UNIVERSIDADE ESTADUAL DE CAMPINAS

MS777 PROJECT

Markov chains in a field of traps: theory and simulation

Author:
Marina VASQUES MOREIRA

Supervisor:
Marina VACHKOVSKAIA

Abstract

In this work we make a review of the main ideas of Markov chains theory, giving some examples as motivation.We then proceed in examining some statements related to Markov chains in quenched and annealed fields. After that, we present computer simulations and compare these results with the ones proved theoretically.

Contents

1 Introduction

"Predictions can be very difficult – especially about the future."
Niels Bohr

In everyday life, we deal with probabilistic phenomena: even those among us without formal mathematical training are acquainted with some features of the theory of probability. People in general develop an intuitive sense of chance and randomness while dealing with such things as tossing coins, lottery games and weather forecasting. The last one is a good example of this intuitive sense: while most of us are not aware of all the mechanisms involved in weather forecasting, we accept that a 95% accurate prediction of a sunny day is a good reason not to wear a raincoat (bear in mind here the difference between accuracy and precision: while accuracy refers to how close the measured value is to the true or accepted value, precision refers to how close together a group of measurement actually are to each other. Precision is usually expressed trough significant digits. In the case of weather forecasting, an accurate forecasting is usually preferable over a precise one: if we are told that the temperature tomorrow may vary from 20°C to 25 °C, and it turns out that we measure it as 23°C, that is just ok. On the other hand, a very precise forecasting may predict a temperature of 23.43°C for tomorrow, and if it turns out to be 32°C, that was a very precise but highly inaccurate – not to mention useless – prediction). More than that: we also carry with us the intuitive notion that, regardless of geographical location or time of the year, we can make a good guess that the weather tomorrow is somehow related to the weather today (of course location and time of the year matter, but we do make statements about the weather tomorrow based on the weather today). In theory, one could say that the current weather is a result of all the climate processes that have taken place on Earth since its origin billions of years ago[1], but even if we had access to this enormous amount of information, that would not make quite a very useful model. Instead, let's try to think of a model where the predictability of the weather today is influenced by the weather observed yesterday, and so on.

With that in mind, we could define some possible weather states (such as "rainy", "sunny", "cloudy") and make up some probabilities. The way we defined our model, the probabilities should be of the kind

"As today is sunny, P is the probability that the weather tomorrow is going to be rainy".

After computing all the probabilities that describe a transition from one weather state to the other, we could predict, for any two given weather states i and j, the probability of the weather – today in state i – being tomorrow in state j. We could also repeat the same calculations and try to predict the weather some days ahead with the information we have today; we could go further and, carrying on this result, have an idea of the long term prediction (which is just a prediction for the weather a big lot of days ahead!).

Strange as it may seem, the apparent simplicity of the model we have just outlined is in fact its great advantage: different phenomena can be modeled in such a way, as this model doesn't carry many assumptions on the nature of the problem, provided only that we can define some states, and the outcome of one experiment is influenced solely by the previous one. This is the heart of the definition of a *Markov chain*, named after the Russian mathematician Andrey Markov. We shall discuss some of the main ideas and definitions of Markov chains in the next section.

[1] This is usually referred to as *determinism*. Laplace is often misquoted as having said"give me the positions and velocities of all the atoms in the universe, and the forces acting upon them, one can predict the future course of the universe." But what he truly said is somehow deeper than that, as it can be seen on this excerpt from one of his works: "Thus, we must consider the present state of the universe as the effect of its previous state and as the cause of those states to follow. An intelligent being which, for a given point in time, knows all the forces acting upon the universe and the positions of the objects of which it is composed, supplied with facilities large enough to submit these data to numerical analysis, would include in the same formula the movements of the largest bodies of the universe and those of the lightest atom. Nothing would be uncertain for it, and the past and future would be known to it. (...) All the efforts of human intellect to search for the truth tend to bring it steadily closer to the intelligent being described above, but this being will remain always infinitely distant."

2 Markov Chains

2.1 Basic ideas and definitions

In this chapter, we define Markov chains and discuss some of its properties. The fundamentals of this theory have been widely discussed on many good books; most of the definitions given here are adpated from "The Theory of Probability", written by B. V. Gnedenko.

A Markov chain is a collection of random variables $\{X_t\}$ (where the index t runs through 0, 1, ...) having the property that, given the present, the future is conditionally independent of the past. Frequently, Markov chains are defined in terms of a given physical system S, which at each instant of time can be in one of the states $A_1, A_2, ..., A_k$ and alters its state only at times $t_1, t_2, ..., t_n, ...$ For Markov chains, the probability of passing to some state $A_i(i = 1, 2, ..., k)$ at a time $\tau(t_s < \tau < t_{s+1})$ depends only on the state the system was in at time $t(t_{s-1} < t < t_s)$ and does not change if we learn what its states were at earlier times. This is the same as saying that Markov chains are memoryless ("future" is independent of "past" given "present").

In the following examples we illustrate this definition:

Example 1. *Gambler's ruin* – Let two players each have a finite number of cents (say, n_A for player A and n_B for player B). Now, flip one of the cents (from either player), with each player having 50% probability of winning, and transfer a cent from the loser to the winner. Repeat the process until one player has all the cents. In this example, we may define a collection of random variables A_t, in which we store the information of the amount of cents with player A at time t, where the index t represents the successive tossing of the coins, $t = 1, 2, ...$ Similarly, define B_t the amount of cents with player B at time t. Then, the possible states for this system are represented by the ordered pairs (A_t, B_t), which we may call $X_t = (A_t, B_t)$. Observe that the game stops when one of the players has all the cents, there is, the game stops when $X_t = (n_A + n_B, 0)$ or $X_t = (0, n_A + n_B)$. [2]

This is an example of a random walk, which is a mathematical formalisation of a trajectory that consists of taking successive random steps. Random walks are usually assumed to be Markov chains. Other examples include the path traced by a molecule as it travels in liquid or gas, the search path of a foraging animal (assuming that the animal takes random steps and does not retrieve information from the past), the price of a fluctuating stock.

Example 2. *Hydrogen atom* – In Bohr's model of the hydrogen atom, the electron can be in one of the allowed orbits. Denote by A_i the event that the electron lies on the *i*th orbit. Further, assume that changes in the state of the atom can only occur at times $t_1, t_2, ...$ (actually, these times are random quantities). The probability of transition from the *i*th orbit to the *j*th at time t_s depends only on i and j (the difference $j - i$ depends on the amount of energy by which the charge of the atom changed at time t_s) and does not depend on the orbits the electron occupied in the past.

