Tossing Coins Until Mostly Tails

18.200 Discrete Mathematics Term Paper

Agustin Garcia

May 15, 2018

Abstract

In this paper, we consider the process of tossing a coin until more tails have been flipped than heads. This process is closely related to other combinatorial ideas, including Dyck paths and the Catalan numbers, random walks, and generating functions. These relations are explored, mostly in the scenario involving a fair coin. Most importantly, we show that in such a scenario, the process described is guaranteed to end. Then, variations with unfair coins and different stopping criteria are considered.

Contents

1 Introduction

Suppose we flip a coin and keep track of the number of heads and tails, stopping when the number of tails exceeds the number of heads. Clearly, there is a one-half chance that the process stops after just a single flip, but if the first flip is heads, what happens next? What is the probability that the process stops after a fixed number of coin tosses, and is the process guaranteed to stop? Furthermore, do these probabilities change if the coin tossed is not fair? In this paper, we consider these questions and relate them to random walks, Dyck paths, and generating functions. Through these relationships, we find the probabilities mentioned above and prove interesting properties of symmetric random walks.

In Section 2.1 of this paper, we restate the fair coin problem using symmetric random walks, which allows us to find the probability that the process ends after $2n+1$ coin tosses for a given n. Later, in Section 2.2, we consider the asymptotics of these probabilities. We find that, as expected, the probability that this process ends converges. However, the expected number of coin tosses it takes does not converge.

Then, in Section 3, we develop generating functions for both Dyck paths and the coin toss problem. This allows us to conclude an important result: that this process is guaranteed to end after some number of coin tosses.

In Section 4, we begin by building on the result of Section 3 to demonstrate that symmetric random walks on the integers eventually reach every integer. We then consider variations of the original problem. We find that if the coin is not fair, the process is *not* guaranteed to end. Furthermore, we consider different stopping criteria for the coin tossing problem, such as stopping after 2 more tails than heads.

2 Probability that the process ends after a given number of steps

We wish to find a closed-form formula for the probability that we first get more tails than heads on the $2n+1^{th}$ toss, and then consider the asymptotics of this probability. To help simplify the problem, we first consider the problem in the context of random walks.

Definition. We define a *random walk on* $\mathbb{Z}$ to be a path starting at 0, and moving 1 in the positive direction with some probability p, and 1 in the negative direction with probability $1-p$ at each step. A random walk is *symmetric* if $p = 0.5$.

2.1 Coin tosses and random walks

In this section, we begin by noticing a relationship between the random coin tossing process and random walks on the integers. This will allow us to discuss the two interchangeably throughout the paper. Assign a head to be $+1$ and a tail to be -1, and construct a random walk using coin tosses. In this section we consider tossing a fair coin, so the random walk is symmetric. It is helpful to visualize the random walk as follows, where the horizontal axis represents the step (or time), and the vertical axis represents the position on $\mathbb{Z}$ at that step. Figure 1 illustrates one such visualization.

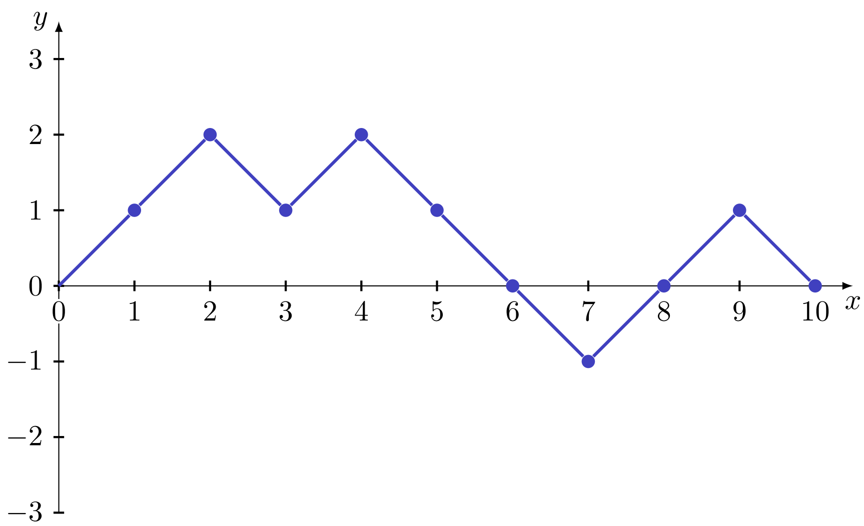

Figure 1: the walk corresponding to HHTHTTTHHT.

