Coins Flips, Fair Casinos and Martingales

An Introduction to the Generalized ABRACADABRA Theorem

Eytan Chong

March 23, 2025

Contents

1 A Simple Coin-Flip Problem

Our journey begins with a classic coin-flip problem:

> *A fair coin is flipped repeatedly until a given sequence of Heads and Tails appears. On average, how many times is the coin flipped?*

Let us unpack this problem and phrase it mathematically. To do so, we introduce the following notation and terminology.

Definition 1. Let $\mathcal{A}$ be an **alphabet**, which is the set of characters from which words are constructed. In the case of coin-flips, $\mathcal{A} = \{H, T\}$, where H represents Heads and T represents Tails.

Definition 2. A **terminator** is a word that terminates the coin-flipping. The set of all terminators is denoted $\mathcal{T}$.

Definition 3. A word w is said to be **immediately terminated** under $\mathcal{T}$ if

- w ends with a terminator $t \in \mathcal{T}$; and
- w contains no other terminators.

The set of all words immediately terminated under $\mathcal{T}$ is denoted $\mathcal{I}_{\mathcal{T}}$.

Example 4. Let $\mathcal{T} = \{HHT, THH\}$. That is to say, we stop flipping the coin the moment we get HHT or THH. The set of words we might get when playing the game is then

$$\mathcal{I}_{\mathcal{T}} = \{HHT, HHHT, HHHHT, \ldots, THH, HTHH, TTHH, \ldots\}.$$

Note that the word $HTHHT$, despite ending with HHT, is not in $\mathcal{I}_{\mathcal{T}}$. This is because it contains the another terminator: $HTHHT$.

Note that in order for our problem to make sense, a terminator cannot contain another terminator. Equivalently, $\mathcal{T} \subseteq \mathcal{I}_{\mathcal{T}}$. This prevents nonsensical scenarios, such as $\mathcal{T} = \{HTT, HT\}$ or $\mathcal{T} = \{HTH, TH\}$.

We now rephrase our original problem:

> *Let $W_{\mathcal{T}}$ be a word constructed by randomly concatenating letters until $W_{\mathcal{T}} \in \mathcal{I}_{\mathcal{T}}$, and let $L_{\mathcal{T}} = |W_{\mathcal{T}}|$ be its length. What is $\mathbb{E}[L_{\mathcal{T}}]$?*

For now, we will simplify the problem and assume $|\mathcal{T}| = 1$. In the following subsections, we will present two common approaches one might take in answering this (simplified) problem.

1.1 A Common (but wrong!) Approach

Suppose $\mathcal{T} = \{T\}$. In context, this means that we stop flipping the coin the moment we get Tails. Intuitively, because the probability of getting T is $1/2$, one might guess that we will, on average, get one T every two flips, so

$$\mathbb{E}\big[L_{\{T\}}\big] = \frac{1}{\mathbb{P}[T]} = \frac{1}{1/2} = 2.$$

This is indeed the correct answer.

Suppose now that $\mathcal{T} = \{TH\}$. Following a similar line of reasoning, one might conclude that

$$\mathbb{E}\big[L_{\{TH\}}\big] = \frac{1}{\mathbb{P}[TH]} = \frac{1}{1/4} = 4,$$

which is once again the correct answer.

However, this argument quickly breaks down once we consider more complicated terminators. For instance, if $\mathcal{T} = \{THT\}$, the above pattern suggests that

$$\mathbb{E}\big[L_{\{THT\}}\big] = \frac{1}{\mathbb{P}[THT]} = \frac{1}{1/8} = 8.$$

However, empirical evidence suggests that $\mathbb{E}\big[L_{\{THT\}}\big]$ is actually 10.

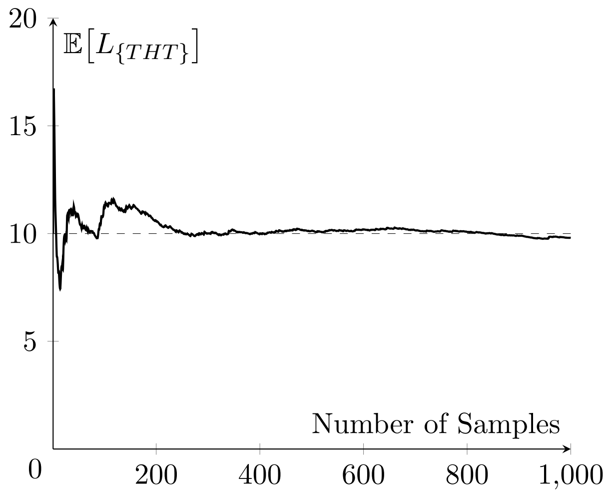

Figure 1: 1 thousand samples of $L_{\{THT\}}$ suggests that $\mathbb{E}[L_{\{THT\}}]$ is 10, not 8.

How, then, do we determine $\mathbb{E}[L_{\mathcal{T}}]$ accurately all the time?

1.2 Case Closed – or is it?

One might observe that depending on which face the coin lands on, the expected number of flips will change accordingly. Thus, by analysing all possible cases, we can form an equation in $\mathbb{E}[L_{\mathcal{T}}]$, which we can then easily solve.

To facilitate further discussion, we first introduce the notion of a left- and right-slice of a word.

Definition 5. The **left-slice** of a word w, denoted $L_n(w)$, refers to the first n characters of w. Analogously, the **right-slice** of w, denoted $R_n(w)$, refers to the last n characters of w.

Example 6. Let $w = HTHH$. The following table gives the left-slices of w for different n.

n	$L_n(w)$	$R_n(w)$
1	H	H
2	HT	HH
3	HTH	THH
4	HTHH	HTHH

To illustrate the method of case-by-case analysis, consider $\mathcal{T} = \{THT\}$.

