POLYA'S URN AND THE MARTINGALE CONVERGENCE THEOREM

SHIRONG LIU

ABSTRACT. This paper is about Polya's Urn and the Martingale Convergence Theorem. I will start with the formal definition, followed by a simple example of martingale and the basic properties of martingale. Then I will explain the Polya's Urn model and how it contributes to proving the Martingale Convergence Theorem. Finally, I will give a full proof of the Martingale Convergence Theorem.

CONTENTS

1. DEFINITION OF MARTINGALE

Definition 1.1. Suppose there exist three series of random variables $X_1, X_2, \cdots, X_n$, $B_1, B_2, \cdots, B_n$ and $M_1, M_2, \cdots, M_n$, that $M_i = X_i \cdot B_i$ for each $1 < i < n$. Then, $\{M_n\}$ is a martingale with respect to $\{X_n\}$ if, for each n, M_n is finite, and B_n depends only on the outcomes of $\{X_1, X_2, \cdots, X_{n-1}\}$.

Furthermore, $E[M_n \mid F_{n-1}] = M_{n-1}$, where F_{n-1} denotes the information in $X_1, X_2, \cdots, X_{n-1}$ (When $n = 1$, $F_{n-1} = F_0$ is empty since there is no previous game before the first game).

Before I go into more depth on the definition of martingale, let's first make the notations clear:

Notation 1.2. In this paper, I use the symbol $E[A \mid B]$ to mean "the conditional expected value of A given B" and the symbol $\mathbf{E}(A)$ to mean "the expected value of A".

To make things even more straightforward, the definition of expected value is:

Date: August 9th, 2016.

Definition 1.3. Let an event A be given, and let $R_1, R_2, \cdots, R_n$ be the n possible outcomes of event A. Define $P_1, P_2, \cdots, P_n$ to be the probability of $R_1, R_2, \cdots, R_n$ happening, respectively. Then, the expected value of A is

$$\mathbf{E}(A) = R_1 \cdot P_1 + R_2 \cdot P_2 + \cdots + R_n \cdot P_n$$

Then, let's get back to where we were. According to the projection rule [1] of conditional expectation,

$$\begin{aligned} E[M_{n+2} \mid F_n] &= E[E[M_{n+2} \mid F_{n+1}] \mid F_n] \\ &= E[M_{n+1} \mid F_n] \\ &= M_n. \end{aligned}$$

More generally, for any $m > n$,

$$E[M_m \mid F_n] = M_n. \tag{1.4}$$

Since M_n is a random variable, we can calculate the expected value of M_n:

$$\mathbf{E}(M_n) = \mathbf{E}(E[M_n \mid F_0]) = \mathbf{E}(M_0) = 0. \tag{1.5}$$

This is the basic property of martingale.

2. Simple Example of a Martingale

A martingale is the mathematical version of a fair game. The simplest example of a martingale, therefore, would be the fair game of flipping a coin. The bettor flips a coin, with probability one half of getting heads and one half of getting tails. Before each flip, the bettor is allowed to place a bet on the outcome of this flip. Let's call the series of events (flips) $X_1, X_2, X_3, \cdots, X_n$ and the series of bets $B_1, B_2, B_3, \cdots, B_n$. Let heads be 1 and tails be -1 (of course, they can be interchanged). Then the total money that the bettor has after n flips is

$$\begin{aligned} M_n &= X_1 \cdot B_1 + X_2 \cdot B_2 + \cdots + X_n \cdot B_n \\ &= \sum_{i=1}^{n} X_i \cdot B_i. \end{aligned}$$

The bettor is allowed to place his bet B_n for X_n knowing the outcomes of the previous games $X_1, \cdots, X_{n-1}$. Thus, B_n is a best guess based on the information collected in $X_1, \cdots, X_{n-1}$. Denote by F_{n-1} the outcomes of $X_1, \cdots, X_{n-1}$. We know that F_0 is empty, since there is no game before the first flip, so

$$E[X_1 \cdot B_1 \mid F_0] = \mathbf{E}(X_1 \cdot B_1) = (\frac{1}{2})(1)(B_1) + (\frac{1}{2})(-1)(B_1) = 0. \tag{2.1}$$