Now we consider the problem at hand using the language of random walks. The problem is equivalent to taking a random walk and stopping as soon as the position is -1. Importantly, this means that before the last step, the position in the random walk was always non-negative.

Definition. We call such a walk of length $2n$ that always stays non-negative and ends on 0 a *Dyck path*. Dyck paths are enumerated by the *Catalan numbers*,

$$C_n = \frac{1}{n+1} \cdot \binom{2n}{n}.$$

We also note that the process always ends after an odd number of steps, since every step changes the parity of the position and we wish to end at an odd position. From here we may formulate an expression for the probability that the process ends after $2n+1$ coin tosses, for any n.

Theorem 2.1. *The probability that this process ends after $2n+1$ coin tosses is*

$$2^{-(2n+1)} \cdot \frac{1}{n+1} \cdot \binom{2n}{n}.$$

Proof. We use our interpretation of the problem using symmetric random walks. We wish to find the number of paths of length $2n+1$ that stay non-negative until the last step. This is equivalent to finding the number of Dyck paths of length $2n$, which is

$$C_n = \frac{1}{n+1} \cdot \binom{2n}{n}.$$

The total number of walks of length $2n+1$ is 2^{2n+1}, and each walk happens with equal probability. Therefore we have that the probability is

$$\mathbb{P}[\text{stop after } 2n+1 \text{ tosses}] = \boxed{2^{-(2n+1)} \cdot \frac{1}{n+1} \cdot \binom{2n}{n}}.$$

Note: From now on we refer to the probability above simply as $\mathbb{P}(n)$. $\square$

2.2 Asymptotics of $\mathbb{P}(n)$

To ensure that the probability found above is mathematically sound, we may wish to show that the sum of the probabilities over all n converges; surely, the sum over all n ought to be no more than 1. In this section we do so using Stirling's approximation.

Theorem 2.2. *The probability that the process ends after $2n+1$ steps approaches*

$$\frac{1}{2n^{3/2}\sqrt{\pi}}$$

for large n.

Proof. The asymptotics of the first two terms in $\mathbb{P}(n)$ are straightforward, so we begin by finding an approximate value of $\binom{2n}{n} = \frac{(2n)!}{n!n!}$. Using Stirling's approximation, that

$$k! \approx \sqrt{2\pi n}\left(\frac{n}{e}\right)^n$$

for large n, we find that

$$\begin{aligned}\binom{2n}{n} &= \frac{(2n)!}{n!n!} \\ &\approx \frac{2\sqrt{\pi n}\left(\frac{2n}{e}\right)^{2n}}{\left(\sqrt{2\pi n}\left(\frac{n}{e}\right)^n\right)^2} \\ &= \frac{2^{2n+1}\sqrt{\pi n}}{2\pi n} \\ &= \frac{2^{2n}}{\sqrt{\pi n}}.\end{aligned}$$

Substituting this into our expression for $\mathbb{P}(n)$, and using the approximation that $\frac{1}{n+1} \approx \frac{1}{n}$ for large n, we find that

$$\begin{aligned}\mathbb{P}(n) &= 2^{-(2n+1)} \cdot \frac{1}{n+1} \cdot \binom{2n}{n} \\ &\approx 2^{-(2n+1)} \cdot \frac{1}{n} \cdot \frac{2^{2n}}{\sqrt{\pi n}} \\ &= \boxed{\frac{1}{2n^{3/2}\sqrt{\pi}}}\end{aligned}$$

as desired. □

From this result we can see that the sum over all n will converge. The terms of the sum behave like $n^{-3/2}$, and it is well known that $\sum n^p$ converges for $p < -1$. However, using this same approximation gives us a surprising result for the expected number of tosses before the process ends.