Example 3. *Branching process* Suppose that a population evolves in generations, and let Z_n be the number of members of the nth generation. Each member of the nth generation giver birth to a family, possibly empty, of members of the $(n + 1)$ th generation; the size of this family is a random variable. We make the following assumptions about these family sizes:

(a) the family sizes of the individuals of the branching process form a collection of independent random variables;

[2]The probability that one of the players will eventually lose all his money approaches 100% as time increases. It can be shown that the chances P_A and P_B that the players A and B, respectively, will end with no cents are: $P_A = n_B/(n_A + n_B), P_B = n_A/(n_A + n_B)$. Thus, the player starting with the smallest number of cents has the greater chance of going bankrupt. The longer you gamble, the longer the chance that the player who started with more cents will win the game. That's why casinos always come out ahead in the long run, as they start with more money than the players.

(b) all family sizes have the same probability mass function f and generating function G.

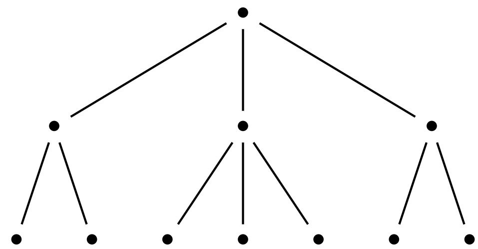

Figure 1: The family tree of a branching process.

We are interested in the random sequence $Z_0, Z_1, ...$ of generation sizes Let $G_n(s) = \mathbf{E}(s^{Z_n})$ be the generatinf function of Z_n. It can be shown that $G_{m+n}(s) = G_m(G_n(s)) = G_n(G_m(s))$ and thus $G_n(s) = G(G(...(G(s))...))$ is the n-fold iterate of G (see Grimmett for a detailed demonstration). So, the branching process random variables Z_n form a Markov chain – since conditional on the number of the generation size at the nth generation, the future values do not depend on the previous values. The Z_n may take values in $S = (0, 1, 2, ...)$ and the transition probabilities p_{ij} are given by the coefficients of S^j in $G(s)^j$.

We say a Markov chain is *homogeneous* when the conditional probability of the occurrence of an event$j$ in the (n+1)st trial, provided that in the nth trial the event j occurred, does not depend on the number of the trial. We call this probability the *transition probability* and denote it by p_{ij}: in this notation, the first subscript always denotes the result of the previous trial, and the second indicates the state into which the system passes in the subsequent instant of time. Mathematically, a Markov chain is homogeneous if:

$$P(X_{n+1} = j | X_n = i) = P(X_1 = j | X_0 = i)$$

That is, the probability $p_{ij} = P(X_{n+1} = j | X_n = i)$ does not depend on n for an homogeneous Markov chain.

The total probabilistic picture of possible change that occur during a transition from one trial to the immediately following is given by the matrix:

$$\pi_1 = \begin{pmatrix} p_{11} & p_{12} & \cdots & p_{1k} \\ p_{21} & p_{22} & \cdots & p_{2k} \\ \dots & \dots & \dots & \dots \\ p_{k1} & p_{k2} & \cdots & p_{kk} \end{pmatrix}$$

This matrix is called *transition matrix (matrix of transition probabilities)*. Being probabilities, all elements of this matrix must be non-negative numbers, i.e., for all i and j,

$$0 \leqslant p_{ij} \leqslant 1$$

This matrix contains useful information about the system: if all the elements in the main diagonal are null, that means the systems never stops at one state, that is to say, once the system reaches state i, it must go to another state in the next step. That is the case in example 1; if at a given time t players A and B each have a given amount of cents, they will both have different amounts at time $t + 1$, for at every play one of the players win and the other loses a coin. It should also be noted that the sum of the elements of each row is equal to unity. This comes from the fact that from state A_i at time t, the system must go to one and only one of the states A_j. This fact can be written as:

$$\sum_{j=1}^{k} p_{ij} = 1(i = 1, 2, ..., k)$$

Sometimes it's useful to analyze these transition probabilities with a diagram. In the diagram below the circles indicate states of the system and the arrows indicate the direction of the transition from one state to another. When $p_{ij} = 0$ the corresponding arrow is omitted.

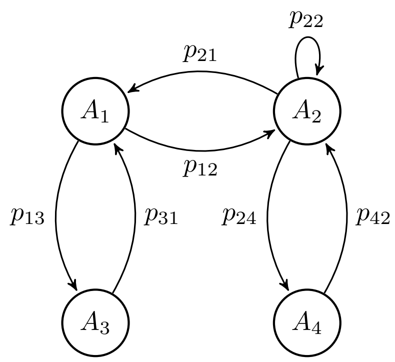

Figure 2: Diagram with transition probabilities for a discrete Markov chain

Let's analyze other interesting aspects of the transition matrix in the following examples.

Example 3. *Bernoulli process* – Take a sequence of independent random variables $X_1, X_2, ..., X_n$ such that for each i, the value of X_i is either o or 1, and for all values of i, the probability that $X_i = 1$ is the same number p. One of the many examples of this process is the tossing of a fair coin, if we define p as the probability of coming up "TAILS", and define "TAILS" as 1 and "HEADS" as 0. Then we have, for all i, the probabilities $P(X_i = 1) = p$ and $P(X_i = 0) = 1 - p = q$. The way we defined the Bernoulli process, we have:

$$P(X_{n+1} = j/X_1, X_2, ..., X_n) = P(X_{n+1} = j) = P(X_{n+1} = j/X_n)$$

Thus, the Bernoulli process is a Markov chain, since the future of the process is independent of the past (incidentally, it is also independent of the present). The transition matrix for this process is given by:

$$\pi = \begin{pmatrix} q & p \\ q & p \end{pmatrix}$$

Any process in which all variables are independent is a Markov chain (that's a special case, since the variables are even independent of the current state). A Bernoulli scheme – a generalization of the Bernoulli process to more than two possible outcomes – is also a special case of Markov chains, just like the Bernoulli process. An interesting property for Bernoulli schemes is that their transition matrices have identical rows. This comes directly from the fact that the next state is not only independent of the past, but, as we mentioned before, it is also independent of the current state.