Given series $M_1, M_2, \cdots, M_n$, $M_i = X_i \cdot B_i$ for each $1 \geq i \leq n$. Then,

$$\begin{aligned} E[M_n \mid F_{n-1}] &= E[(M_{n-1} + X_n \cdot B_n) \mid F_{n-1}] \\ &= E[M_{n-1} \mid F_{n-1}] + E[X_n \cdot B_n \mid F_{n-1}] \\ &= M_{n-1} + 0 \\ &= M_{n-1}. \end{aligned}$$

[1] See Greg Lawler's *Random Walk and the Heat Equation* page 115

Then,

$$E[M_n \mid F_{n-1}] = M_{n-1}$$
$$E[M_{n-1} \mid F_{n-2}] = M_{n-2}$$
$$\vdots$$
$$E[M_1 \mid F_0] = \mathbf{E}(M_1) = 0.$$

Thus, $E[M_n \mid F_{n-1}] = 0$. Therefore, the example of flipping coins satisfies the property of being a martingale.

3. Polya's Urn

After understanding what martingale is and its basic property, let's look at another statistical model called *Polya's Urn.* Suppose there is an urn that contains red and green balls that are different only by color. At the beginning of the game, the urn only contains 1 red ball and 1 green ball. At each discrete time (trial) n, the player takes out a ball randomly from the urn, and returns the ball along with a new ball of the same color to the urn. Let $X_1, X_2, X_3, \cdots, X_n$ denote the outcome of the color of the ball taken out at each trial. Specifically, let X_i be 1 if the color is red and 0 if the color is green for each $1 \leq i \leq n$. Since the player knows the outcome of the previous $n-1$ trials before he plays the nth trial, he can calculate the probability and the expected value of the outcome of his nth trial. Now the question is: given that the player has picked out k red balls in his first n trials, what is the probability that he is going to pick out the $(k+1)$th red ball?

We know that the original two balls, 1 red and 1 green, will never leave the urn, and the balls added to the urn later in the game is a total of n balls. So the total number of balls in the urn after n trials is $n+2$. Since for each i, X_i equals 1 if the chosen ball is red and 0 if the ball is green, the number of red balls chosen after n trials equals the sum of $X_1, X_2, \cdots, X_n$. Let $Y_n = \sum_{i=1}^{n} X_i$, then the number of red balls after n trials is $Y_n + 1$. Now we can answer the question.

$$\mathbf{P}(Y_{n+1} = k+1 | X_1, X_2, ..., X_n) = \frac{Y_n + 1}{n+2} = \frac{k+1}{n+2}$$

Then, the expected value of Y_{n+1} is

$$\begin{aligned} E[Y_{n+1} \mid Y_n = k] &= \frac{k+1}{n+2}(k+1) + (1 - \frac{k+1}{n+2})k \\ &= \frac{k^2+2k+1}{n+2} + \frac{nk - k^2 + k}{n+2} \\ &= \frac{nk + 2k + k + 1}{n+2} \\ &= \frac{k(n+2)}{n+2} + \frac{k+1}{n+2} \\ &= k + \frac{k+1}{n+2}. \end{aligned}$$

Next, we'll show that Polya's Urn is a martingale.

Let's denote a new variable M_n to be the fraction of red balls in the urn after n trials, that is, $M_n = \frac{Y_n+1}{n+2}$. Then, the expected value of M_{n+1} given the information in the previous n trials is

$$\begin{aligned} E[M_{n+1} \mid Y_n = k] &= E[\frac{Y_{n+1}}{n+3} \mid Y_n = k] \\ &= \frac{k + \frac{k+1}{n+2} + 1}{n+3} \\ &= \frac{k(n+2) + k + 1 + (n+2)}{(n+2)(n+3)} \\ &= \frac{nk + 3k + n + 3}{(n+2)n+3} \\ &= \frac{(k+1)(n+3)}{(n+2)(n+3)} \\ &= \frac{k+1}{n+2} = M_n. \end{aligned}$$

Therefore, Polya's Urn satisfies the property of a martingale.

4. Using Polya's Urn to Prove the Martingale Convergence Theorem

In this section, we will use Polya's Urn model to prove a lemma, which we will use later in the proof of the Martingale Convergence Theorem.

