Course:	Introduction to Stochastic Processes
Term:	Fall 2019
Instructor:	Gordan Žitković

Lecture 9
Absorption and Reward

Caveat: From now on, all Markov chains will have **finite** state spaces.

9.1 Absorption

Remember the "Tennis" example from a few lectures ago and the question we asked there, namely, how does the probability of winning a single point affect the probability of winning the overall game? An algorithm that will help you answer that question will be described in this lecture.

The first step is to understand the structure of the question asked in the light of the canonical decomposition of the previous lecture. In the "Tennis" example, all the states except for the winning ones are transient, and there are two one-element recurrent classes $\{$*Serena wins*$\}$ and $\{$*Roger wins*$\}$. The chain starts from a transient state $(0,0)$, moves around a bit, and, eventually, gets absorbed in one of the two. The probability we are interested in is not the probability that the chain will eventually get absorbed. That probability is always 1. We are, instead, interested in the probability that the absorption will occur in a particular state - the state $\{$*Serena wins*$\}$ (as opposed to $\{$*Roger wins*$\}$) in the "Tennis" example.

A more general version of the problem above is the following: let $i \in S$ be any state, and let j be a recurrent state. If the set of all recurrent states is denoted by C, and if τ_C is the first hitting time of the set C, then X_{τ_C} denotes the first recurrent state visited by the chain. Equivalently, X_{τ_C} is the value of X at (random) time τ_C; its value is the name of the state in which it happens to find itself the first time it hits the set of all recurrent states. For any two states $i, j \in S$, the **absorption probability** u_{ij} is defined as

$$u_{ij} = \mathbb{P}_i[X_{\tau_C} = j] = \mathbb{P}_i[\text{ the first recurrent state visited by } X \text{ is } j\ ].$$

There are several boring situations to discard first:

1. *j is transient:* in this case $u_{ij} = 0$ for any i because j cannot possibly be the first recurrent state we hit - it is not even recurrent.

2. *j is recurrent, and so is i.* Since i is recurrent, i.e., $i \in C$, we clearly have $\tau_C = 0$. Therefore $u_{ij} = \mathbb{P}_i[X_0 = j]$, and this equals to either 1 or 0, depending on whether $i = j$ or $i \neq j$.

That leaves us with the situation where $i \in T$ and $j \in C$ as the interesting one. In many calculations related to Markov chains, the method of *first-step decomposition* works miracles. Simply, we cut the probability space according to what happened in the first step and use the law of total probability (assuming $i \in T$, $j \in C$)

$$\begin{aligned} u_{ij} = \mathbb{P}_i[X_{\tau_C} = j] &= \sum_{k \in S} \mathbb{P}[X_{\tau_C} = j | X_0 = i, X_1 = k] \mathbb{P}[X_1 = k | X_0 = i] \\ &= \sum_{k \in S} \mathbb{P}[X_{\tau_C} = j | X_1 = k] p_{ik} \end{aligned}$$

The conditional probability $\mathbb{P}[X_{\tau_C} = j | X_1 = k]$ is an absorption probability, too. If $k = j$, then $\mathbb{P}[X_{\tau_C} = j | X_1 = k] = 1$. If $k \in C \setminus \{j\}$, then we are already in C, but in a state different from j, so $\mathbb{P}[X_{\tau_C} = j | X_1 = k] = 0$. Therefore, the sum above can be written as

$$u_{ij} = \sum_{k \in T} p_{ik} u_{kj} + p_{ij}, \tag{9.1.1}$$

which is a system of linear equations for the family $(u_{ij}, i \in T, j \in C)$. Linear systems are typically better understood when represented in the matrix form. Let U be a $T \times C$-matrix $U = (u_{ij}, i \in T, j \in C)$, and let Q be the portion of the transition matrix P corresponding to the transitions from T to T, i.e. $Q = (p_{ij}, i \in T, j \in T)$, and let R contain all transitions from T to C, i.e., $R = (p_{ij})_{i \in T, j \in C}$. If P_C denotes the matrix of all transitions from C to C, i.e., $P_C = (p_{ij}, i \in C, j \in C)$, then the canonical form of P looks like this:

$$P = \begin{bmatrix} P_C & 0 \\ R & Q \end{bmatrix}.$$

The system (9.1.1) now becomes:

$$U = QU + R, \text{ i.e., } (I - Q)U = R.$$

If the matrix $I - Q$ happens to be invertible, we are in business, because we then have an explicit expression for U:

$$U = (I - Q)^{-1} R.$$

So, is $I - Q$ invertible? It is when the state space S is finite; here is the argument, in case you are interested:

Proposition 9.1.1. *The matrix $I - Q$ is invertible and*

$$(I-Q)^{-1} = \sum_{n=0}^{\infty} Q^n. \tag{9.1.2}$$

Moreover, the entry at the position i, j in $(I-Q)^{-1}$ is the expected total number of visits to the state j, for a chain started at i.