Example 4. Suppose we have a system S with possible states A_1, A_2, A_3 whose probability transitions are described by the transition matrix below.

$$\pi = \begin{pmatrix} p_{11} & p_{12} & p_{13} \\ p_{21} & p_{22} & p_{23} \\ p_{31} & p_{32} & p_{33} \end{pmatrix}$$

Suppose also that the Markov chain is *homogeneous*, that is, the transition probabilities have no relation with the number of the trial. We are interested in knowing the probability that the system will be at state A_3 two steps ahead if it starts at state A_1. We shall denote this probability as $p_{13}^{(2)}$. A quick look shows us that state A_3 can be reached through states A_1 and A_2, and the probability we look for can be found considering the possible ways we may get from state 1 to 3 passing through the other states. Then, the probability $p_{13}^{(2)}$ is:

$$p_{13}^{(2)} = p_{11}p_{13} + p_{12}p_{23} + p_{13}p_{33}$$

In general, for a system with k states, we have:

$$p_{ij}^{(2)} = \sum_{r=1}^{k} p_{ir}p_{rj}$$

It can be shown that if a Markov chain has a transition matrix π, the *ij*th element $p_{ij}^{(n)}$ of the matrix π^n gives the probability that the Markov chain will be in state j, starting from state i, after n steps. An outline of the demonstration is given: Suppose we have a systems which has states $A_1, A_2, ..., A_k$.Let's denote as $A_i^{(s)}$ the event that the system is in state i at time s. We want to find the probability that the system will be at state j after n steps, that is, we want to find the probability of the event $A_j^{(s+n)}$. This probability was denoted above as $p_{ij}^{(n)}$. We now examine some intermediate trial with the number $(s+m)$. One of the possible events $A_r^{(s+m)}$, com $(1 \leqslant r \leqslant k)$ will occur in this trial, with probability $p_{ir}^{(m)}$. The probability of transition from state $A_r^{(s+m)}$ to state $A_j^{(s+n)}$ is $p_{rj}^{(n-m)}$. By the formula of total probability,

$$p_{ij}^{(n)} = \sum_{r=1}^{k} p_{ir}^{(m)} p_{rj}^{(n-m)} \tag{2.1}$$

This is the *Kolmogorov-Chapman equation.*

We denote by π_n the transition matrix after n trials:

$$\pi_n = \begin{pmatrix} p_{11}^{(n)} & p_{12}^{(n)} & \dots & p_{1k}^{(n)} \\ p_{21}^{(n)} & p_{22}^{(n)} & \dots & p_{2k}^{(n)} \\ \dots & \dots & \dots & \dots \\ p_{k1}^{(n)} & p_{k2}^{(n)} & \dots & p_{kk}^{(n)} \end{pmatrix}$$

According to the formula of the total probability given above, the following relation holds between the matrices π_s with different subscripts:

$$\pi_n = \pi_m \cdot \pi_{n-m} (0 < m < n)$$

In particular, for $n = 2$, we find:

$$\pi_2 = \pi_1 \cdot \pi_1 = \pi_1^2$$

For n=3:

$$\pi_3 = \pi_1 \cdot \pi_2 = \pi_1 \cdot \pi_1^2 = \pi_1^3$$

And generally, for any n:

$$\pi_n = \pi_1^n$$

We note a special case of formula (1) for $m = 1$:

$$p_{ij}^{(n)} = \sum_{r=1}^{k} p_{ir}^{(m)} p_{rj}^{(n-1)} \tag{2.2}$$

2.2 Classification of states

A state A_i of a Markov chain is called *absorbing* if it is impossible to leave it (i.e., $p_{ii} = 1$ and $p_{ij} = 0$ for $i \neq j$). A Markov chain is absorbing if it has at least one absorbing state, and if from every state it is possible to go to an absorbing state (not necessarily in one step). In Example 1 (the gambler's ruin), we have two absorbing states, one corresponding to the situation when player A gets all the cents and wins the game, and the other is the situation when player B wins. Note that this is an example of an absorbing Markov chain: at any moment of the game, we can say that, with some probability, one of the players will get all the cents at a time (t_n), thus winning the game and reaching an absorbing state. We may not all argue that this will eventually happen – otherwise, there wouldn't be so many ruined gamblers! – but it is clear that it is possible, at any given moment, that one of the players will win the game after a numbers of plays. In fact, it can be proven that for an absorbing Markov chain, the probability that the process will be absorbed is 1.

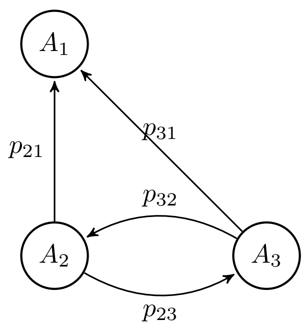

Figure 3: State A_1 is an absorbing state.

We say a state A_j is accessible from a state A_i ($A_i \longleftrightarrow A_j$) if a system started in state A_i has a non-zero probability of transitioning into state A_j at some point; that is, for some $n > 0$:

$$p_{ij}^{(n)} > 0$$

A state A_i has period d if any return to state A_i must occur in multiples of d time steps. If $d = 1$, then the state is said to be aperiodic: returns to state A_i can occur at irregular times.

A state A_i is called *unessential* (or *transient*) if there exist A_j and n such that $p_(ij)^{(n)} > 0$, but $p_(ji)^{(m)} = 0$ for all m. Thus, an unessential state possesses the property that it is possible, with positive probability, to pass from it to another state, but it is no longer possible to return from that state to the original (unessential) state.

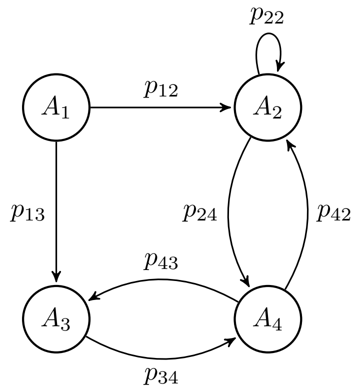

Figure 4: State A_1 is a transient (unessential) state. All other states are recurrent.

All states not transient are called *recurrent* or *persistent*. From the definition it follows that if the states A_i and A_j are recurrent, then for all j such that ($A_i \longleftrightarrow A_j$), it is also true that ($A_j \longleftrightarrow A_i$). Another way of saying this is that a state is recurrent if there exist positive m and n such that along with the inequality $p_{ij}^{(m)} > 0$, the inequality $p_{ji}^{(n)} > 0$ also holds. Recurrent states have finite hitting time with probability 1.