Let's consider a more general case of Polya's Urn. Let's assume there are J red balls and K green balls in the urn at the beginning. Let $b = \frac{J}{J+K}$ be the fraction of red balls in the urn. Then, we define a series of random variables $\{M_n\}$ to be the fraction of red balls in the urn at each discrete time n, so $M_0 = b$. From here on, we will examine two special cases of the movements of this martingale.

Case 1: Let $a < b$, and let T be the smallest n such that $M_n \leq a$, that is, T is the first time that M_T becomes less than or equal to a. Let T be a stopping time; in other words, we are only going to observe the movement of M_n when n is between 1 and T. One thing to note is that it is possible for T to be ∞, that is to say that $\{M_n\}$ is a nondecreasing sequence, since that is the only way to guarantee that M_n will never become less than or equal to a. Next, we will bound the probability that T will be finite.

There are basically two options for T:

1). $T > n$ for all n, which means that T is infinite, or

2). $T \leq n$, which means that T is finite.

Let $T \wedge n$ denotes minimum$\{T, n\}$. Using the optional sampling theorem [2], we can get

$$\begin{aligned} b &= \mathbf{E}(M_0) = \mathbf{E}(M_{T\wedge n}) \\ &= \mathbf{P}(T \leq n)E[M_T \mid T \leq n] + (1 - \mathbf{P}(T \leq n))E[M_T \mid T > n]. \end{aligned}$$

Therefore,

$$\mathbf{P}(T \leq n) = \frac{E[M_T \mid T > n] - b}{E[M_T \mid T > n] - E[M_T \mid T \leq n]}.$$

Now, let's bound $E[M_T \mid T > n]$ and $E[M_T \mid T \leq n]$.

If $T > n$, then as explained before, $\{M_n\}$ has to be a nondecreasing sequence because it has to guarantee that M_i is never less than or equal to a for any $0 < i \leq \infty$. Since $M_0 = b$, and the sequence $\{M_n\}$ is nondecreasing, therefore M_n can

[2]See Greg Lawler's *Random Walk and the Heat Equation*, page 117

never be less than b. Also, M_n can never exceed 1 since it represents the fraction of red balls in the urn. Therefore, it's clear that

$$b \leq E[M_T \mid T > n] \leq 1.$$

Next, if $T \leq n$, then for some finite T, M_T becomes less than or equal to a at time T. Then we can bound $E[M_T \mid T \leq n]$ as

$$0 \leq E[M_T \mid T \leq n] \leq a,$$

since the game stops at time T when $M_T \leq a$.

Therefore,

$$\begin{aligned}\mathbf{P}(T \leq n) &= \frac{E[M_T \mid T > n] - b}{E[M_T \mid T > n] - E[M_T \mid T \leq n]} \\ &\leq \frac{1-b}{1-a}.\end{aligned}$$

Therefore, the probability that given $M_0 = b$, M_n will ever get as small as a is $\frac{1-b}{1-a}$.

Case 2: Let's assume that for some fixed value j, $M_j = a$, then what's the probability that M_n will get as large as b, for some $n > j$?

To answer this question, let S denote the time when M_n becomes greater than or equal to b. Then,

$$\begin{aligned}a &= \mathbf{E}(M_j) = \mathbf{E}(M_{S\wedge n}) \\ &= \mathbf{P}(S \leq n) \cdot E[M_S \mid S \leq n] + (1 - \mathbf{P}(S \leq n)) \cdot E[M_S \mid S > n].\end{aligned}$$

Then,

$$E[M_S \mid S \leq n] = \frac{E[M_S \mid S > n] - a}{E[M_S \mid S > n] - E[M_S \mid S \leq n]}.$$

Again, let's bound $E[M_S \mid S > n]$ and $E[M_S \mid S \leq n]$:

$$0 \leq E[M_S \mid S > n] \leq a$$
$$b \leq E[M_S \mid S \leq n] \leq 1$$

Therefore,

$$E[M_S \mid S \leq n] \leq \frac{a - E[M_S \mid S > n]}{b - E[M_S \mid S > n]} \leq \frac{a}{b}.$$