Proof. For $k \in \mathbb{N}$, the matrix Q^k is the same as the submatrix corresponding to the transient states of the full k-step transition matrix P^k. Indeed, going from a transient state to another transient state in k steps can only happen via other transient states (once we hit a recurrent class, we are stuck there forever).

Using the same idea as in the proof of Proposition **??** from the set of notes on classification of states, we can conclude that for any two transient states i and j, we have (remember $\mathbb{E}_i[\mathbf{1}_{\{X_n=j\}}] = \mathbb{P}_i[X_n = j] = p_{ij}^{(n)}$)

$$\mathbb{E}_i[\sum_{n=0}^{\infty} \mathbf{1}_{\{X_n=j\}}] = \sum_{n\in\mathbb{N}_0} p_{ij}^{(n)} = \sum_{n\in\mathbb{N}_0} q_{ij}^{(n)} = (\sum_{n\in\mathbb{N}_0} Q^n)_{ij}.$$

On the other hand, the left hand side above is simply the expected number of visits to the state j, if we start from i. Since both i and j are transient, this number will either be 0 (if the chain never even reaches j from i), or a geometric random variable (if it does). In either case, the expected value of this quantity is finite, and, so

$$\sum_{n\in\mathbb{N}_0} q_{ij}^{(n)} < \infty.$$

Therefore, the matrix sum $F = \sum_{n\in\mathbb{N}_0} Q^n$ is well defined, and it remains to make sure that $F = (I-Q)^{-1}$, which follows from the following simple computation:

$$QF = Q \sum_{n\in\mathbb{N}_0} Q^n = \sum_{n\in\mathbb{N}_0} Q^{n+1} = \sum_{n\in\mathbb{N}} Q^n = \sum_{n\in\mathbb{N}_0} Q^n - I = F - I. \qquad \square$$

When the inverse $(I-Q)^{-1}$ exists (like in the finite case), it is called the **fundamental matrix** of the Markov chain.

Example 9.1.2. Before we turn to the "Tennis" example, let us analyze a simpler case of Gambler's ruin with $a = 3$. The states 0 and 3 are absorbing, and all the others are transient. Therefore $C_1 = \{0\}$, $C_2 = \{3\}$ and $T = T_1 = \{1, 2\}$. The transition matrix P in the canonical form (the rows and columns represent the states in the order $0, 3, 1, 2$)

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1-p & 0 & 0 & p \\ 0 & p & 1-p & 0 \end{bmatrix}$$

Therefore,

$$R = \begin{bmatrix} 1-p & 0 \\ 0 & p \end{bmatrix} \text{ and } Q = \begin{bmatrix} 0 & p \\ 1-p & 0 \end{bmatrix}.$$

The matrix $I - Q$ is a 2×2 matrix so it is easy to invert:

$$(I-Q)^{-1} = \frac{1}{1-p+p^2} \begin{bmatrix} 1 & p \\ 1-p & 1 \end{bmatrix}.$$

So

$$U = \frac{1}{1-p+p^2} \begin{bmatrix} 1 & p \\ 1-p & 1 \end{bmatrix} \begin{bmatrix} 1-p & 0 \\ 0 & p \end{bmatrix} = \begin{bmatrix} \frac{1-p}{1-p+p^2} & \frac{p^2}{1-p+p^2} \\ \frac{(1-p)^2}{1-p+p^2} & \frac{p}{1-p+p^2} \end{bmatrix}.$$

Therefore, for example, if the initial "wealth" is 1, the probability of getting rich before bankruptcy is $p^2/(1-p+p^2)$.