1. If the first coin is H, we have effectively wasted one flip since starting with H does not contribute to getting THT. Thus,

$$\mathbb{E}[L_{\mathcal{T}} \mid L_1(W_{\mathcal{T}}) = H] = \mathbb{E}[L_{\mathcal{T}}] + 1.$$

2. If the first coin is T, we have two subcases to consider:
   a) If the second coin is T, we have effectively "gone back" to the case where our first coin is T. Thus,

$$\mathbb{E}[L_{\mathcal{T}} \mid L_2(W_{\mathcal{T}}) = TT] = \mathbb{E}[L_{\mathcal{T}}].$$

   b) If the second coin is H, we have two more subcases to consider:

i. If the third coin is T, we have reached the terminator. Thus,

$$\mathbb{E}[L_{\mathcal{T}} \mid L_3(W_{\mathcal{T}}) = THT] = 3.$$

ii. If the third coin is H, we have effectively "gone back" to the case where our first coin is H. This means that we wasted 3 flips, so

$$\mathbb{E}[L_{\mathcal{T}} \mid L_3(W_{\mathcal{T}}) = THH] = \mathbb{E}[L_{\mathcal{T}}] + 3.$$

Since all words must start with either H, TT, THT or THH, by the law of total expectation,

$$\begin{aligned}\mathbb{E}[L_{\mathcal{T}}] = \quad & \mathbb{E}[L_{\mathcal{T}} \mid L_1(W_{\mathcal{T}}) = H]\,\mathbb{P}[L_1(W_{\mathcal{T}}) = H] \\ + & \,\mathbb{E}[L_{\mathcal{T}} \mid L_2(W_{\mathcal{T}}) = TT]\,\mathbb{P}[L_2(W_{\mathcal{T}}) = TT] \\ + & \,\mathbb{E}[L_{\mathcal{T}} \mid L_3(W_{\mathcal{T}}) = THT]\,\mathbb{P}[L_3(W_{\mathcal{T}}) = THT] \\ + & \,\mathbb{E}[L_{\mathcal{T}} \mid L_3(W_{\mathcal{T}}) = THH]\,\mathbb{P}[L_3(W_{\mathcal{T}}) = THH].\end{aligned}$$

Because the coin is fair, the probability that $L_n(W_{\mathcal{T}}) = w$ for some arbitrary word w is simply $1/2^n$. Substituting the values we found,

$$\mathbb{E}[L_{\mathcal{T}}] = \frac{\mathbb{E}[L_{\mathcal{T}}] + 1}{2^1} + \frac{\mathbb{E}[L_{\mathcal{T}}]}{2^2} + \frac{3}{2^3} + \frac{\mathbb{E}[L_{\mathcal{T}}] + 3}{2^3}.$$

After simplification, we get $\mathbb{E}[L_{\mathcal{T}}] = 10$, which aligns with the results obtained from our simulation.

Of course, this is a perfectly sound solution to the problem, and one can always calculate the correct value of $\mathbb{E}[L_{\mathcal{T}}]$ using this algorithm. However, it becomes incredibly inefficient and tedious when the terminators become more complicated, rendering it effectively useless.

1.3 All In!

Fortunately for us, there is a simple and elegant way to calculate $\mathbb{E}[L_{\mathcal{T}}]$.

To set the stage, imagine that you work as a dealer at D'Casino. Unlike most casinos, D'Casino is a *fair casino*; it neither wins nor loses money in the long run. Furthermore, there is only one game available for play at D'Casino:

> *Each round, you, the dealer, flip a fair-coin. Gamblers go all-in, betting on the outcome of this coin-flip. If they win, they double their money, and they play again. If they lose, the casino takes everything and they go home empty-handed. This repeats until a terminator (say THT) appears, at which point the casino closes and everybody goes home.*

As luck would have it, a group of gamblers, obsessed with the sequence THT, frequents D'Casino. Every flip, a new gambler from this group arrives with \$1 and plays the game, hoping that the subsequent flips appear T, H, T in that order.

To illustrate this, suppose the coins come up $HTTHT$. Let R_n and C_n be the total revenue earned and total cost incurred by the gamblers at the nth flip, respectively.

n	**Event**	R_n	C_n
1	Gambler #1 walks into the casino and bets \$1 that the first coin-flip is a T. He loses and leaves the casino.	0	1
2	Gambler #2 walks into the casino and bets \$1 that the second coin-flip is a T. He wins and doubles his money. He then bets \$2 that the third coin-flip is an H. He loses the bet, and leaves the casino with a net loss of \$1.	0	2
3	Gambler #3 walks into the casino and bets \$1 that the third coin-flip is a T. He wins and doubles his money. He then bets \$2 that the fourth coin-flip is an H. He wins and doubles his money once again. Finally, he bets \$4 that the fifth coin-flip is a T. Miraculously, he wins, earning \$7 overall. For consistency, we will record this as a win of \$8, and a loss of \$1.	8	3
4	Gambler #4 walks into the casino and bets \$1 that the fourth coin-flip is a T. He immediately loses.	8	4
5	Lastly, Gambler #5 walks into the casino and bets \$1 that the fifth coin-flip is a T. He wins the bet, pocketing \$2. However, after the fifth coin-flip, the casino closes, since the terminator THT has just been flipped. Like before, we record this as a win of \$2 and a loss of \$1.	10	5

The total revenue earned, R, is hence \$10, while the total cost incurred, C, is \$5. Let $X = R - C = \$5$ be the profit made by the gamblers.

We now make three key observations:

- Because of the way we recorded losses, each gambler contributes only \$1 to C. Additionally, by our set-up, each gambler corresponds to exactly one coin-flip. Thus, *the total cost incurred is equal to the number of coin-flips made.* Taking expectations, we have
$$\mathbb{E}[C] = \mathbb{E}\left[L_{\{THT\}}\right]. \tag{1.1}$$
- Only the last three gamblers stand a chance to earn money; R depends solely on the last three gamblers. Because the last three coin-flips must always be THT, it follows that R *is a constant*, so
$$\mathbb{E}[R] = R = 10. \tag{1.2}$$
- Because the casino is fair, the expected profit should be 0. Since $X = R - C$, we have
$$\mathbb{E}[R] = \mathbb{E}[C]. \tag{1.3}$$

Chaining (1.1), (1.2) and (1.3) yields
$$\mathbb{E}\left[L_{\{THT\}}\right] = \mathbb{E}[C] = \mathbb{E}[R] = R = 10,$$
which is indeed what we got using case-by-case analysis earlier.

As we will see in a later section, this “fair casino” method, as compared to case-by-case analysis, can be generalized and applied to different problems much more easily. Before we get ahead of ourselves, we first iron out the details and formalize our argument using the language of martingales.

2 Martingales and the Optional Stopping Theorem

2.1 Martingales

Put simply, a martingale is a random process, typically represented by a sequence of random variables $\{X_n\}_{n\in\mathbb{N}}$, which models a gambler's fortune in a fair game.