Consider a type of special event, that is, when M_n becomes less than or equal to a at some time T and then increases to be greater than or equal to b again for some time after T. This event is called an $(a,\, b)$ fluctuation, or, more formally, *a downcrossing of a followed by an upcrossing of b*. Since the events involve only picking balls from the urn, and therefore each $\{M_i\}$ variable is independent, we know that the probability of an $(a,\, b)$ fluctuation happening is bounded by the product of the probability of Case 1 and the probability of Case 2:

$$\mathbf{P}(\text{(a, b) fluctuation}) \leq \left(\frac{1-b}{1-a}\right)\frac{a}{b}.$$

So the probability of k $(a,\, b)$ fluctuation happening is bounded by

$$\left(\left(\frac{1-b}{1-a}\right)\frac{a}{b}\right)^k. \tag{4.1}$$

Then, as $k \to \infty$, Equation (4.1) goes to 0. So the probability of infinitely many $(a,\, b)$ fluctuations is 0.

$$\mathbf{P}(\text{infinitely many (a, b) fluctuations}) = \lim_{k\to\infty} \left(\left(\frac{1-b}{1-a} \right) \frac{a}{b} \right)^k = 0 \tag{4.2}$$

Since the probability of infinitely many $(a,\, b)$ fluctuations is 0, it is of probability 1 that there are finitely many $(a,\, b)$ fluctuations. This leads to the conclusion of Lemma 4.3.

Lemma 4.3. *Given a sequence of real numbers $X_1, X_2, X_3, \cdots$ such that for every rational interval $(a,\, b)$, $\{X_n\}$ makes only finitely many upcrossings of $(a,\, b)$. Then, there exist $C \in [-\infty, \infty]$ such that*

$$\lim_{n\to\infty} X_n = C.$$

If the sequence is bounded, $C \in (-\infty, \infty)$.

Proof. Proof by contradiction:

Suppose that the sequence $\{X_n\}$ does not converge to finite nor infinite limit $\in [-\infty, \infty]$, then

$$\lim_{n\to\infty} \inf[X_n] < \lim_{n\to\infty} \sup[X_n],$$

by the properties of infimum and supremum.

Then, we can find rational numbers a and b, $a < b$, such that

$$\lim_{n\to\infty} \inf[X_n] < a < b < \lim_{n\to\infty} \sup[X_n].$$

Since $\{X_n\}$ does not converge to finite nor infinite limit, the sequence has to be less than a infinitely many times and greater than b infinitely many times. So the sequence $\{X_n\}$ makes infinitely many upcrossings of rational interval $(a,\, b)$. This contradicts the assumption of this lemma.

Therefore, $\exists\ C \in [-\infty, \infty]$ such that $\lim_{n\to\infty} X_n = C$. □

Applying Lemma 4.3 to Polya's Urn model, we can establish the following:

With probability 1, the sequence $\{M_n\}$ converges to a M_∞, for some random variable $M_\infty \in [0, 1]$.

$$\lim_{n\to\infty} M_n = M_\infty$$

Next, I will use Lemma 4.3, which we concluded from Polya's Urn model, to give a proof of the Martingale Convergence Theorem.

5. Martingale Convergence Theorem

Theorem 5.1. *(Martingale Convergence Theorem). Let a martingale $M_1, M_2, M_3, \cdots$ with respect to independent random variables $X_1, X_2, X_3, \cdots$ be given. Then if there exists $C < \infty$ such that for all n,*

$$\mathbf{E}(|M_n|) \leq C,$$

there exists a random variable M_∞ such that, with probability 1,

$$\lim_{n\to\infty} M_n = M_\infty.$$

Proof. To make this problem simpler, we can assume that $M_0 = 0$. This is because if $M_0 \neq 0$, we can consider the sequence $M_0 - M_0, M_1 - M_0, M_2 - M_0, \cdots$ as the martingale under examination. We proved in the previous section that, for any rational interval (a, b), the martingale makes only finitely many upcrossings of (a, b). We will show that the probability of the sequence $\{M_n\}$ in this case making infinitely many (a, b) fluctuations is zero.

To prove this, let us write the sequence $\{M_n\}$ differently:

$$M_n = \delta_1 + \delta_2 + \delta_3 + \cdots + \delta_n,$$

where $\delta_i = M_i - M_{i-1}$ for each $1 \leq i \leq n$.