Theorem 2.3. *The expected number of tosses before the process terminates diverges.*

Proof. From Theorem 2.2, we have that $\mathbb{P}(n)$ grows like $\frac{1}{2n^{3/2}\sqrt{\pi}}$. Then the expected number of coin tosses is given by

$$\sum_{n=0}^{\infty} \frac{2n+1}{2n^{3/2}\sqrt{\pi}} \approx \sum_{n=0}^{\infty} \frac{1}{\sqrt{\pi n}}.$$

In this case we have terms that grow like $n^{-1/2}$, which results in a divergent series. □

Remark. In this section we found that the probability that the process eventually ends converges. To further explore this question, we may wish to compute

$$\sum_{n=0}^{\infty} \mathbb{P}(n) = \sum_{n=0}^{\infty} 2^{-(2n+1)} \cdot \frac{1}{n+1} \cdot \binom{2n}{n}$$

exactly, to determine what the probability that the process ever ends is. At the moment this is difficult to compute; however, we now develop a generating function that makes this problem much easier.

3 Generating functions

In this section, we develop a generating function for the Catalan numbers using Dyck paths. We use this to motivate a similar formulation for $\mathbb{P}(n)$ and then to show that the sum of $\mathbb{P}(n)$ over all n is 1 (so the process is guaranteed to stop).

Theorem 3.1. *The generating function $G(x)$ for the Catalan numbers C_n can be expressed as*

$$G(x) = \frac{1 - \sqrt{1-4x}}{2x}.$$

Proof. Suppose there is some generating function for the Catalan numbers,

$$G(x) = \sum_{n=0}^{\infty} C_n x^n.$$

Using the fact that Catalan numbers enumerate the number of Dyck paths of length $2n$, we can find an equation relating $G(x)$ to itself.

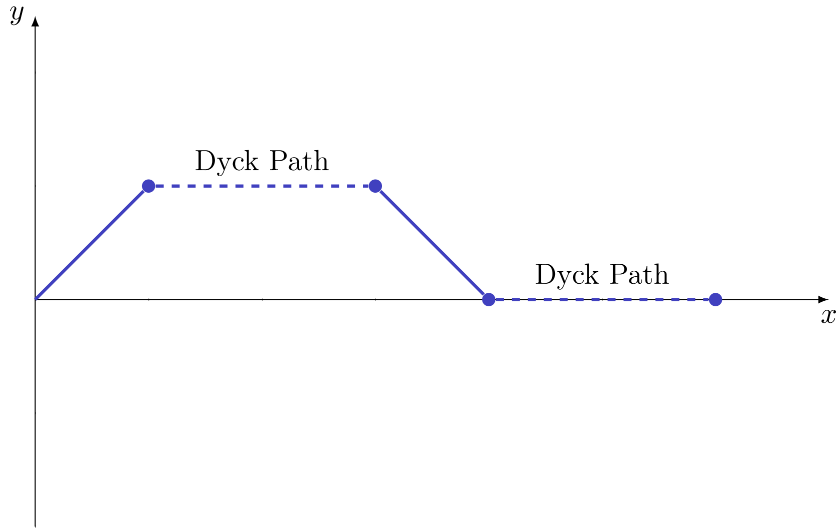

Figure 2: an illustration of a Dyck path broken down into 4 components.

Note that either a Dyck path has length of 0, or it has a positive length. The singular Dyck path of length 0 contributes a term of $1 \cdot x^0 = 1$ to $G(x)$. If the Dyck path has positive length, we know that it begins with a $+1$ step, then stays at or above 1 for some number of steps, then returns to 0 for the first time and stays at or above 0 for some number of steps. In this way, we can decompose every non-zero Dyck path as

$$\text{[first instance of 0 to 1]} \to \text{[Dyck path]} \to \text{[first 1 to 0]} \to \text{[Dyck path]}$$

as illustrated in Figure 2. As such, the generating function for non-zero Dyck paths is $G(x)^2$, since each path includes two smaller Dyck paths. Since the additional $+1$ and -1 steps contribute an additional length of 2, we multiply $G(x)^2$ by x so that the exponents correspond to the correct lengths of Dyck paths. As such, our equation is

$$G(x) = 1 + xG(x)^2.$$

We can solve for $G(x)$ using the quadratic formula to find

$$G(x) = \boxed{\frac{1 \pm \sqrt{1-4x}}{2x}}.$$

By noting that $G(0)$ should equal 1, we discard the positive root, since it gives $G(0)$ undefined. The negative solution works by L'Hôpital's rule. □

Remark. We could continue by expanding this into a power series to extract the coefficients of each x^n term, but we would simply recover the expression for C_n from earlier. Interested readers may find a complete proof in [1].