Example 9.1.3. In the "Tennis" example, the transition matrix is 20×20, with only 2 recurrent states (each in its own class). The matrix given in Lecture 8 is already in the canonical form (the recurrent states correspond to the first two rows/columns). In order to get $(I-Q)^{-1}$, we need to invert an 18×18 matrix. This is a job for a computer, and we use the Python package markov362m (you may need to consult its documentation to understand all the functions used below) for the particular case $p = 0.45$:

```
In [1]: from markov362m import *
        import numpy as np

In [2]: m = tennis(p=0.45)

In [3]: a0_T = m.dict_to_T_row({"0-0" : 1})

In [4]: a0_T @ m.F @ m.R

Out[4]: array([0.377, 0.623])
```

The numbers 0.377 and 0.623 are the probabilities for wining the entire game for Serena and Roger (in that order).

If we perform the same computation for a range of values of p and plot the results, we get the picture below:

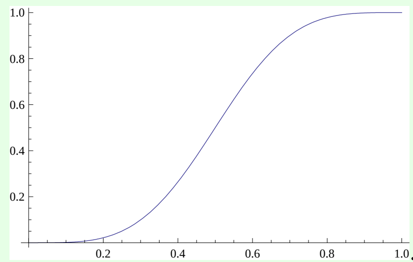

Using a symbolic software package (like *Mathematica*) we can even get an explicit expression for the win probability in this case:

$$u_{(0,0)\,S_{win}} = p^4 + 4p^4q + 10p^4q^2 + \tfrac{20p^5q^3}{1-2pq}.$$

Actually, you don't really need computers to derive the expression above. Can you do it by finding all the ways in which the game can be won in $n = 4, 5, 6, 8, 10, 12, \dots$ rallies, computing their probabilities, and then adding them all up?

9.2 Expected reward

Suppose that each time you visit a transient state $j \in T$ you receive a *reward* $g(j) \in \mathbb{R}$. The name "reward" is a bit misleading since the negative $g(j)$ corresponds more to a fine than to a reward; it is just a name, anyway. Can we compute the expected total reward before absorption

$$v_i = \mathbb{E}_i[\sum_{n=0}^{\tau_C - 1} g(X_n)]?$$

And if we can, what is it good for? Many things, actually, as the following two special cases show:

1. If $g(j) = 1$ for all $j \in T$, then v_i is the expected time until absorption. We will calculate $v_{(0,0)}$ for the "Tennis" example to compute the expected duration of a tennis game.
2. If $g(k) = 1$ and $g(j) = 0$ for $j \neq k$, then v_i is the expected number of visits to the state k before absorption. In the "Tennis" example, if $k = (40, 40)$, the value of $v_{(0,0)}$ is the expected number of times the score $(40, 40)$ is seen in a tennis game.

We compute v_i using the first-step decomposition:

$$\begin{aligned} v_i &= \mathbb{E}_i[\sum_{n=0}^{\tau_C-1} g(X_n)] = g(i) + \mathbb{E}_i[\sum_{n=1}^{\tau_C-1} g(X_n)] \\ &= g(i) + \sum_{k\in S} \mathbb{E}_i[\sum_{n=1}^{\tau_C-1} g(X_n)|X_1 = k]\mathbb{P}_i[X_1 = k] \\ &= g(i) + \sum_{k\in S} p_{ik}\mathbb{E}_i[\sum_{n=1}^{\tau_C-1} g(X_n)|X_1 = k] \end{aligned} \tag{9.2.1}$$

If $k \in T$, then the Markov property implies that

$$\mathbb{E}_i[\sum_{n=1}^{\tau_C-1} g(X_n)|X_1 = k] = \mathbb{E}_k[\sum_{n=0}^{\tau_C-1} g(X_n)] = v_k.$$

When $k \notin T$, then

$$\mathbb{E}_i[\sum_{n=1}^{\tau_C-1} g(X_n)|X_1 = k] = 0,$$

because we have "arrived" and no more rewards are going to be collected. Therefore, for $i \in T$ we have

$$v_i = g(i) + \sum_{k\in T} p_{ik}v_k.$$

If we organize all v_i and all $g(i)$ into column vectors $v = (v_i, i \in T)$, $g = (g(i), i \in T)$, we get

$$v = Qv + g, \text{ i.e., } v = (I - Q)^{-1}g = F = Fug.$$

Having derived the general formula for various rewards, we can provide another angle to the interpretation of the fundamental matrix given in Proposition 9.1.1 above. Let us pick a transient state j and use the reward function g given by

$$g(k) = \mathbf{1}_{\{k=j\}} = \begin{cases} 1, & k = j \\ 0, & k \neq j. \end{cases}$$