To motivate our formal definition of a martingale, consider the following game:

> *Flip a fair coin. If it comes up Heads, we win \$1, but if it comes Tails, we lose \$1. Repeat this process forever.*

There are two key properties that this game possesses:

- The coin is *fair*: The probability of getting Heads will always remain at 50%. Likewise, the probability of getting Tails will always remain at 50%.
- Each coin-flip is *independent*: The outcomes of past flips will not influence the outcome of future flips.[1]

As a consequence of this fairness, ***our expected wealth after the next flip, given everything we currently know*** (e.g. outcomes of previous flips, any observed patterns, etc.)***, is exactly our current wealth.***

We can generalize this idea. Let X_n represent some quantity that we are interested in (e.g. total wealth) at round n of a fair game. Let $\mathcal{F}_n$ represent all the information we know up to round n. Then, because of fairness and independence, ***given full knowledge of the past, the future expectation of*** X ***is equal to the current value of*** X. Mathematically,

$$\mathbb{E}[X_{n+1} \mid \mathcal{F}_n] = X_n.$$

This brings us to the formal definition of a martingale.

Definition 7. A sequence of random variables $\{X_n\}_{n\in\mathbb{N}}$ is a **martingale** with respect to a **filtration** $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ if,

- $\mathbb{E}[X_n]$ is finite, and
- $\mathbb{E}[X_{n+1} \mid \mathcal{F}_n] = X_n$ (the "fairness/independence" principle).

Remark. Typically, the filtration $\mathcal{F}_n$ is simply the information we get from observing the past outcomes $X_0, X_1, \ldots, X_n$. This is commonly referred to as the **natural filtration**.

Apart from gambling, another context in which martingales can be applied in is the stock market. Let the sequence $\{X_n\}$ represents the price of a stock traded on the stock exchange, and let the filtration $\{\mathcal{F}_n\}$ represent the price history of the stock. In this context, the fairness principle states "the price of the stock tomorrow, given its price history, is equal to the price of the stock today." Intuitively, this should be a reasonable assumption: if one expects the price of the stock to double from \$10 today to \$20 tomorrow, they would be incentivized to buy the stock today and sell tomorrow. Meanwhile, those that own the stock would be incentivized to hold on to the stock and sell it tomorrow. This increase in demand and fall in supply bids up the price of the stock to \$20 *today*, which is exactly what we expected the price to be *tomorrow*. Hence, the sequence $\{X_n\}$ is a martingale.

[1] Though these properties may seem trivial, it is nevertheless important to highlight them as it may sometimes run against our human intuition. For example, the probability that the next flip is Heads, given that the previous 100 flips were all tails, will still remain at 1/2. This misguided belief that we are more likely to win after a series of losses is commonly known as the Gambler's Fallacy.

Exercise A. Determine if the following sequences are martingales.

i. $X_n = 10$,

ii. $X_n = n$,

iii. Let $X_1, X_2, X_3, \dots$ be a sequence of independent random variables, each equal to -1 with probability $1/2$ and 1 with probability $1/2$. Let $Y_0 = 0$ and $Y_n = X_1 + X_2 + \cdots + X_n$ for $n > 0$. Is Y_n a martingale with respect to X_n?

2.2 Stopping Strategies and Stopping Times

Given a martingale $\{X_n\}$, we might ask: what is our expected payout under a given strategy? To answer this question, we must first define what it means to stop playing a game.

Definition 8. A **stopping time** τ with respect to a filtration $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$, is a random variable taking values in $\mathbb{N} \cup \{\infty\}$ such that for all $n \geq 0$, the event $\{\tau = n\}$ is a member of $\mathcal{F}_n$. In other words, $\{\tau = n\}$ can be written as an event depending solely on $X_0, X_1, \dots, X_n$. This event is called the gambler's **stopping strategy**.

In layman terms, τ can be thought of as the round at which a gambler quits playing the game. In addition, the condition $\{\tau = n\} \in \mathcal{F}_n$ means that the gambler quits using only information available to him before round n; he cannot see into the future (view the outcome of $X_{n+1}, X_{n+2}, \dots$) to decide when to stop playing.

Example 9. Suppose a gambler employs a stopping strategy where he quits after playing 10 games. Then his stopping time is simply $\tau = 10$.

Another gambler may employ a different stopping strategy, opting to quit after losing three times in a row. If the gambler plays a fair game with a \$1 stake, the event $\{\tau = n\}$ can be expressed as

$$\Bigg\{\underbrace{X_1 = -1,\ X_2 = -1,\ X_3 = -1,\ \dots,\ X_{n-3} = -1}_{\text{all losses}},\ \underbrace{X_{n-2} = 1,\ X_{n-1} = 1,\ X_n = 1}_{\text{3 wins in a row}}\Bigg\}.$$

Exercise B. Suppose a gambler plays a fair game with a \$1 stake per round. Determine if the following events are stopping strategies.

i. The third time the gambler loses in a row.

ii. Two rounds before the gambler profits \$50.

iii. The first time the gambler profits \$50 or goes bankrupt.

iv. The first time the gambler starts a sequence of 10 losses in a row.

2.3 The Optional Stopping Theorem

With a broader range of stopping strategies available to us, we extend our original question and ask: is there a stopping strategy that allows us to expect a payout greater than our initial amount? As it turns out, the answer is generally a no.

Theorem 10 (Doob's Optional Stopping Theorem)**.** Let $\{X_n\}_{n\in\mathbb{N}}$ be a martingale and let τ be a stopping time, both with respect to a filtration $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$. Then $\mathbb{E}[X_\tau] = X_0$ if at least one of the following holds:

1. The martingale is bounded.
2. The stopping time τ is bounded.
3. The stopping time τ has finite expectation, and all increments of X are bounded, i.e. there exists a constant C such that for all n,

$$|X_{n+1} - X_n| \leq C.$$

The Optional Stopping Theorem (OST) tells us that as long as our stopping strategy is *reasonable enough*, our expected payout, $\mathbb{E}[X_\tau]$, must be equal to the amount we started with, X_0. One can think of the OST as the mathematician's version of the adage "*there is no edge in a fair game*".

To see why all reasonable strategies obey the OST, suppose we somehow came up with a profitable strategy. That is, we managed to force $\mathbb{E}[X_\tau] > X_0$. Then this strategy either

- breaks the validity of our stopping time, or
- breaks the conditions of the OST.

An invalid stopping time implies that we can somehow look into the future, which is clearly impossible. Furthermore, the three conditions of the OST are hard to break in real life:

- Casinos typically have bet limits, so $|X_n|$ is typically bounded and we cannot break the first condition.
- Gamblers have finite lifespans, so τ must also be bounded and we cannot break the second condition.
- For similar reasons (bet limits and finite lifespans), the third condition also cannot be broken.

Thus, a profitable strategy is nigh impossible to come up with, and so for all practical purposes, any reasonable strategy we come up with obeys the OST.