If A_i and A_j are such that for both of them these inequalities hold, given certain m and n, then they are called *communicating* (they are said to communicate). Every state communicates with itself, that is, for all i we have $p_{ii}^{(0)} = 1$. It is clear that if A_i communicates with A_j and A_j communicates with A_k, then A_i also communicates with A_k. Thus, all essential states can be partitioned into *classes* such that all states belonging to a single class communicate and those belonging to different classes do not communicate. All states of one and the same class have one and the same period.

2.3 Long term behavior

Usually, we are interested in the long term behavior of a system. In this section we are going to state some results using the previous definitions. We want to find the probability that the state is going to be in a certain state j after n steps – for a very large n – starting from a initial state i, that is, we want to find $\lim_{x\to\infty} p_{ij}^{(n)}$. If this probability is independent of the initial state, it converges to a stationary probability $p_j > 0$. For a Markov chain with k states and transition matrix π, all states are recurrent if and only if there is a solution for:

$$p_j = \sum_{i=1}^{k} p_i \cdot p_{ij}, j = 1, 2, ..., k$$

$$\sum_{j=1}^{k} p_j = 1$$

If the system above has a solution, it is unique, and $p_j > 0 \forall j$. We can also interpret p_j as the relative frequency of occurrence of state j in the long-term behavior of the system.

Example 5. *A strange climate* – In the Land of Oz, there are never two nice days in a row. A nice day may be followed by a rainy or a snowy day with equal probabilities; a snowy day has an even chance of being followed by another snowy day (and the same happens for rainy days). If there is change from snow or rain, only half of the time this a change to a nice day. Dorothy is thinking about moving from Kansas to the Land of Oz, and as she was raised in Western Kansas – which exhibits a semiarid climate – she is very curious about the strange climate of the Land of Oz, so she asked the wonderful wizard of Oz to do some calculations for her – because matrix multiplication is one of the wonderful things he does. What is the long term behavior of the climate in the Land of Oz?

The states for this system are the kinds of weather: rainy (R), nice (N) and snowy (S). The transition matrix may be written as:

$$\pi = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 1/2 & 1/4 & 1/4 \\ 1/2 & 0 & 1/2 \\ 1/4 & 1/4 & 1/2 \end{pmatrix} \end{array}$$

We showed in the previous section that the probabilities of reaching states from a initial state in n steps can be computed by calculating the nth power of the transition matrix. Thus, for the Land of Oz we have:

$$\pi_2 = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 0.4375 & 0.1875 & 0.3750 \\ 0.3750 & 0.2500 & 0.3750 \\ 0.3750 & 0.1875 & 0.4375 \end{pmatrix} \end{array}$$

$$\pi_3 = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 0.4063 & 0.2031 & 0.3906 \\ 0.4063 & 0.1875 & 0.4063 \\ 0.3906 & 0.2031 & 0.4063 \end{pmatrix} \end{array}$$

$$\pi_4 = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 0.4023 & 0.1992 & 0.3984 \\ 0.3984 & 0.2031 & 0.3984 \\ 0.3984 & 0.1992 & 0.4023 \end{pmatrix} \end{array}$$

$$\pi_5 = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 0.4004 & 0.2002 & 0.3994 \\ 0.4004 & 0.1992 & 0.4004 \\ 0.3994 & 0.2002 & 0.4004 \end{pmatrix} \end{array}$$

At this point, Dorothy started to become cheerful, as it seemed the matrices were converging to a certain matrix, so she asked the wizard to run some more calculations:

$$\pi_6 = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 0.4001 & 0.2000 & 0.3999 \\ 0.3999 & 0.2002 & 0.3999 \\ 0.3999 & 0.2000 & 0.4001 \end{pmatrix} \end{array}$$

$$\pi_7 = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 0.4000 & 0.2000 & 0.4000 \\ 0.4000 & 0.2000 & 0.4000 \\ 0.4000 & 0.2000 & 0.4000 \end{pmatrix} \end{array}$$

$\vdots$

$$\pi_{23} = \begin{array}{c} \\ R \\ N \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & N & S \end{array} \\ \begin{pmatrix} 0.4000 & 0.2000 & 0.4000 \\ 0.4000 & 0.2000 & 0.4000 \\ 0.4000 & 0.2000 & 0.4000 \end{pmatrix} \end{array}$$

At this point, she was convinced that she had found the long-term behavior of the climate of the Land of Oz. The probabilities for states R, N, and S, are $p_R = 0.4, p_N = 0.2,$ and $p_S = 0.4$ after n days for a large n. These probabilities are independent of the initial state.

The same results could have been found in a more elegant way – and sparing the wizard from a lot of calculation – by solving:

$$\begin{cases} (p_R \quad p_S \quad p_N) = (p_R \quad p_S \quad p_N) \begin{pmatrix} 1/2 & 1/4 & 1/4 \\ 1/2 & 0 & 1/2 \\ 1/4 & 1/4 & 1/2 \end{pmatrix} \\ p_R + p_S + p_N = 1 \end{cases}$$

3 Markov chains in a field of traps

3.1 Introduction

In 1998, Robin Pemantle and Stanilav Volkov wrote an article entitled *"Markov chains in a field of traps."* Some articles had been written before on the subject of random walks in a random field of traps, but their article encompass results obtained by others and provide some useful generalizations on the subject. Most of the definitions and theorems in this section have been taken directly from this source. We also use the same notation they used, calling the attention of the reader to the occasions when this notation might be a little different from the one we used earlier.

Having established the fundamentals of Markov chains in the previous section, we now make a quick review of their article, focusing on its principal theorems and results. The reader wishing to see all the demonstrations in details may refer to their article. Following this section we present some computer simulations with the purpose to observe the long-term behavior of Markov chains and reproduce the results formally proved by Pemantle and Volkov.

A site is defined as being a trap with probability p. A trap – as the name indicates – is an absorbing state. In the classical problem of the drunkard's walk, we could imagine a manhole as being the trap: once the drunkard falls down the manhole, he stops his walk (a less dangerous example could be the bar or his house, provided that once he reaches any of these places, he stays there indefinitely). Let's now define our variables and notation:

Let S be a countable space and let $p : S^2 \to [01]$ be a set of transition probabilities on S, i.e.,$\sum_y p(x,y) = 1$ for all $x \in S$. Let P_x denote any probability measure such that the random sequence $X_0, X_1, X_2, ...$ is a Markov chain with transition probabilities $\{p(x,y)\}$ starting from $X_0 = x$. Take a function $q : S \to [0,1)$, representing a set of trapping probabilities.