By the discussion above, the i^{th} entry in $v = (I - Q)^{-1}g$ is the expected reward when we start from the state i. Given the form of the reward function, v_i is the expected number of visits to the state j when we start from i. On the other hand, as the product of the matrix $F = (I - Q)^{-1}$ and the vector $g = (0, 0, \ldots, 1, \ldots, 0)$, v_i is nothing but the (i, j)-entry in $F = (I - Q)^{-1}$.

Example 9.2.1. We continue the analysis of the "Tennis" chain from Example 9.1.3. We set $g(i) = 1$ for all transient states i to find the expected duration of a tennis game. We use markov362m to perform the computation in the case $p = 0.45$:

```
In [1]: from markov362m import *
        import numpy as np

In [2]: m = tennis(p=0.45)

In [3]: a0_T = m.dict_to_T_row({"0-0" : 1})
        ones = m.dict_to_T_column({}, value_for_omitted=1)

In [4]: a0_T @ m.F @ ones

Out[4]: array([6.681])
```

.

Therefore, a game between fairly equally matched opponents lasts 6.68 rallies on average. The same computation for a range of values of p yields the following picture

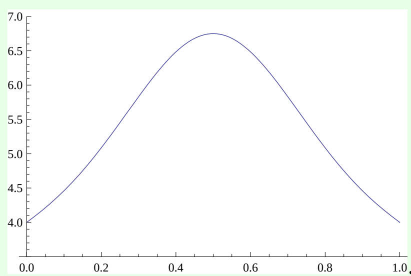

.

When $p = 0$, the course of the game is totally predictable and Serena wins in 4 points. The same holds when $p = 1$, only it is Roger who wins with probability 1 this time. In between, we see that the expected game-length varies between 4 and about 7 (actually, the exact number is 6.75), and it longest when $p = \frac{1}{2}$.
How about the expected number of deuces (scores $(40, 40)$)? We can compute that too by setting $g(i) = 0$ if $i \neq (40, 40)$ and $g((40, 40)) = 1$. Using markov362m again we obtain that the expected value is 0.6 when $p = 0.45$:

```
In [1]: from markov362m import *
        import numpy as np

In [2]: m = tennis(p=0.45)

In [8]: a0_T = m.dict_to_T_row({"0-0" : 1})
        g = m.dict_to_T_column({"40-40" : 1}, value_for_omitted=0)

In [7]: a0_T @ m.F @ g

Out[7]: array([0.6])
```

As a function of p, this expected value looks like this:

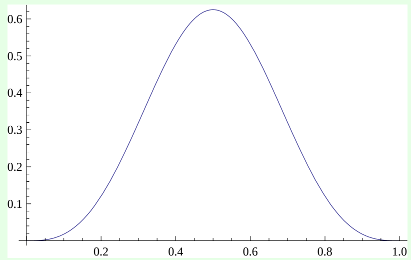

The plot of the obtained expressions, as a function of p, is given on the right. Therefore, the expected number of deuces varies between 0 and a bit more than 0.6 (the exact number is 0.625 and corresponds to the case $p = \frac{1}{2}$). When asked, people would usually give a higher estimate for this probability. The reason is that the distribution of the number of deuces looks something like the picture below (a histogram of the number of deuces in a simulation of 1000 tennis games with $p = \frac{1}{2}$). We see that most of the games have no deuces. However, in the cases where a deuce happens, it is quite possible it will be repeated.

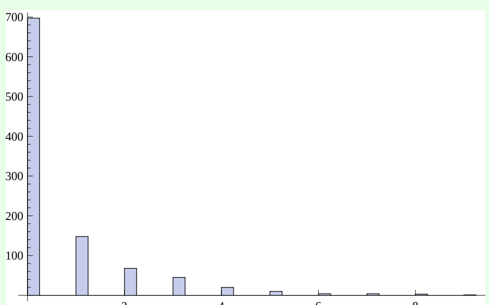

In fact, the expected number of deuces *given that the game contains at least one deuce* is approximately equal to 2.1.