3 ABRACADABRA

We now formalize our elegant solution using martingales.

As before, define R_n and C_n to be the total revenue earned and total cost incurred by the gamblers at the nth flip. Define $X_n = R_n - C_n$ to be the combined wealth of the gamblers at the nth flip. Let our stopping strategy be the event that a terminator appears, and let τ be the stopping time under this strategy.

We now show that $\{X_n\}$ satisfies the two defining criteria of a martingale:

- The X_n is at a maximum when all n gamblers win. Likewise, X_n is at a minimum when all n gamblers lose. We thus have the bounds

  $$|X_n| \leq n \cdot 2^n,$$

  so $\mathbb{E}[X_n]$ must also be bounded and hence finite.

- Let A_n be the total wealth of gamblers that have lost before the nth flip. Correspondingly, let B_n be the total wealth of gamblers that are still playing at the nth flip. Since the coin is fair and independent, and the gamblers bet double-or-nothing, we have

  $$\mathbb{E}[B_{n+1} \mid \mathcal{F}_n] = \frac{1}{2}(2B_n) + \frac{1}{2}(0) = B_n.$$

  Since $X_n = A_n + B_n$, it follows that

  $$\mathbb{E}[X_{n+1} \mid \mathcal{F}_n] = \mathbb{E}[A_{n+1} \mid \mathcal{F}_n] + \mathbb{E}[B_{n+1} \mid \mathcal{F}_n] = A_n + B_n = X_n.$$

We now wish to show that our wealth martingale $\{X_n\}$, along with the stopping time τ, obeys Doob's OST. Specifically, we will show that τ and $|X_{n+1} - X_n|$ are both bounded (the third scenario). To do so, suppose D'Casino opens a new game:

> *Suppose the terminator has length n. Each round, n coins are flipped. If these n coins matches the terminator, we stop flipping. If not, we continue on with another round.*

Let the random variable Y represent the number of rounds played under this game. Since Y measures the number of failures (rounds) until a success (we obtain the terminator), we see that Y follows a geometric distribution with probability of success $p = 1/2^n$. It follows that

$$\mathbb{E}[Y] = \frac{1}{p} = 2^n.$$

Quite clearly, we expect this game to take a longer time to conclude than our original game. Since a total of nY coin-flips are made in this game, it follows that

$$\mathbb{E}[\tau] = \mathbb{E}[L_{\mathcal{T}}] \le \mathbb{E}[nY] = n \cdot 2^n,$$

so τ is bounded. Similarly, the maximum change in X is bounded above by

$$|X_{k+1} - X_k| \le n \cdot 2^n.$$

Thus, the third scenario of Doob's OST is satisfied, whence we conclude that

$$\mathbb{E}[X_\tau] = \mathbb{E}[X_0] = 0.$$

By the same key observations made earlier, we have

$$\mathbb{E}[L_{\mathcal{T}}] = \mathbb{E}[C_\tau] = \mathbb{E}[R_\tau] = R_\tau.$$

Since R_τ depends solely on the last few gamblers, we now have an easy way of calculating the expected number of coin flips $\mathbb{E}[L_{\mathcal{T}}]$.

Example 11. To illustrate, consider yet again the example where $\mathcal{T} = \{THT\}$. Our goal is to calculate R_τ. To do so, we simply imagine that the terminator THT has already been flipped and then work backwards.

- The third-last gambler wins \$$2^3$, since he sees THT.
- The second-last gambler wins \$0, since he sees H and immediately loses.
- The last gambler wins \$$2^1$, since he sees T before the casino closes.

Hence, we have $R_\tau = 10$, whence $\mathbb{E}[L_{\mathcal{T}}] = 10$.

We now introduce a more efficient way to calculate R_τ. First, we introduce the notion of *correlation* between two strings, as defined in [1]:

Definition 12. Let X and Y be two words. The **correlation** of X and Y, denoted $(X : Y)$, is a string of 1's and 0's with the same length as X. The ith bit (from the left) of $(X : Y)$ is determined as follows: place Y under X so that its leftmost character is under the ith character of X (from the left). Then, if all pairs of characters in the overlapping segment are identical, the ith bit of $(X : Y)$ is 1, else it is 0.

Example 13. Let $X = HTHTTH$ and $Y = HTTHT$. Then $(X : Y) = 001001$, as depicted below:

X:	H	T	H	T	T	H					
Y:	H	T	T	H	T						0
		H	T	T	H	T					0
			H	T	T	H	T				1
				H	T	T	H	T			0
					H	T	T	H	T		0
						H	T	T	H	T	1

Note that in general, $(X : Y) \neq (Y : X)$. For instance, using the same X and Y as the above example, we have $(Y : X) = 00010$.

Definition 14. We often wish to interpret the correlation $(X : Y)$ as a number in some base z, in which case we write $(X : Y)_z$. We call this the **correlation polynomial** of X and Y. For instance, in the above example, we have

$$(HTHTTH : HTTHT)_2 = 2^4 + 2^1.$$

With this new terminology, one can easily see that $R_{\mathcal{T}}$ is simply $(t, t)_2$, where t is the terminator.

Example 15. Once again, suppose $\mathcal{T} = \{THT\}$. Notice that $(THT : THT)$ is 101, as illustrated below:

X:	T	H	T			
Y:	T	H	T			1
		T	H	T		0
			T	H	T	1

Thus, $(THT : THT)_2 = 2^3 + 2^1 = 10$, which is precisely $R_{\mathcal{T}}$!
We can summarize this data using a matrix:

	T	H	T
THT	2^3		2^1

Unlike the case-by-case method we explored earlier, this method can easily be applied to terminators of longer lengths, as demonstrated in the following example.

Example 16. Let $\mathcal{T} = \{THHTHHTHH\}$. Computing correlations, we obtain the following matrix:

	T	H	H	T	H	H	T	H	H
$THHTHHTHH$	2^9			2^6			2^3		

Hence, $\mathbb{E}[L_{\mathcal{T}}] = 2^9 + 2^6 + 2^3 = 584$.
If we change the final character to a T, i.e. $\mathcal{T} = \{THHTHHTHT\}$, then our matrix becomes

	T	H	H	T	H	H	T	H	T
$THHTHHTHH$	2^9								2^1

Hence, $\mathbb{E}[L_{\mathcal{T}}] = 2^9 + 2^1 = 514$.

From the above examples, one can see that it is the "self-repetition" of the terminators that determines how long it takes to reach them. For instance, $THHTHHTHH$ self-repeats many times (at the sixth-last and third-last characters), while $THHTHHTHT$ only repeats itself at the last character.