When dealing with traps, two kinds of problems can be formulated as Markov chains:

- Quenched problem – "Site x is a trap with probability $q(x)$ forever". Let $T \subseteq S$ be the random set of traps, i.e., $P_x(\{x_1, ..., x_n\} \subset T) = \Pi_{i=1}^n q(x_i)$ for all finite subsets of S, and T is independent of $X_0, X_1, ...$ We are interested in the quantities

$$\pi(x) = P_x(X_n \in T \text{ for some } n \geqslant 0) \tag{3.1}$$

  We say that the quenched field is trapping or non-trapping, according to whether $\pi(x) = 1$ for all x, or whether $\pi(x) < 1$ for some x.

  In our analogy of the drunkard's walk, let's imagine the city as being a square grid (because S is a countable space). The quenched problem represent the situation of the drunkard wandering through the streets of this city with a fixed amount of open manholes.

- Annealed problem – Site x is a trap with probability $q(x)$, but each unit of time, this probability is updated. This gives us an IID [3] sequence $\{T_n, n \geqslant 0\}$ of trap sets, each distributed as T in the quenched problem. Returning to our drunkard's walk, he is now walking through a city where each manhole may be either open or closed with some probability which is updated each unit of time. The same way we defined for the quenched problem, we are interested in the quantity:

$$\tilde{\pi}(x) = P_x(X_n \in T_n \text{ for some } n \geqslant 0) \tag{3.2}$$

  We say that the quenched field is trapping or non-trapping, according to whether $\tilde{\pi}(x) = 1$ for all x, or whether $\tilde{\pi}(x) < 1$ for some x.

It is worth noticing that here $\pi(x)$ and $\tilde{\pi}(x)$ indicate trapping probabilities, while in the previous section we used the letter π to indicate the transition matrix for a Markov chain. The first result obtained by them is the equivalence of the two problems:

Define $g(x,y) = \sum_{n=0}^{\infty} P_x(X_n = y)$ (Greens function). Then we have:

Theorem 1 – *Assume there is a constant K such that*

$$g(x,x) \leqslant K \text{ for all } x \in S. \tag{3.3}$$

Then the annealed field is trapping if and only if the quenched field is trapping.

As $\pi(x) \leqslant \tilde{\pi}(x)$, we are only interested in one direction of the theorem.

3.2 A criterion for trapping

A criterion for trapping is established for the annealed case, and under the assumptions of Theorem 1, the criterion also holds for the quenched case. Let $\tau(x) = \inf\{k : X_k = x\}$ denote the first hitting time of the

[3] Independent and identically distributed

point x. Define the quantities:

$$R_n = \sum_{i=0}^{n} -\log(1 - q(X_i))\mathbb{1}_{\tau(X_i)=i}; \quad (3.4)$$

$$\tilde{R}_n = \sum_{i=0}^{n} -\log(1 - q(X_i)).$$

Then the probability of no trapping in the annealed chain up to time n is given by:

$$P_x(X_i \notin T_i \text{ for all } i \leqslant n \| X_0, ..., X_n) = exp(-\tilde{R}_n)$$

Similarly, for the quenched problem we have:

$$P_x(X_i \notin T \text{ for all } i \leqslant n \| X_0, ..., X_n) = exp(R_n)$$

Thus, the trapping probabilities are:

$$\pi(x) = \mathbf{E}[1 - exp(-R_{infty})], \quad (3.5)$$

$$\tilde{\pi}(x) = \mathbf{E}[1 - exp(-\tilde{R}_{infty})],$$

We want to determine a trapping criterion for the annealed case, so we need to find $\tilde{R}_{infty}$ such that $\tilde{\pi}(x) = 1$. From the equations above, we can see that if $\tilde{R}_\infty = \infty$ we get:

$$\tilde{\pi}(x) = \mathbf{E}\big[1 - \frac{1}{e^\infty}\big] = \mathbf{E}[1] = 1,$$

So, our problem now is to determine when $\tilde{R}_\infty = \infty$ with probability 1. Define:

$$S_n = \sum_{i=0}^{n} q(X_i)\mathbb{1}_{\tau(X_i)=i}$$

$$\tilde{S}_n = \sum_{i=0}^{n} q(X_i).$$

The way we defined $\tilde{R}_n$, we see that $\tilde{R}_\infty = \infty$ if and only if $\tilde{S}_\infty = \infty$. Therefore, a necessary condition for trapping in the annealed case is that for all x_0,

$$\mathbf{E}\tilde{S}_\infty = \sum_{x \in S} g(x_0, x)q(x) = \infty \quad (3.6)$$

This is a necessary, but not sufficient condition for trapping.

Definition 1 A subset $A \subseteq S$ is transient for x if $P_x(X_n \in A \text{ for some } n) < 1$. A is transient if it is transient for some x. When all states communicate in A^c, then A must be transient for all x or for no x.

The example below depicts a situation when $\sum_{x \in S} g(x_0, x)q(x) = \infty$ but the chain is not trapping.

Example 6. Let A be any transient set satisfying $\sum_{x \in A} g(x_0, x) = \infty$. An example of such a set A for a simple random walk in $\mathbf{Z}^3$ is the set

$$\{x : \|x - (0, 0, 2^n)\| \leqslant 2^{n/2} \text{ for some } n \geqslant 1\}$$

Let $q(x) = c \in (0, 1)$ for $x \in A$ and 0 for $x \notin A$. Then clearly $\tilde{\pi}(x_0) < 1$ (so the chain is non-trapping), while $\sum_{x \in S} g(0, x)q(x) = \infty$.

Theorem 2 – *Suppose that the annealed field is non-trapping, i.e., for some $x_0, \tilde{\pi}(x_0) < 1$. Then there exists a subset $A \subseteq S$ such that:*

(i) *A is transient for x_0 and*

(ii) $\sum_{x \in S \setminus A} g(x_0, x)q(x) < \infty$

3.3 A necessary and sufficient condition for trapping

We now discuss under which conditions $\sum_{x \in S} g(x_0, x)q(x) = \infty$ is not only a necessary but also a sufficient condition for trapping. We do so by imposing some conditions on the geometry inherent in the Greens function. Assume here that the Greens function bound (3.3) holds.

