 1 and -.2, but over time, the eigenvector associated with -.2 will be multiplied by $(-.2)^n \to 0$. We consider the eigenvector with eigenvalue 1.

$$(A - I)\mathbf{x} = 0 \Longrightarrow \text{ (normalized) } \mathbf{x} = \begin{bmatrix} \frac{2}{3} \\ \frac{1}{3} \end{bmatrix}$$

We are twice as likely to end up with a BROKEN elevator.

MONOPOLY®

Ian Stewart, a math professor at the University of Warwick, wrote a column in the April, 1996 issue of Scientific American seeking to answer the question: 'Is Monopoly fair?' In other words, is every MONOPOLY® square equally likely to be occupied? His initial analysis was only a mathematical exercise, and his model abstracted many of the realistic playing rules.

Initial Analysis

We abstract away the rules about rolling doubles, Chance/Community Chest squares, and the complications involving going to Jail. Then on each roll of our dice the number of steps we could possibly take (sum of rolling two dice) is distributed as follows:

Number on the two dice	Probability
7	$\frac{6}{36}$
6, 8	$\frac{5}{36}$
5, 9	$\frac{4}{36}$
4, 10	$\frac{3}{36}$
3, 11	$\frac{21}{36}$
2, 12	$\frac{1}{36}$

The initial position probability vector, $\mathbf{H_0}$, is a 40-dimensional-vector with 1 as the 0th element and 0 everywhere else. If we index the board from 0 through 39, beginning at 'Go', then after the first turn, the vector of position probabilties $\mathbf{H_1}$ would be:

$$[0, 0, \frac{1}{36}, \frac{2}{36}, \frac{3}{36}, \frac{4}{36}, \frac{5}{36}, \frac{6}{36}, \frac{5}{36}, \frac{4}{36}, \frac{3}{36}, \frac{2}{36}, \frac{1}{36}, \underbrace{0, 0, 0, \ldots, 0}_{\text{27 zeros}}]'.$$

The 1st column of our 40×40 Markov transition matrix M would be $\mathbf{H_1}$. The second column of the matrix would be this vector with 3 zeros before the beginning of the fractions and 27 zeros after, in essence, the same vector 'shifted' over by 1. The third row would be shifted over again, so on until we've completed all 40 rows of the matrix.

To get our state probability vector after moving n times, we would just calculate $A^n\mathbf{H_0}$. For sufficiently large n, our state probability vector would be the eigenvector of A that corresponds to eigenvalue 1. The eigenvector for the simple model of Monopoly has all entries equal to 1, because the transition from any square on the board is the same.

The Real Game

We made a few simplifying assumptions in the last section: tossing doubles, the "Go to Jail" square, and Chance/Community Chest cards.

To deal with the rule about tossing doubles, we can modify the initial vector and the transition matrix. The maximum number of spaces a player can move is 35, (if a player rolls a $\{(6{,}6),(6{,}6),(6{,}5)\}$). If a player lands on 'Go to Jail', then he would go immediately to Jail and end his turn there. If a player rolls three doubles in row, then he would also have to go to jail.

The transition at each turn would be a vector of probabilities allowing up to 2 tosses with doubles. The transition matrix will have this vector for each column (with appropriate offsets). Then we adjust the probabilities for the Jail square, which is the sum of what it had from routine tossing, the probability of the "Go to Jail" square, and probability of tossing 3 doubles in a row ($6^3/36^3 = 1/6^3$).

The Chance/Community Chest cards are actually not as complicated as they seem. There are 16 Chance cards, 10 of which tell the player to move to another square. The probability of staying in Chance is thus $\frac{1}{8}$ the probability it had before, and each of the 10 destinations is increased by $\frac{1}{10} \times P$(probability of landing on Chance). The same goes for Community Chest, which only has 2 cards that send players to other squares. The probabilities would be adjusted accordingly.

As a consequence of the "Go to Jail" square, tossing doubles, and Chance and Community Chest cards sending players to different squares, the probability distribution is no longer uniformly distributed over all 40 squares. Instead, it is skewed toward certain squares. Players are almost twice as likely to be in Jail than in any other square; the next-most-frequented square is Illinois Avenue, and GO is the third most likely square. B&O Railroad is the most-often occupied railroad.

Where and When to Build?

Rent-collecting is when things actually start to get interesting. After all, the whole point of the game is to bankrupt the other players. What strategy should we take in building houses and hotels? What can we use from our probabilistic analysis? If we take the actual decimal values of the probabilities and analyze the time of the break-event point (total cost of buildings divided by expected earnings from property per turn; how would you calculate the expected earnings?) which is when rents collected becomes greater than the cost of building the houses and hotels, we find that with 2 houses or fewer, it typically takes 20 moves or more to break even. With 3 houses, the chances are significantly better. It is even better than building 4 houses or a hotel. This is preferable strategy because one of the principle strategies of MONOPOLY® is to deplete accounts of other players fast while accumulating fast yourself (so that you can purchase more property and build more buildings). If the break-even point takes too long, then we are wasting valuable resources that could have been allocated to buildings on other properties and raising the rents of those properties.

Remarks

Other similar board games can be modeled in the same way. The premesis of Markov chains is that the next state is independent of all previous states–it ony depends on the current state.