Even if the alphabet $\mathcal{A}$ changes, the core idea remains the same:

Example 17. Consider the following problem:

> *An immortal monkey types one random character on a typewriter every second. How long would it take this monkey to type the word "ABRACADABRA"?*

In this context, our alphabet now contains 26 characters (A, B, C, etc.). To maintain the fairness of the casino, the payout for each win should now be 26 times the bet. Hence, the base of our correlation should be 26.

Comparing the correlation "ABRACADABRA" with itself, we see that our matrix is

	A	B	R	A	C	A	D	A	B	R	A
$ABRACADABRA$	26^{11}							26^4			26^1

The expected time taken is thus $26^{11} + 26^4 + 26$ seconds, or 116.4 million years.

More generally, we can state our result as follows:

Theorem 18 (ABRACADABRA Theorem)**.** Let $\mathcal{T} = \{t\}$ with alphabet $\mathcal{A}$. Then

$$\mathbb{E}[L_{\mathcal{T}}] = (t : t)_{|\mathcal{A}|}.$$

This result is known in the literature as the ABRACADABRA Theorem, named after the problem posed in Example 17.

As a bonus, one can express $(t : t)_{|\mathcal{A}|}$ in terms of left- and right-slices:

$$(t : t)_{|\mathcal{A}|} = \sum_{i=1}^{|t|} |\mathcal{A}|^i \, \mathbf{1}\{R_i(t) = L_i(t)\},$$

where the indicator function $\mathbf{1}(P)$ returns 1 if the statement P is true and 0 otherwise.

4 Extending our Results

We now turn our attention to solving the problem in its most general form. First, let us introduce one more piece of notation:

Definition 19. Let $t \in \mathcal{T}$. We define $[t]$ to be the set of all immediately terminated words w that end with t. Mathematically,

$$[t] = \left\{ w \in \mathcal{I}_{\mathcal{T}} : R_{|t|}(w) = t \right\}.$$

Example 20. If $\mathcal{T} = \{HHT, THH\}$, then

$$[HHT] = \{HHT, HHHT, HHHHT, \ldots\}$$

and

$$[THH] = \{THH, HTHH, TTHH, \ldots\}.$$

Suppose now that we have n terminators $\mathcal{T} = \{t_1, t_2, \ldots, t_n\}$. Since $[t_1]$, $[t_2]$, ..., $[t_n]$ form a partition of $\mathcal{I}_\mathcal{T}$, by the law of total expectation,

$$\mathbb{E}[L_\mathcal{T}] = \sum_{i=1}^{n} \mathbb{E}[L_\mathcal{T} \mid W_\mathcal{T} \in [t_i]]\, \mathbb{P}[W_\mathcal{T} \in [t_i]].$$

This is the approach we will take in calculating $\mathbb{E}[L_\mathcal{T}]$.

4.1 Probabilities

We now seek a formula for $\mathbb{P}[W_\mathcal{T} \in [t_i]]$.

To build our intuition, we first look at the case where $\mathcal{T} = \{THT, HTT\}$. Once again, suppose you are a dealer at D'Casino, whose job is to flip a fair coin until either THT or HTT appears. This time, D'Casino has introduced a slight modification to the coin-flipping game:

> *Each round, a fair coin is flipped, and a game is played between two parties, A and B:*
>
> *A goes all-in, betting \$$n$ on the outcome of the coin-flip. If A wins, they win \$$n$ from B, and they play again. If A loses, B takes everything, and the two stop playing.*
>
> *This repeats until a terminator appears, at which point the casino closes and everybody goes home.*

When compared to the original game, we see that A plays the role of the "gambler", while B plays the role of the "dealer".

Now, suppose we have two groups of gamblers, Group 1 and Group 2, that frequent D'Casino. The gamblers in Group 1 are obsessed with the sequence THT, while those in Group 2 are obsessed with HTT. Every flip, a new gambler from each group arrives with \$1. The two gamblers then play two games simultaneously:

- In the first game, the Group 1 gambler is A and the Group 2 gambler is B. The Group 1 gambler bets that the next few coin-flips will be THT.
- In the second game, the Group 2 gambler is A and the Group 1 gambler is B. The Group 2 gambler bets that the next few coin-flips will be HTT.

To illustrate, suppose the coin-flips come up $HHTHT$. Let $R_n(i, j)$ be the total revenue earned in Game i by Group j at the nth flip. Define also

$$X_n = \underbrace{(R_n(1,1) + R_n(2,1))}_{\text{Group 1's revenue}} - \underbrace{(R_n(1,2) + R_n(2,2))}_{\text{Group 2's revenue}}$$

to be the difference in revenue between the two groups.

We will focus on Game 1 first. Recall that Group 1 is A, betting on THT, while Group 2 is B. As this is almost identical to what we have seen before, we will keep it brief.

n	**Event**	$R_n(1,1)$	$R_n(1,2)$
1	Gambler #1 loses his first bet.	0	1
2	Gambler #2 loses his first bet.	0	2
3	Gambler #3 wins all three bets. He hence earns \$7 overall. For consistency, we record this as a gain of \$8 to Group 1, and a gain of \$1 to Group 2.	8	3
4	Gambler #4 loses his first bet.	8	4
5	Gambler #5 wins his first bet before the casino closes. Like before, we record this as a gain of \$2 to Group 1, and a gain of \$1 to Group 2.	10	5

We now do the same thing for Game 2. Here, Group 2 is A, betting on HTT, while Group 1 is B.

n	**Event**	$R_n(2,1)$	$R_n(2,2)$
1	Gambler #1 loses his second bet.	1	0
2	Gambler #2 loses his third bet.	2	0
3	Gambler #3 loses his first bet.	3	0
4	Gambler #4 wins his first and second bet, but the casino closes before his third. We record this as a gain of \$4 to Group 2, and a gain of \$1 to Group 1.	4	4
5	Gambler #5 loses his first bet.	5	4

Like before, we make three key observations:

- In both games, due to the way we recorded revenues, the revenue of the group playing B increases by one every round. That is, $R_n(1,2) = R_n(2,1) = n$. Hence,

$$X_n = (R_n(1,1) + R_n(2,1)) - (R_n(1,2) + R_n(2,2)) = R_n(1,1) - R_n(2,2). \quad (4.1)$$

- In both games, only the last three gamblers stood a chance to earn money. Thus, $R_n(1,1)$ and $R_n(2,2)$ *depend solely on the last three gamblers in each game.*

- Since the games are fair, the difference in revenue, $\{X_n\}$ is a martingale.[2] Hence, by the Optional Stopping Theorem,

$$\mathbb{E}[X_\tau] = \mathbb{E}[X_0] = 0. \quad (4.2)$$

From (4.1) and (4.2), it follows that

$$\mathbb{E}[R_\tau(1,1)] = \mathbb{E}[R_\tau(2,2)]. \quad (4.3)$$

We now calculate $\mathbb{E}[R_\tau(1,1)]$ and $\mathbb{E}[R_\tau(2,2)]$.