Definition 2 – For $x_0 \in S, L \leqslant K$ and $\alpha \in (0, 1)$ the Greens function annulus $H_\alpha(L, x_0)$ is the set

$$\{x \in S : L \geqslant g(x_0, x) \geqslant \alpha L\}$$

We say the Markov chain has reasonable annuli [4]if for some $\alpha \in (0, 1)$ and for every $L \leqslant K, x_0 \in S$ and $A \subseteq$ transient with respect to x_0, the annulus $H_\alpha(L, x_0)$ has finite cardinality and:

$$\limsup_{L \to 0} \frac{|H_\alpha(L, x_0) \bigcap A|}{|H_\alpha(L, x_0)|} < 1 \tag{3.7}$$

Theorem 3 *Suppose the Markov chain on S with transitions $p(x, y)$ has reasonable annuli, and suppose that, for some $C, C\prime > 1$, the function $q : S \to [0, 1)$ satisfies the following regularity condition:*

$$\frac{1}{C} g(x_0, x) \leqslant g(x_0, y) \leqslant C g(x_0, x) \Rightarrow \frac{1}{C\prime} q(x) \leqslant q(y) \leqslant C\prime q(x) \tag{3.8}$$

Then $\tilde{\pi}(x) = 1$ if and only if $\sum_x g(x_0, x)q(x) = \infty$.

4 Simulation

We ran some simulations for the problems of absorbing Markov chains (using the example of the gambler's ruin) and also for quenched and annealed Markov chains in a field of traps, comparing both cases. All simulations were made using Matlab. It must be emphasized that the purpose of the simulations is neither to prove or confirm anything – since everything we are simulating has already been proven (in a very fashion way), we are only entertaining ourselves while illustrating the theory reviewed in the previous section.

[4] Annuli is a plural of annulus.

4.1 Gambler's ruins

As mentioned in example (number), two players each have a finite number of coins, say n_1 and n_2. A coin is flipped. If the result is TAILS, player 1 gives one of his coins to player 2. If the result is HEADS, player 2 gives one of his coins to player 1. The game stops when one of the players runs out of coins. The probability that a player will go bankrupt is equal to the ratio of the amount of coins his opponent have to the total amount of coins. Calling P_i the probability that player i will go bankrupt (which we shall refer as "failure"), we have:

$$P_1 = \frac{n_2}{n_1 + n_2}$$

$$P_2 = \frac{n_1}{n_1 + n_2}$$

We wrote a small code in Matlab to examine the probabilities above. We do this by running the code several times and comparing the frequencies F_{N1} and F_{N2} of failure for each player with the calculated probabilities. Matlab has special commands to generate pseudo-random numbers and arrays. Some of them worth mentioning are:

- rand – Basic syntax: rand(n) returns an n-by-n matrix containing pseudo-random values drawn from the standard uniform distribution on the open interval (0,1).
- randi – Basic syntax: r = randi(imax,n)returns an n-by-n matrix containing pseudo-random integer values drawn from the discrete uniform distribution on 1:imax.
- randn – Basic syntax: r = randn(n) returns an n-by-n matrix containing pseudo-random values drawn from the standard normal distribution.
- binornd – Basic syntax: binornd(N,P). This command generates random numbers from the binomial distribution with parameters specified by the number of trials, N, and probability of success for each trial, P. N and P can be vectors, matrices, or multidimensional arrays that have the same size, which is also the size of R. A scalar input for N or P is expanded to a constant array with the same dimensions as the other input.
- sprandsym – Basic syntax: sprandsym(S) returns a symmetric random matrix whose lower triangle and diagonal have the same structure as S. Its elements are normally distributed, with mean 0 and variance 1.
- sprand – Basic syntax: sprand(S)has the same sparsity structure as S, but uniformly distributed random entries.
- sprandn – Basic syntax: sprandn(S)has the same sparsity structure as S, but normally distributed random entries with mean 0 and variance 1.
- randsrc – Basic syntax: randsrc(m)generates an m-by-m matrix, each of whose entries independently takes the value -1 with probability 1/2, and 1 with probability 1/2. This command can also be used to generate a random matrix using prescribed alphabet. This is done by the elements of the alphabet, and the probabilities of occurrence of each element. Then, a $m \times n$ random matrix is generated with the command randsrc(m,n,[alphabet; prob]).

Each of these commands have many possibilities as to parameters input; ore information on the syntax of each command and some examples can be found at the MathWorks website (see references).

Below is the pseudo-code for our program. The inputs are the quantities of coins n_1 and n_2 each player has at the beginning of the game and p, the probability of a coin turning HEADS (the probability that player 1 wins a coin from player 2) and m, the desired number of matches. The output is: N_1, the number of failures for player

1 out of m matches; N_2, the number of failures for player 2; v, a vector of length m which stores information of the number of steps at each match until one of the players wins the game, and mv, which is just the mean of vector v. When mv is not an integer, we simply take the closest integer, since the purpose of this part of the code is just to give a glimpse at the magnitude of the number of steps until the end of the game as the initial conditions vary.

In the pseudo-code below, only the main features of the algorithm are presented, as the Matlab code has auxiliary variables used to store the initial values of n_1 and n_2 (which are the same from the first to the mth match), and so on. The results obtained are summarized below, where each line contains the information of $m = 100$ matches.

```
Input: n1,n2,p,m
Output: N1, N2, v, mv
while i < m + 1 do
    while (n1 ≠ 0) AND (n2 ≠ 0) do
        R = binomial (p);
        if R=1 then
            n1 := n1 + 1;
            n2 := n2 − 1;
            k := k + 1;
        end
        else
            n1 := n1 − 1;
            n2 := n2 + 1;
            k := k + 1;
        end
    end
    if n1 = 0 then
        %Player 1 has lost the game N1 := N1 + 1;
        $number of defeats for player 1
    end
    else
        %Player 2 has lost the game N2 := N2 + 1;
        $number of defeats for player 2
    end
    i := i + 1 ;
    v(m) := k ;
    m := m + 1 ;
end
mv := mean(v);
```

Algorithm 1: The gambler's ruin

The results obtained are summarized in the next tables, where each line contains the information of $m = 100$ matches.