We end with another example from a different area:

Example 9.2.2. Alice plays the following game. She picks a pattern consisting of three letters from the set $\{H, T\}$, and then tosses a fair coin until her patterns appears for the first time. If she has to pay \$1 for each coin toss, what is the expected cost she is going to incur? What pattern should she choose to minimize that cost?

We start by choosing a pattern, say HTH, and computing the number of coin tosses Alice expects to make before it appears.

This is just the kind of computation that can be done using our absorption-and-reward techniques, if we can find a suitable Markov chain. It turns out that the following will do:

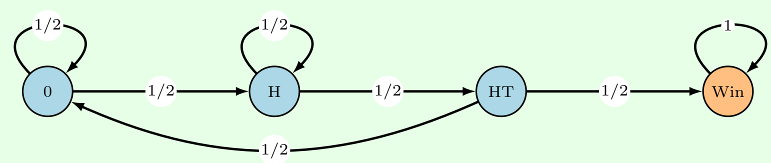

As Alice tosses the coin, she keeps track of the largest initial portion of her pattern that appears at last several places of the sequence of past tosses. The state 0 represents no such portion (as well as the intial state), while HT means that the last two coin tosses were H and T (in that order) so that it is possible to end the game by tossing a H next. On the other hand, if the last toss was a T, there is no need to keep track of that - it is as good as 0.

Once we have this chain, all we have to do is perform the absorption and reward computation with the reward function $g \equiv 1$. The Q-matrix of this chain (with the transient states ordered as $0, H, HT$) is

$$Q = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 \\ 1/2 & 0 & 0 \end{bmatrix}$$

and the fundamental matrix F turns out to be

$$F = \begin{bmatrix} 4 & 4 & 2 \\ 2 & 4 & 2 \\ 2 & 2 & 2 \end{bmatrix}.$$

Therefore, the required expectation is the sum of all the elements on the first row, i.e., 10.

Let us repeat the same for the pattern HHH. We build a similar Markov chain:

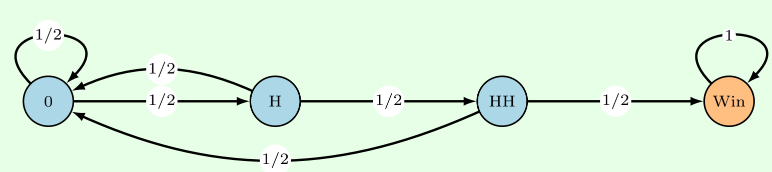

We see that there is a subtle difference. One transition from the state H, instead of going back to itself, is directed towards 0. It is clear from here, that this can only increase Alice's cost. Indeed, the fundamental matrix is now given by

$$F = \begin{bmatrix} 8 & 4 & 2 \\ 6 & 4 & 2 \\ 4 & 2 & 2 \end{bmatrix},$$

and the expected number of tosses before the first appearance of HHH comes out as 14.

Can you do this for other patterns? Which one should Alice choose to minimize her cost?

9.3 Problems

Problem 9.3.1. In a Markov chain with a finite number of states, the fundamental matrix is given by

$$F = \begin{bmatrix} 3 & 4 \\ \frac{3}{2} & 4 \end{bmatrix}.$$

The initial distribution of the chain is uniform on all transient states. Compute the expected value of

$$\tau_C = \min\{n \in \mathbb{N}_0 \,:\, X_n \in C\},$$

where C denotes the set of all recurrent states.

Problem 9.3.2. Let $\{X_n\}_{n\in\mathbb{N}_0}$ be a Markov chain with the following transition matrix

$$P = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 1/3 & 1/3 & 1/3 \\ 0 & 0 & 1 \end{bmatrix}$$

Suppose that the chain starts from the state 1.

1. What is expected time that will pass before the chain first hits 3?
2. What is the expected number of visits to state 2 before 3 is hit?
3. Would your answers to 1. and 2. change if we replaced values in the third row of P by any other values (as long as P remains a stochastic matrix)? Would 1 and 2 still be transient states?
4. Use the idea of part 3. to answer the following question. What is the expected number of visits to the state 2 before a Markov chain with transition matrix

$$P = \begin{bmatrix} 17/20 & 1/20 & 1/10 \\ 1/15 & 13/15 & 1/15 \\ 2/5 & 4/15 & 1/3 \end{bmatrix}$$

hits the state 3 for the first time (the initial state is still 1)? Remember this trick for the next exam.