- Suppose $W_\mathcal{T} \in [THT]$. That is to say, the last three coin-flips are THT. Then Group 1's revenue in Game 1 is given by the correlation $(THT : THT)_2$:

$$\mathbb{E}[R_\tau(1,1) \mid W_\mathcal{T} \in [THT]] = (THT : THT)_2 = 2^3 + 2^1 = 10.$$

  Similarly, Group 2's revenue in Game 2 is given by the correlation $(THT : HTT)_2$:

$$\mathbb{E}[R_\tau(2,2) \mid W_\mathcal{T} \in [THT]] = (THT : HTT)_2 = 2^2 = 4.$$

[2] One can adapt the argument used in Section 3 to rigorously show that $\{X_n\}$ is indeed a martingale.

- Now suppose $W_\mathcal{T} \in [HTT]$. With completely analogous arguments, we see that

$$\mathbb{E}[R_\tau(1,1) \mid W_\mathcal{T} \in [HTT]] = (HTT : THT)_2 = 2^1 = 2$$

and

$$\mathbb{E}[R_\tau(2,2) \mid W_\mathcal{T} \in [HTT]] = (HTT : HTT)_2 = 2^3 = 8.$$

By the law of total expectation, it follows that

$$\mathbb{E}[R_\tau(1,1)] = 10\mathbb{P}[W_\mathcal{T} \in [THT]] + 2\mathbb{P}[W_\mathcal{T} \in [HTT]]$$

and

$$\mathbb{E}[R_\tau(2,2)] = 4\mathbb{P}[W_\mathcal{T} \in [THT]] + 8\mathbb{P}[W_\mathcal{T} \in [HTT]]$$

By (4.3), the two are equal, giving us the equation

$$10\mathbb{P}[W_\mathcal{T} \in [THT]] + 2\mathbb{P}[W_\mathcal{T} \in [HTT]] = 4\mathbb{P}[W_\mathcal{T} \in [THT]] + 8\mathbb{P}[W_\mathcal{T} \in [HTT]].$$

Further, by the law of total probability,

$$\mathbb{P}[W_\mathcal{T} \in [THT]] + \mathbb{P}[W_\mathcal{T} \in [HTT]] = 1.$$

This gives us a system of two linear equations in two unknowns, which we can easily solve:

$$\mathbb{P}[W_\mathcal{T} \in [THT]] = \mathbb{P}[W_\mathcal{T} \in [HTT]] = \frac{1}{2}.$$

4.2 Correlation Matrices and Probability Vectors

We now introduce a more elegant way of calculating such probabilities, using the language of linear algebra.

Once again, consider $\mathcal{T} = \{THT, HTT\}$. We begin by summarizing the data using a matrix, just like we did in Section 3:

	T	H	T	H	T	T
THT	2^3		2^1			2^1
HTT		2^2		2^3		

Because of its importance, we give this matrix a special name: the correlation matrix of $\mathcal{T}$.

Definition 21. The **correlation matrix** of $\mathcal{T} = \{t_1, \ldots, t_n\}$ with alphabet $\mathcal{A}$ is an $n \times n$ matrix $\mathbf{M}_\mathcal{T} = (m_{ij})$, where

$$m_{ij} = (t_j : t_i)_{|\mathcal{A}|}.$$

We also define another object called the probability vector of $\mathcal{T}$:

Definition 22. The **probability vector** of $\mathcal{T} = \{t_1, \ldots, t_n\}$, denoted $\mathbf{p}_\mathcal{T}$, is defined as

$$\mathbf{p}_\mathcal{T} = \begin{pmatrix} \mathbb{P}[W_\mathcal{T} \in [t_1]] \\ \vdots \\ \mathbb{P}[W_\mathcal{T} \in [t_n]] \end{pmatrix}.$$

Then, to get $\mathbb{E}[R_\tau(i,i)]$, we simply multiply the correlation matrix $\mathbf{M}_\mathcal{T}$ with the probability vector $\mathbf{p}_\mathcal{T}$ and read off the ith row of the resulting vector.

In context, we have

$$\mathbf{M}_{\mathcal{T}}\mathbf{p}_{\mathcal{T}} = \begin{pmatrix} 10 & 2 \\ 4 & 8 \end{pmatrix} \begin{pmatrix} \mathbb{P}[W_{\mathcal{T}} \in [THT]] \\ \mathbb{P}[W_{\mathcal{T}} \in [HTT]] \end{pmatrix} = \begin{pmatrix} 10\mathbb{P}[W_{\mathcal{T}} \in [THT]] + 2\mathbb{P}[W_{\mathcal{T}} \in [HTT]] \\ 4\mathbb{P}[W_{\mathcal{T}} \in [THT]] + 8\mathbb{P}[W_{\mathcal{T}} \in [HTT]] \end{pmatrix},$$

which perfectly aligns with what we calculated above.

To find the probabilities, we simply form our system of equations (by equating the two rows) and solve!

With this method in mind, let us compute some probabilities:

Example 23. Let $\mathcal{T} = \{TTT, THH\}$. For convenience, let T_1 and T_2 denote the events $W_{\mathcal{T}} \in [TTT]$ and $W_{\mathcal{T}} \in [THH]$ respectively. Our correlation matrix is of the form

	T	T	T	T	H	H
TTT	2^3	2^2	2^1			
THH			2^1	2^3		

The revenue earned by each group is hence

$$\begin{pmatrix} 14 & 0 \\ 2 & 8 \end{pmatrix} \begin{pmatrix} \mathbb{P}[T_1] \\ \mathbb{P}[T_2] \end{pmatrix} = \begin{pmatrix} 14\mathbb{P}[T_1] \\ 2\mathbb{P}[T_1] + 8\mathbb{P}[T_2] \end{pmatrix}.$$

Equating the two, we have the system of equations

$$\begin{cases} \mathbb{P}[T_1] + \mathbb{P}[T_2] = 1, \\ 14\mathbb{P}[T_1] = 2\mathbb{P}[T_1] + 8\mathbb{P}[T_2]. \end{cases}$$

Solving, we obtain $\mathbb{P}[T_1] = 2/5$ and $\mathbb{P}[T_2] = 3/5$.