Table 1: Game with $n_1 = n_2 = 10$, $p = 0.5$, $m = 100$. Calculated probabilities for failure are $P_1 = P_2 = 0.5$.

$n_1 = n_2 = 10, p = 0.5$				
N_1	F_{N1}	N_2	F_{N2}	mv
52	0.52	48	0.48	111
53	0.53	47	0.47	82
48	0.48	52	0.52	86
59	0.59	41	0.41	112
47	0.47	53	0.53	99
44	0.44	56	0.56	100
50	0.50	50	0.50	98
54	0.54	46	0.46	106
55	0.55	45	0.45	99

The simulation seems to indicate that the probabilities of failure are related to the calculated ratio between your opponents coins and the total amount of coins, and the number of steps until the game reaches and end is reasonable – around one hundred steps. Let's now examine a situation when one of the players has far more coins then the other, aside with a better probability at his side (similar to casinos, which always have more money than individual players and usually also have better chances of winning).

Table 2: Game with $n_1 = 10000, n_2 = 10$, $p = 0.7$, $m = 100$. Calculated probabilities for failure are $P_1 = 0.01$ and $P_2 = 0.99$.

$n_1 = 10000, n_2 = 10, p = 0.7$				
N_1	F_{N1}	N_2	F_{N2}	mv
0	0	100	1	26
0	0	100	1	24
0	0	100	1	27
0	0	100	1	26
0	0	100	1	24
0	0	100	1	27
0	0	100	1	27
0	0	100	1	24
0	0	100	1	24

Repetitive as it seem, the lines of this table show us what we already knew – the player starting with less coins has greater chance of losing. But still, one could argue that even at casinos there are games in which the odds are not so good for the house, and thus players have a chance of (eventually) winning. So, let's examine a situation in which one of the players has more coins, but the other has better odds:

Table 3: Game with $n_1 = 10000, n_2 = 10$, $p = 0.3$ (so the player with less coins has some advantage), $m = 100$. Calculated probabilities for failure are $P_1 = 0.01$ and $P_2 = 0.99$.

$n_1 = 100000, n_2 = 10, p = 0.3$				
N_1	F_{N1}	N_2	F_{N2}	mv
0	0	100	1	2510
0	0	100	1	2507
0	0	100	1	2477
0	0	100	1	2497
0	0	100	1	2518
0	0	100	1	2499
0	0	100	1	2484
0	0	100	1	2512
0	0	100	1	2513

Analyzing table 3, once again the result is obvious: since the probabilities of failure depend on the amount of coins you start, even when the poorer player has better odds, it will lose. Plating the game with these initial conditions is just like running my Matlab program for this case: it takes longer, but nobody doubts what the result is going to be. This example was intended to show two things: first, an absorbing Markov chain reaches an absorption state after a given number of steps – which may not be a really big number, after all – with probability 1. as Table 1 indicates. Second, when considered the long-term profits among several games, casinos *always* win. That's a mathematical fact. [5]

4.2 Drunkard walk – one dimensional random walk

We now turn to the comparison of trapping probabilities for annealed and quenched fields, bearing in mind the result obtained by Pemantle and Volkov: *assuming that the bound on the Greens function holds, than the annealed field is trapping if and only if the quenched field is trapping.* Suppose a drunkard is currently at position $d = 0$. At position $d = -b$ and $d = h$ (where b and h are integers) are the bar and his house, respectively. He takes discrete steps with equal probabilities to left or right, and once he reaches the bar or his house he stays there with fixed probabilities p_b and b_h. The bar and the house represent traps with fixed probabilities, hence this is a quenched problem. We want to investigate whether this is a trapping field. A Matlab program was written to simulate this situation, and the pseudo-code is given below. The input data is the positions of the bar (a negative integer) and the house (a positive integer), and their respective probabilities of trapping. The output is the final position d (either the bar or his house) and the number k of steps taken.

[5] Besides that, casinos make even more money by altering their payouts from true odds to casino odds. The following excerpt tells us the difference between true and casino odds: " Roulette offers a prime example of the difference between true odds and casino odds. On an American roulette wheel, there are 38 pockets. The odds of any given number coming up on any given spin are 37 to 1. But when you bet on a given number, the casino only pays out 35 to 1 if you win. Say you bet $1 on each number, for a total wager of $38. For the winning number, you'll be paid $35, plus you'll get your $1 back on that particular number, for a total of $36. So, even when you bet on every available number, you lose $2. That's the house edge: 2/38, or 5.26 percent the difference between true odds and casino odds."

Input: b, h, p_b, p_h
Output: d, k
while $i < m + 1$ **do**
  step = randsrc(1,1); %the result is a step left or right with prob. 1/2;
  k := k+1; ;
  %number of steps position = position + step; % and he walks!;
  if $d = b$ **then**
    c = binornd (1, p_b);
    if *c=0* **then**
      d := b+1 % a step away from the bar;
      k := k+1 ;
    end
  end
  if $d = h$ **then**
    c = binornd (1, p_h);
    if *c=0* **then**
      d := h-1 % a step away from the bar;
      k := k+1 ;
    end
  end
end

4.2.1 Quenched field

As we have proceeded for the gambler's ruin, we run the algorithm above with different parameters. Once again, we have ran the program 100 times and calculated the mean number of steps until reaching an end (each line is the result of this set of a hundred walks). We have also included a stop criterion based on the number of steps at the beginning of the program, just in case the walk never reaches one of the traps. The results are summarized over the next tables:

Table 4: Drunkard walk with $b = -10, p_b = 0.5, h = 10, p_h = 0.5$ – symmetric random walk

Quenched problem: $b = -10, p_b = 0.5, h = 10, p_h = 0.5$		
d=bar	d=house	mean number of steps
48	52	117
42	58	126
52	48	112
60	40	115
45	55	111
48	52	116
48	52	117
40	60	118
47	53	121

It seems this field is trapping. But we could try and make the walk a little more interesting by making it asymmetric – the probabilities of taking a step left or a step right are no longer equal. This can be done with some small changes in our code. Let's also put the bar and the house at considerably different distances, that is, either $|b| >> |h|$ or $|h| >> |b|$. Will the field still be trapping? In the next simulation, the probability of taking a step left is 0.2 (and consequently, the probability of taking a step right is 0.8).