Problem 9.3.3. A basketball player is shooting a series of free throws. The probability of hitting any single one is $1/2$, and the throws are independent of each other. What is the expected number of throws the player will attempt before hitting 3 free throws in a row (including those 3)? **Note:** Instead of using (and computing) the fundamental matrix, write down a system of equations and solve it directly.

Problem 9.3.4. A fair 6-sided die is rolled repeatedly, and for $n \in \mathbb{N}$, the outcome of the n-th roll is denoted by Y_n (it is assumed that $\{Y_n\}_{n\in\mathbb{N}}$ are independent of each other). For $n \in \mathbb{N}_0$, let X_n be the remainder (taken in the set $\{0, 1, 2, 3, 4\}$) left after the sum $\sum_{k=1}^n Y_k$ is divided by 5, i.e. $X_0 = 0$, and

$$X_n = \sum_{k=1}^{n} Y_k \pmod 5, \text{ for } n \in \mathbb{N},$$

making $\{X_n\}_{n\in\mathbb{N}_0}$ a Markov chain on the state space $\{0, 1, 2, 3, 4\}$ (no need to prove this fact).

1. Write down the transition matrix of the chain.
2. Classify the states, separate recurrent from transient ones, and compute the period of each state.
3. Compute the expected number of rolls before the first time $\{X_n\}_{n\in\mathbb{N}_0}$ visits the state 2, i.e., compute $\mathbb{E}[\tau_2]$, where

$$\tau_2 = \min\{n \in \mathbb{N}_0 : X_n = 2\}.$$

4. Compute $\mathbb{E}[\sum_{k=0}^{\tau_2-1} X_k]$.

(*Note:* For parts 3. and 4., you can phrase your answer as $(B^{-1}C)_{ij}$, where the matrices B and C, as well as the row i and the column j have to be given explicitly. You don't need to evaluate B^{-1}.)

Problem 9.3.5. Let $S = \{0, 1, 2, \dots, a\}$ be the state space of a *Gambler's ruin*, where a gambler starts with the "wealth" $\$x \in \{\$1, \$2, \dots, \$a - 1\}$ and tries to make $\$a$ before going bankrupt (hitting \$0). Unlike in the classical Gambler's ruin, the player does not bet \$1 each time. She plays boldly and bets the maximum possible amount that makes sense. More precisely, if her current wealth is x, she bets either x or $a - x$, whatever is smaller. Assume that each game is a coin toss with probability $p \in (0, 1)$ of getting *heads*. If that happens, the gambler's wealth increases by the amount she bet. If the coin shows *tails* she loses the bet, and her wealth decreases by the same amount. As in the classical Gambler's ruin, the gambler stops when either 0 or a is reached (and we assume that her wealth remains at that level forever).

<u>ASSUME THAT $a = 5$</u>.

1. write down the transition matrix of this Markov Chain and sketch its transition graph. Identify transient and recurrent states, and find periods of all states.

2. compute the probability of winning (hitting 5 before hitting 0) for the gambler who starts with \$1. **Note:** Write equations directly - do not use (and invert) matrices.

Problem 9.3.6. The fundamental matrix associated to a finite Markov chain is $F = \begin{pmatrix} 3 & 3 \\ 3/2 & 3 \end{pmatrix}$, with the first row (and column) corresponding to the state A and the second to B. Some of the following statements are true and the others are false. Find which ones are true and which are false; give explanations for your choices.

(a) The chain has 2 recurrent states.

(b) If the chain starts in A, the expected number of visits to B before hitting the first recurrent state is 3.

(c) If the chain is equally likely to start from A or B, the expected number of steps it will take before it hits its first recurrent state is $\frac{21}{4}$.

(d) $\mathbb{P}_A[X_1 = C] = 0$ for any recurrent state C.

Problem 9.3.7. Let $\{Y_n\}_{n\in\mathbb{N}_0}$ be a sequence of die-rolls, i.e., a sequence of independent random variables with distribution

$$Y_n \sim \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \end{pmatrix}.$$

Let $\{X_n\}_{n\in\mathbb{N}_0}$ be a stochastic process defined by $X_n = \max(Y_0, Y_1, \dots, Y_n)$. In words, X_n is the maximal value rolled so far.