3.1 Generating function for $\mathbb{P}(n)$

We may employ a similar strategy to find an equation for $H(x)$, where

$$H(x) = \sum_{n=0}^{\infty} \mathbb{P}(n) \cdot x^n.$$

Theorem 3.2. *The generating function $H(x)$ satisfies*

$$H(x) = \frac{1 - \sqrt{1-x}}{x}.$$

Proof. We wish to relate $H(x)$ to itself in a similar manner to the previous proof. Note that $2H(x)$ is the probability that the first $2n$ steps form a Dyck path, since the final step is -1 with probability $1/2$. We will use the generating function for this value throughout the proof.

If $n = 0$, we trivially get a Dyck path of length 0, so this contributes a 1 to the generating function $2H(x)$. Otherwise, the Dyck path follows a pattern identical to the one above,

$$[\text{first } 0 \text{ to } 1] \to [\text{Dyck path}] \to [\text{first } 1 \text{ to } 0] \to [\text{Dyck path}].$$

If we group the first two pieces together, we have a $1/2$ chance that the sequence starts with a $+1$, followed by the probability that the sequence is a Dyck path, giving us $(1/2)\cdot 2H(x) = H(x)$. Similarly, the later two pieces can be grouped to get the same generating function $H(x)$. From this we attain the generating function

$$2H(x) = 1 + xH(x)^2$$

which we can solve to get a closed form for the generating function,

$$H(x) = \frac{2 - \sqrt{4-4x}}{2x} = \boxed{\frac{1-\sqrt{1-x}}{x}}.$$

We once again discard the positive solution by noting $H(0) = 1/2$ and using L'Hôpital's rule. □

3.2 Probability that the process ends

The generating function $H(x)$ gives us the necessary tools to find the probability that the coin toss process ends.

Theorem 3.3. *The coin tossing process is guaranteed to end. Equivalently,*

$$\sum_{n=0}^{\infty} \mathbb{P}(n) = 1.$$

Proof. We can write

$$\sum_{n=0}^{\infty} \mathbb{P}(n) = \sum_{n=0}^{\infty} \mathbb{P}(n) \cdot 1^n = \sum_{n=0}^{\infty} \mathbb{P}(n) \cdot x^n \Big|_{x=1}$$

which is equivalent to the generating function just found, evaluated at 1. We can now evaluate this sum as:

$$\sum_{n=0}^{\infty} \mathbb{P}(n) = H(1) = \frac{1-\sqrt{1-1}}{1} = \boxed{1}.$$

So we have found that the sum over all n is 1; that is, it is certain that the process stops after some number of odd steps. □

Remark. In the language of random walks, this means that a random walk starting at 0 will necessarily pass through -1, a remarkable fact. We can extend this logic to show that a one-dimensional symmetric random walk will pass through all points; this is formalized in Section 4.1.

4 Extensions

4.1 Implication for Random Walks

It is well known that a symmetric random walk on $\mathbb{Z}$ will reach every point with probability 1. Proving this is usually done analytically by showing that the expected number of returns to any point is infinite [2]. Here, we use the result from Theorem 3.3 to prove this.

Theorem 4.1. *The symmetric random walk on $\mathbb{Z}$ is guaranteed to reach every point in $\mathbb{Z}$.*

Proof. We will prove by induction that the symmetric random walk reaches every negative point. Then, by symmetry, it is clear that it will reach every positive point as well.

First, by Theorem 3.3, a random walk starting at 0 on $\mathbb{Z}$ will reach -1 after some finite number of steps. (Note that this number of steps could be arbitrarily large, but because there are only countably many possible number of steps, this number of steps must eventually be reached.) This is our base case.

Now, suppose we are at some point n on the number line after some finite number of steps. We may now shift the number line by $-n$; we are now at 0 and guaranteed to reach -1 after some number of steps. Equivalently, in the original number line, we are guaranteed to reach $n-1$ after finitely many steps. As such, we are guaranteed to reach every negative integer after some finite number of steps, as desired. □

4.2 Using Unfair Coins

Throughout this paper, we have only discussed the situation in which the coin tossed is fair. Now we consider the use of an unfair coin. In the language of random walks, this walk is no longer symmetric as $p \neq 0.5$. First, we generalize the probability of ending after $2n+1$ coin tosses.