Table 5: Drunkard walk with $b = -10, p_b = 0.5, h = 15000, p_h = 0.5$ – asymmetric random walk

Quenched problem: $b = -10, p_b = 0.5, h = 15000, p_h = 0.5$		
d=bar	d=house	mean number of steps
0	100	5007
0	100	5012
0	100	5002
0	100	5023
0	100	5007
0	100	5010
0	100	5014
0	100	5009
0	100	5018

Just for a comparison, we ran the program for the symmetric random walk with the values $b = -10, p_b = 0.5, h = 15000, p_h = 0.5$ used above, and the mean number of steps until reaching a trap is not considerably different from the above. The mean number of steps until reaching an absorbing state is 6408 steps, with the difference that the preferential absorbing state is now the bar (because it is closer, and the probabilities of taking a step left or right are equal). In fact, it can be seen in the tables above that the probabilities of getting trapped in trap i are close to the ratio between the distance from the origin to trap j and the distance from one trap to the other (which is quite intuitive, but still interesting to observe!). This is actually very similar to the problem of the gambler's ruin, which is natural since the gambler's ruin problem is itself a random walk.

The simulations above suggest that the field in this problem is trapping (the field here is the set of integers $S = [b, h]$ with traps at positions $x = b$ and $x = h$, with fixed probabilities).

4.2.2 Annealed field

We consider now the annealed problem (where the probabilities that the bar and the house are traps aren't fixed). Once the drunkard reaches the bar or the house, the probability of trapping is a random number, chosen from a uniform distribution on the interval(0,1). The annealed field is trapping if and only if the quenched field is trapping, we expect to observe trapping for all cases considered. Over the next table, each line contains the results of a hundred walks. We do this for both the symmetric and the asymmetric case. The results are:

Table 6: Drunkard walk with $b = -10, h = 10, p_h$ and p_b not fixed – symmetric random walk

Annealed problem: $b = -10, h = 10$		
d=bar	d=house	mean number of steps
59	41	114
44	59	125
51	49	136
56	44	120
54	46	109
44	56	119
35	65	118
56	44	129
47	53	117

Now we consider the asymmetric random walk, with the probability of taking a step left as 0.2. We have also put the bar and the house at considerably different distances.

Table 7: Drunkard walk with $b = -10, h = 15000p_h$ and p_b not fixed – asymmetric random walk

Annealed problem: $b = -10, h = 10$		
d=bar	d=house	mean number of steps
0	100	4991
0	100	5006
0	100	5013
0	100	5008
0	100	4987
0	100	4995
0	100	5017
0	100	5010
0	100	4990

From the tables above, we can see the field is annealed, as we expected. In the previous section, we stated in which conditions $\sum_{x \in S} g(x_0, x)q(x) = \infty$ is a necessary and sufficient condition for trapping.

We are taking only homogeneous Markov chains (the probabilities of going from one state to another don't depend on the number of the step). Let's examine the probabilities $P(X_n) = x$ by looking at the transition probability matrix. Let's refer to our example of the drunkard: he never stays at a given position for more than one unit of time – he must always take a step left or right. The only exception is when he reaches one of the trapping (absorbing) states – namely, the bar or his house, for which x=-b or x=h – where he stays with fixed probabilities (quenched problem) or varying probabilities (annealed problem). Anyway, the probability of trapping is a function $q(x)$ such that $0 \leqslant 1q(x) \leqslant$; besides that, he walks only from one state to its immediate neighbor – he can't jump from the origin to state 3, for example. Setting the origin as zero, the bar at x=-b and the house at x=h, the transition probability matrix has the special property that each row has at most two non-null elements. Just to give an idea of how this matrix looks like, below is an example in which the x = -3 is an absorbing state with probability 0.4; x = 2 is an absorbing state with probability 0.4 and each step has equal probabilities for both right and left sides.

$$\pi = \begin{array}{c} \\ -3 \\ -2 \\ -1 \\ 0 \\ 1 \\ 2 \end{array} \begin{array}{c} \begin{array}{cccccc} -3 & -2 & -1 & 0 & 1 & 2 \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0.4 & 0 & 0.5 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0.5 & 0 & 0.4 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

We computed some powers of this matrix with Matlab to see how the values evolve. The nth power of this matrix, for large n, is:

$$\pi = \begin{array}{c} \\ -3 \\ -2 \\ -1 \\ 0 \\ 1 \\ 2 \end{array} \begin{array}{c} \begin{array}{cccccc} -3 & -2 & -1 & 0 & 1 & 2 \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0.64 & 0 & 0.5 & 0 & 0 & 0.16 \\ 0.48 & 0.5 & 0 & 0.5 & 0 & 0.32 \\ 0.32 & 0 & 0.5 & 0 & 0.5 & 0.48 \\ 0.16 & 0 & 0 & 0.5 & 0 & 0.64 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

We can see from the matrix above that all higher powers have the same values for each and every element, so, now matter where we start, clearly $\sum_{x \in S} g(x_0, x)q(x) = \infty$. For our problem the results would be similar, the only difference being the size of the matrix. As most Markov chains, a simple random walk has reasonable annuli (see [7]), and so our condition is sufficient and enough for trapping.

5 Conclusions

The simulations all showed the same results that were proved theoretically. Also, while writing the code used to run the simulations we confirmed the simplicity and power of modeling phenomena as Markov chains: the algorithms are very simple and can easily be changed to study more complex phenomena (provided that this is still a Markov chain, of course), as all it takes to run these programs is essentially a random number generator.

6 References

[1] Drie, John H. Van. "What Laplace *really* said - the myth of Laplacian determinism"

[2] Gnedenko, B. V. *The Theory of Probability*. Mir Publishers, Moscow, 1969.

[3] Grimmett, G. , David Stirzaker *Probability and Random Processes*. Oxford University Press, New York, 3rd edition 2001.

[4] Hannun, Robert C. "A Guide to Casino Mathematics".

[5] MATLAB version 7.0.0. The MathWorks Inc, Natick, Massachusetts, 2004.

[6] MathWorks Product Documentation

[7] Pemantle, R. , Stanislav Volkov,*"Markov Chains in a Field of Traps"*.Journal of Theoretical Probability, 1 April 1998, 561-569.

[8] Shiryaev, A. N. *Probability*. Springer 2nd. edition 1996.

[9] Weisstein, Eric W. "Gambler's Ruin." From MathWorld–A Wolfram Web Resource.

[10] Schneider, Meg E. 'Playing Against the House".