This method even works for terminators of differing lengths.

Example 24. Let $\mathcal{T} = \{TT, THT\}$. Let T_1 and T_2 denote the events $W_{\mathcal{T}} \in [TT]$ and $W_{\mathcal{T}} \in [THT]$ respectively. Our correlation matrix is of the form

	T	T	T	H	T
TT	2^2	2^1			2^1
THT		2^1	2^3		2^1

The revenue earned by each group is hence

$$\begin{pmatrix} 6 & 2 \\ 2 & 10 \end{pmatrix} \begin{pmatrix} \mathbb{P}[T_1] \\ \mathbb{P}[T_2] \end{pmatrix} = \begin{pmatrix} 6\mathbb{P}[T_1] + 2\mathbb{P}[T_2] \\ 2\mathbb{P}[T_1] + 10\mathbb{P}[T_2] \end{pmatrix}.$$

Equating the two, we have the system of equations

$$\begin{cases} \mathbb{P}[T_1] + \mathbb{P}[T_2] = 1, \\ 6\mathbb{P}[T_1] + 2\mathbb{P}[T_2] = 2\mathbb{P}[T_1] + 10\mathbb{P}[T_2]. \end{cases}$$

Solving, we obtain $\mathbb{P}[T_1] = 2/3$ and $\mathbb{P}[T_2] = 1/3$.

This method also easily generalizes when we consider more terminators. For instance, if $|\mathcal{T}| = 3$, then our set-up consists of three groups that play three "cyclical" games:

- Group 1 (A) vs Group 2 (B)
- Group 2 (A) vs Group 3 (B)

- Group 3 (A) vs Group 1 (B)

Additionally, we will obtain the equation

$$\mathbb{E}[R_\tau(1,1)] = \mathbb{E}[R_\tau(2,2)] = \mathbb{E}[R_\tau(3,3)]$$

using an argument completely analogous to $|Tc| = 2$ case (consider the balance between any two groups and invoke the OST).

Example 25. Let $\mathcal{T} = \{THT, TTT, HHH\}$. Let T_1, T_2 and T_3 denote the events $W_\mathcal{T} \in [THT]$, $W_\mathcal{T} \in [TTT]$ and $W_\mathcal{T} \in [HHH]$ respectively. Computing our correlation matrix, we get

	T	H	T	T	T	T	H	H	H
THT	2^3		2^1			2^1			
TTT			2^1	2^3	2^2	2^1			
HHH							2^3	2^2	2^1

The revenue earned by each group is thus given by the rows of the following vector:

$$\begin{pmatrix} 10 & 2 & 0 \\ 2 & 14 & 0 \\ 0 & 0 & 14 \end{pmatrix} \begin{pmatrix} \mathbb{P}[T_1] \\ \mathbb{P}[T_2] \\ \mathbb{P}[T_3] \end{pmatrix} = \begin{pmatrix} 10\mathbb{P}[T_1] + 2\mathbb{P}[T_2] \\ 2\mathbb{P}[T_1] + 14\mathbb{P}[T_2] \\ 14\mathbb{P}[T_3] \end{pmatrix}.$$

For reasons entirely analogous to the $|\mathcal{T}| = 2$ case, all three rows must be equal. This gives us two equations: row 1 = row 2 and row 1 = row 3.[3] This gives us the following systems of equations:

$$\begin{cases} \mathbb{P}[T_1] + \mathbb{P}[T_2] + \mathbb{P}[T_3] = 1, \\ \quad 10\mathbb{P}[T_1] + 2\mathbb{P}[T_2] = 2\mathbb{P}[T_1] + 14\mathbb{P}[T_2], \\ \quad 10\mathbb{P}[T_1] + 2\mathbb{P}[T_2] = 14\mathbb{P}[T_3], \end{cases}$$

Solving, we have $\mathbb{P}[T_1] = 21/52$, $\mathbb{P}[T_2] = 14/52$ and $\mathbb{P}[T_3] = 17/52$.[4]

Now that we are fully acquainted with the correlation matrices, we proceed to derive a general formula for $\mathbf{p}_\mathcal{T}$.

Theorem 26. Let $\mathcal{T} = \{t_1, t_2, \dots, t_n\}$. Let $\mathbf{1} = (1, 1, \dots, 1)$ be the $n \times 1$ vector that is all ones. Then

$$\mathbf{p}_\mathcal{T} = \frac{1}{\mathbf{1}^\intercal \mathbf{M}_\mathcal{T}^{-1} \mathbf{1}} \mathbf{M}_\mathcal{T}^{-1} \mathbf{1}.$$

Proof. Let A_i be the ith row of $\mathbf{M}_\mathcal{T} \mathbf{p}_\mathcal{T}$. Since $A_1 = A_2 = \dots A_n = \lambda$, we have

$$\mathbf{M}_\mathcal{T} \mathbf{p}_\mathcal{T} = \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_n \end{pmatrix} = \lambda \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} = \lambda \mathbf{1}.$$

Pre-multiplying by $\mathbf{M}_\mathcal{T}^{-1}$, we obtain

$$\mathbf{p}_\mathcal{T} = \lambda \mathbf{M}_\mathcal{T}^{-1} \mathbf{1}.$$

[4] Note that row 1 = row 2 and row 1 = row 3 implies row 2 = row 3, so we do not need to include row 2 = row 3 in our system.

By the law of total probability, we must have $\mathbf{1}^\mathsf{T}\mathbf{p}_\mathcal{T} = 1$. Pre-multiplying the above equation by $\mathbf{1}^\mathsf{T}$,

$$1 = \lambda \mathbf{1}^\mathsf{T}\mathbf{M}_\mathcal{T}^{-1}\mathbf{1} \implies \lambda = \frac{1}{\mathbf{1}^\mathsf{T}\mathbf{M}_\mathcal{T}^{-1}}.$$

Thus,

$$\mathbf{p}_\mathcal{T} = \frac{1}{\mathbf{1}^\mathsf{T}\mathbf{M}_\mathcal{T}^{-1}\mathbf{1}}\mathbf{M}_\mathcal{T}^{-1}\mathbf{1}.$$

□

Remark. $\mathbf{1}^\mathsf{T}\mathbf{M}_\mathcal{T}^{-1}\mathbf{1}$ is the sum of all entries of the matrix $\mathbf{M}_\mathcal{T}^{-1}$.

This expression allows us to easily calculate the probabilities using software.