Theorem 4.2. *The probability that this process ends after* $2n+1$ *coin tosses is*

$$\mathbb{P}_p(n) = \frac{p^n \cdot (1-p)^{n+1}}{n+1}\binom{2n}{n}.$$

Proof. The proof is similar to that of Theorem 2.1. The number of paths that end after $2n+1$ is the number of Dyck paths of length $2n$, of which there are $\frac{1}{n+1}\cdot\binom{2n}{n}$. Each of these paths requires n heads and $n+1$ tails, which happens with probability $p^n \cdot (1-p)^{n+1}$. □

Naturally, we may ask about the asymptotics of this procedure; does it still behave like $n^{-3/2}$, or some other rational function? Applying Stirling's approximation, we see that it actually decreases exponentially in n.

Theorem 4.3. *The probability that the process ends after* $2n+1$ *steps approaches*

$$P_p(n) = \Big(p\cdot(1-p)\cdot 4\Big)^n \cdot \frac{1-p}{n^{3/2}\sqrt{\pi}}$$

for large n.

Proof. From Theorem 2.2 we have that $\binom{2n}{n}$ approaches $\frac{2^{2n}}{\sqrt{\pi n}}$. Using the approximation that $\frac{1}{n+1} \approx \frac{1}{n}$ for large n, we have that

$$\begin{aligned}\mathbb{P}_p(n) &= \frac{p^n \cdot (1-p)^{n+1}}{n+1}\binom{2n}{n}\\ &\approx \Big(p\cdot(1-p)\cdot 4\Big)^n \cdot \frac{1-p}{n^{3/2}\sqrt{\pi}}\end{aligned}$$

as desired.

From this expression, the function will behave exponentially for large n, as an exponential dominates the $n^{-3/2}$ term. It is also necessary to show that $p\cdot(1-p)\cdot 4 < 1$ for $0 < p < 1$ (and $p \neq 1/2$) to see that the probability does, in fact, tend to 0. We can see this by determining that the function $f(p) = p\cdot(1-p)\cdot 4$ has a maximum of 1 at $p = 1/2$. □

Another worthwhile question to consider is whether this process is guaranteed to terminate, as it was with a fair coin. Here, we show that this process is *not* guaranteed to terminate for $p < 0.5$. Before doing so, however, we must find a generating function for the probability $\mathbb{P}_p(n)$ to allow us to use an argument analogous to that of Theorem 3.3.

Theorem 4.4. *The generating function $H_p(x) = \sum_{n=0}^{\infty} \mathbb{P}_p(n) \cdot x^n$ satisfies*

$$H_p(x) = \frac{1 - \sqrt{1 - 4xp(1-p)}}{2xp}$$

Proof. We begin by noting that the probability of the process ending after $2n+1$ steps is equivalent to the probability of the first $2n$ steps forming a Dyck path, multiplied by $1-p$, the probability of the $2n+1^{th}$ step being -1. From this fact, it is natural to define

$$H_p'(x) = \frac{1}{1-p} H_p(x).$$

Now $H_p'(x)$ gives the probability that the first $2n$ steps form a Dyck path. We may now proceed with an argument analogous to the generating function arguments we have used before.

We wish to relate $H_p'(x)$ to itself. Once again, we consider two separate cases, the case of a Dyck path with length 0 and the case of a Dyck path with positive length. For the former, the path of length 0 is trivially a Dyck path, which contributes a 1 to $H_p'(x)$. For paths of positive length, we may partition the path into four segments as earlier,

$$\text{[first 0 to 1]} \to \text{[Dyck path]} \to \text{[first 1 to 0]} \to \text{[Dyck path]}.$$

The first piece, a single $+1$ step, happens with probability p. The second is a Dyck path, represented by $H_p'(x)$. The third is a single -1 step which happens with probability $1-p$, and the final piece is yet another Dyck path represented by $H_p'(x)$. Combining this, we get the relation

$$H_p'(x) = 1 + x(p)(1-p)(H_p'(x))^2.$$

Solving for $H_p'(x)$, we have

$$H_p'(x) = \frac{1 - \sqrt{1 - 4xp(1-p)}}{2xp(1-p)},$$

where the positive solution is discarded as before. Now, to find $H_p(x)$, the original generating function, we multiply by $1-p$ and attain the desired expression. □

This generating function will now allow us to find the probability that this process terminates.