1. Explain why X is a Markov chain, and find its transition matrix and the initial distribution.

2. Supposing that the first roll of the die was 3, i.e., $X_0 = 3$, what is the expected time until a 6 is reached?

3. Under the same assumption as above ($X_0 = 3$), what is the probability that a 5 will not be rolled before a 6 is rolled for the first time?

4. Starting with the first value $X_0 = 3$, each time a die is rolled, the current record (the value of X_n) is written down. When a 6 is rolled for the first time all the numbers are added up and the result is called S (the final 6 is not counted). What is the expected value of S?

(*Hint:* You can eliminate 1 and 2 from the state space for parts 2. and 3. Also, use the fact that

$$A^{-1} = \begin{bmatrix} \frac{1}{a} & \frac{r}{ab} & \frac{rb+r^2}{abc} \\ 0 & \frac{1}{b} & \frac{r}{bc} \\ 0 & 0 & \frac{1}{c} \end{bmatrix} \text{ when } A = \begin{bmatrix} a & -r & -r \\ 0 & b & -r \\ 0 & 0 & c \end{bmatrix},$$

for all $a, b, c, r > 0$.)

Problem 9.3.8. Go back to Problem **??** (the one with the professor and umbrellas in the first lecture on Markov chains). One way to model it using a Markov Chain is like this:

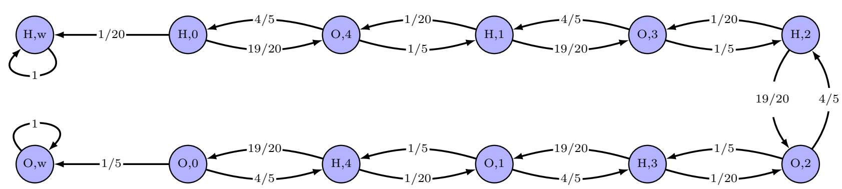

Consult the solution to Problem **??** for further explanation.

Assuming that on day 0, before the professor leaves home in the morning, there are two umbrellas at home and two in the office.

1. What is the expected number of days before she gets wet?
2. What is the probability that the first time the professor does get wet, it will be on her way home from the office?

Note: This problem is best done using numerical methods (software). If you choose to use the package markov362m, you will find the chain already implemented under the name professor.

Problem 9.3.9. Go back to Problem **??** (the one with airline reservation system and the computer-repair facility). Assuming that $p = 0.3$, and that the system starts with both computers operational, compute the expected time before both computers break down. Assuming that each day with only one computer operational costs the company $\$10,000$, what is the total cost the company is expected to incur before both computers break down? **Note:** This problem is best done using numerical methods (software). If you choose to use the package markov362m, you will find the chain already implemented under the name facility.

Problem 9.3.10. A zoologist, Dr. Gurkensaft, claims to have trained Basil the Rat so that it can aviod being shocked and find food, even in highly confusing situations. Another scientist, Dr. Hasenpfeffer does not agree. She says that Basil is stupid and cannot tell the difference between food and an electrical shocker until it gets very close to either of them.

The two decide to see who is right by performing the following experiment. Basil is put in the compartment 3 of a maze that looks like this:

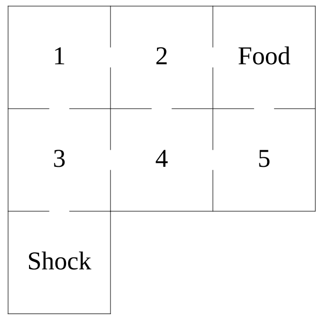

Dr. Gurkensaft's hypothsis is that, once in a compartment with k exits, Basil will prefer the exits that lead him closer to the food. Dr. Hasenpfeffer's claim is that every time there are k exits from a compartment, Basil chooses each one with probability $1/k$.

After repeating the experiment 100 times, Basil got shocked before getting to food 52 times and he reached food before being shocked 48 times. Compute the theoretical probabilities of being shocked before getting to food, under the assumption that Basil is stupid (all exits are equally likely). Compare those to the observed data. Whose side is the evidence on (no need for any statistical analysis; simply state what you think) ? **Note:** This problem is best done using numerical methods (software). Unlike in the previous problems, even if you choose the package markov362m, you will have to build the chain yourself.