Then we define a sequence of variables $\{W_n\}$ such that

$$W_n = \sum_{i=1}^{n} B_i \cdot \delta_i.$$

Here, imagine that each B_i is an investment or bet we place on the outcome of the event δ_i. This investment decision is made as a best guess from knowing the outcomes of the previous $(i-1)$ games, but not knowing the outcome of δ_i. So if we denote the information in the outcomes of the previous $(i-1)$ games as F_{i-1}, then the expectation of each W_i can be written as:

$$\begin{aligned} E[W_i \mid F_{i-1}] &= E[W_{i-1} + B_i\delta_i \mid F_{i-1}] \\ &= E[W_{i-1} \mid F_{i-1}] + E[B_i\delta_i \mid F_{i-1}] \\ &= W_{i-1} + 0 \\ &= W_{i-1}. \end{aligned}$$

More specifically, let's see why $E[B_i\delta_i \mid F_{i-1}] = 0$.

To calculate for the conditional expectation of $E[B_i\delta_i \mid F_{i-1}]$, let's list the possible outcomes of δ_i:

a). δ_i is possible: $\delta_i = |\delta_i|$.

b). δ_i is negative: $\delta_i = -|\delta_i|$.

Each of these possibilities has a probability of 0.5 of happening. Therefore,

$$\begin{aligned} E[B_i\delta_i \mid F_{i-1}] &= (\frac{1}{2})(B_i|\delta_i|) + (\frac{1}{2})(-B_i|\delta_i|) \\ &= 0. \end{aligned}$$

Since this is true for all i, in general

$$\mathbf{E}(W_n) = \mathbf{E}(W_0).$$

Then, all we need to do is strategically choose the B_i. Remember we want to prove that M_n can only make finitely many (a, b) fluctuations.

Without loss of generality, let's assume $a \geq 0$ (The case $a < 0$ is done similar to this one). Then we let $B_i = 0$ for all $i < T$, where T is the smallest n such that $M_{n-1} \geq a$. When $i = T$, we let $B_i = 1$. We keep $B_i = 1$ until the first time S, $S > T$, such that $M_S \geq b$. When $i = S$, we change B_i back to 0 and we keep it that way until M_i drops back below a again, for some $i > S$. This way, if W_n is the total amount of money in some financial investment, we gain at least $(b-a)$ for every (a, b) fluctuation in $\{W_n\}$. We denote J_n as the number of (a, b) fluctuations by the nth game, then we have

$$W_n \geq J_n \cdot (b-a) + (M_n - a) \cdot \mathbb{1}_{M_n \leq a} \geq J_n \cdot (b-a) - |M_n|, \tag{5.2}$$

where $\mathbb{1}_{M_n \le a}$ is the indicator function of the set $(M_n \le a)$. In addition, $(M_n - a) \cdot \mathbb{1}_{M_n \le a}$ is a nonpositive term that accounts for the amount we lost when we set the bet to be 1 at the last drop below a before we end this process (right before time n).

Since $\mathbf{E}(W_n) = \mathbf{E}(W_0) = 0$, when we take the expectation of both sides, we get:

$$\begin{aligned} 0 &\ge \mathbf{E}(J_n \cdot (b-a) - |M_n|) \\ &\ge \mathbf{E}(J_n) \cdot (b-a) - \mathbf{E}(M_n) \\ \mathbf{E}(J_n) &\le \frac{\mathbf{E}(M_n)}{b-a} \\ \lim_{n\to\infty} \mathbf{E}(J_n) &\le \lim_{n\to\infty} \frac{\mathbf{E}(M_n)}{b-a} \\ \frac{\mathbf{E}(M_n)}{b-a} &\le \frac{C}{b-a}. \end{aligned}$$

As $n \to \infty$, $J_n \to J_\infty$,

$$\mathbf{E}(J_\infty) \le \frac{C}{b-a} \le \infty,$$

since both C and $(b-a)$ are finite. Therefore, with probability one, there are finitely many (a, b) fluctuations made by $\{M_n\}$.

Then, by Lemma 4.3, we can see that with probability one, there exists a limit M_∞ such that

$$\lim_{n\to\infty} M_n = M_\infty$$

□

Acknowledgments. It is a pleasure to thank my mentor, Redmond McNamara, for his help and guidance in studying Random Walk and Martingale. I would also like to thank my program director, Peter May, for his hard work in arranging and running the Math REU Program.

References

[1] Gregory F. Lawler. Random Walk and the Heat Equation. Chapter 3. Martingales