Theorem 4.5. *The probability that the process terminates is*

$$\sum_{n=0}^{\infty} \mathbb{P}_p(n) = \frac{1-p}{p}.$$

Proof. We use a similar approach as for Theorem 3.3. We may write

$$\sum_{n=0}^{\infty} \mathbb{P}_p(n) = \sum_{n=0}^{\infty} \mathbb{P}_p(n) \cdot x^n \Big|_{x=1}.$$

This is equivalent to evaluating the generating function above at $x = 1$. Doing so, we find

$$H_p(1) = \frac{1 - \sqrt{1 - 4p(1-p)}}{2p} = \frac{1 - \sqrt{(2p-1)^2}}{2p} = \frac{1-p}{p}$$

as desired. □

Remark. From this result, we see that this process is *not* guaranteed to end for unfair coins with $p > 1/2$, and equivalently that random walks on $\mathbb{Z}$ only reach every point when they are symmetric.

4.3 Different Stopping Criteria

In this final section, we consider another variation of the original problem, in which we flip a fair coin but wish to get 2 more tails than heads. Note that in this case we require an even number of steps to reach -2.

Theorem 4.6. *The probability that this revised process ends after $2n$ steps is equivalent to*

$$\sum_{m=0}^{n-1} \mathbb{P}(m) \cdot \mathbb{P}(n-(1+m)).$$

Proof. Clearly, for this process to end, there will be some point at which the number of tails flipped is one more than the number of heads. We will consider the process as two separate parts; the first part ending after the *first* time there is one more tail than heads, and the second part lasting until there are two more tails than heads.

Now, each of the two separate parts can be modeled exactly as the original problem in Section 2.1. (Note that the second part is the same as the original problem shifted down by 1; as we start at -1 and end at -2.) If we suppose the first part of the process has a length of $2m+1$, then the second half must have length

$$2n-(2m+1) = 2(n-(1+m))+1.$$

The probabilities of these happening are given by $\mathbb{P}(m)$ and $\mathbb{P}(n-(1+m))$, respectively. Therefore, the probability of both happening is the product of these probabilities.

To finish, we consider these probabilities for every possible length of the first part, which range from $m=0$ to $m=n-1$. Summing over these values, we get the probability being

$$\sum_{m=0}^{n-1} \mathbb{P}(m) \cdot \mathbb{P}(n-(1+m))$$

as desired. □

Suppose we wished to find the probability that this process ends at all; that is, the sum of the probability above over all n. This is guaranteed to happen, in agreement with Theorem 4.1.

Theorem 4.7. *The probability that this process ends is* 1.

Proof. We wish to find the sum

$$\sum_{n=1}^{\infty}\sum_{m=0}^{n-1} \mathbb{P}(m) \cdot \mathbb{P}(n-(1+m)).$$

This probability is equal to

$$\begin{aligned}
&[\mathbb{P}(0)\cdot\mathbb{P}(0)]+\\
&[\mathbb{P}(1)\cdot\mathbb{P}(0)+\mathbb{P}(0)\cdot\mathbb{P}(1)]+\\
&[\mathbb{P}(2)\cdot\mathbb{P}(0)+\mathbb{P}(1)\cdot\mathbb{P}(1)+\mathbb{P}(0)\cdot\mathbb{P}(2)]+\\
&[\mathbb{P}(3)\cdot\mathbb{P}(0)+\mathbb{P}(2)\cdot\mathbb{P}(1)+\mathbb{P}(1)\cdot\mathbb{P}(2)+\mathbb{P}(0)\cdot\mathbb{P}(3)]+...
\end{aligned}$$

An application of the distributive property allows us to see that this is equal to

$$[\mathbb{P}(0)+\mathbb{P}(1)+\mathbb{P}(2)+...]\cdot[\mathbb{P}(0)+\mathbb{P}(1)+\mathbb{P}(2)+...]$$

or equivalently

$$\Big(\sum_{x=0}^{\infty}\mathbb{P}(x)\Big)^2.$$

From Theorem 3.2, we have that this sum is 1, so the sum squared is also 1. □

References

[1] Jarvis, Frazer. *Catalan numbers.* Retrieved March 26, 2018 from `http://www.afjarvis.staff.shef.ac.uk/maths/jarvisspec01.pdf`

[2] Johnston, Derek. *An Introduction to Random Walks.* Retrieved March 29, 2018 from `http://www.math.uchicago.edu/ may/VIGRE/VIGRE2011/REUPapers/Johnston.pdf`