4.3 The Expected Length

Using our formula for $\mathbf{p}_\mathcal{T}$, we now aim to find a closed form for $\mathbb{E}[L_\mathcal{T}]$. We will do so using yet another "fair casino" argument.

To build our intuition, we once again consider the case where we have only two terminators, say $\mathcal{T} = \{THT, HTT\}$. As usual, you, as the dealer at D'Casino, keep flipping a coin until either terminator is achieved. This time, D'Casino only offers its original game, which is restated below:

> *Each round, a fair-coin is flipped. Gamblers go all-in, betting on the outcome of this coin-flip. If they win, they double their money, and they play again. If they lose, the casino takes everything and they go home empty-handed. This repeats until a terminator appears, at which point the casino closes and everybody goes home.*

Suppose now that we have two groups of gamblers that frequent D'Casino. Group 1 is obsessed with THT, while Group 2 is obsessed with HTT. Every flip, a new gambler from each group enters D'Casino and bets \$1 on their respective sequence.

To illustrate, suppose the coin-flips come up $HHTHT$. Let $R_n(i)$ and $C_n(i)$ be the total revenue earned and total cost incurred by Group i gamblers at the nth flip. Let

$$X_n = \underbrace{(R_n(1) - C_n(1))}_{\text{Group 1's profit}} + \underbrace{(R_n(2) - C_n(2))}_{\text{Group 2's profit}}$$

be the total profit made by the two groups at the nth flip.

We begin by analysing Group 1's profits.

n	**Event**	$R_n(1)$	$C_n(1)$
1	Gambler #1 loses his first bet.	0	1
2	Gambler #2 loses his first bet.	0	2
3	Gambler #3 wins all three bets. He hence earns \$7 overall. For consistency, we record this as a revenue of \$8 and a loss of \$1.	8	3
4	Gambler #4 loses his first bet.	8	4
5	Gambler #5 wins his first bet before the casino closes. Like before, we record this as a revenue of \$2 and a loss of \$1.	10	5

We now analyse Group 2's profit.

n	**Event**	$R_n(2)$	$C_n(2)$
1	Gambler #1 loses his second bet.	0	1
2	Gambler #2 loses his third bet.	0	2
3	Gambler #3 loses his first bet.	0	3
4	Gambler #4 wins his first and second bet, but the casino closes before his third. We record this as a revenue of \$4 and a loss of \$1.	4	4
5	Gambler #5 loses his first bet.	4	5

We now make the same key observations as we did back in Section 1.3.

- Due to the way we record losses, $C_\tau(i)$ is equal to the number of coin-flips made. Taking expectations
$$\mathbb{E}[C_\tau(i)] = \mathbb{E}[L_\mathcal{T}]. \tag{4.4}$$
- Only the last three gamblers in each group stand a chance to win money.
- Because the casino is fair, the total profit $\{X_n\}$ is a martingale. Hence, by the OST,
$$\mathbb{E}[X_\tau] = \mathbb{E}[X_0] = 0. \tag{4.5}$$

Chaining (4.4) and (4.5), we have

$$\mathbb{E}[R_\tau(1) + R_\tau(2)] = \mathbb{E}[C_\tau(1) + C_\tau(2)] = 2\mathbb{E}[L_\mathcal{T}].$$

By the law of total expectation, one can write this as

$$\begin{aligned} 2\mathbb{E}[L_\mathcal{T}] = \mathbb{E}[R_\tau(1) + R_\tau(2) \mid W_\mathcal{T} \in [THT]]\mathbb{P}[W_\mathcal{T} \in [THT]] \\ + \mathbb{E}[R_\tau(1) + R_\tau(2) \mid W_\mathcal{T} \in [HTT]]\mathbb{P}[W_\mathcal{T} \in [HTT]]. \end{aligned} \tag{4.6}$$

Now consider our correlation matrix:

	T	H	T	H	T	T
THT	2^3		2^1			2^1
HTT		2^2		2^3		

Reading off the first column (which corresponds to the event $W_\mathcal{T} \in [THT]$), we see that

$$\mathbb{E}[R_\tau(1) \mid W_\mathcal{T} \in [THT]] = 2^3 + 2^1 \quad \text{and} \quad \mathbb{E}[R_\tau(2) \mid W_\mathcal{T} \in [THT]] = 2^2.$$

Similarly, looking at the second column, we see that

$$\mathbb{E}[R_\tau(1) \mid W_\mathcal{T} \in [HTT]] = 2^1 \quad \text{and} \quad \mathbb{E}[R_\tau(2) \mid W_\mathcal{T} \in [HTT]] = 2^3.$$

From the previous section, we found that both probabilities were $1/2$. Substituting all of these values back into (4.6), we have

$$2\mathbb{E}[L_\mathcal{T}] = (10+4)\left(\frac{1}{2}\right) + (2+8)\left(\frac{1}{2}\right) = 12,$$

so $\mathbb{E}[L_\mathcal{T}] = 6$.

Generalizing this is quite simple. If we have n terminators $\mathcal{T} = \{t_1, \ldots, t_n\}$, then

$$\mathbb{E}[L_\mathcal{T}] = \sum_{i=1}^{n} \frac{1}{n} \left(\text{sum of } i\text{th column of } \mathbf{M}_\mathcal{T}\right) \mathbb{P}[W_\mathcal{T} \in [t_i]]. \tag{4.7}$$

We can phrase this more neatly using linear algebra:

Theorem 27. Let $\mathcal{T} = \{t_1, \ldots, t_n\}$. Then

$$\mathbb{E}[L_\mathcal{T}] = \frac{1}{\mathbf{1}^\intercal \mathbf{M}_\mathcal{T}^{-1} \mathbf{1}}.$$

Proof. The sum of the ith column is given by the ith entry of $\mathbf{M}_\mathcal{T}^\intercal \mathbf{1}$. (4.7) can hence be written as an inner product:

$$\mathbb{E}[L_\mathcal{T}] = \frac{1}{n} \left(\mathbf{M}_\mathcal{T}^\intercal \mathbf{1}\right)^\intercal \mathbf{p}_\mathcal{T}.$$

Invoking Theorem 26, this simplifies to

$$\mathbb{E}[L_\mathcal{T}] = \frac{1}{n} \left(\mathbf{M}_\mathcal{T}^\intercal \mathbf{1}\right)^\intercal \frac{1}{\mathbf{1}^\intercal \mathbf{M}_\mathcal{T}^{-1} \mathbf{1}} \left(\mathbf{M}_\mathcal{T}^{-1} \mathbf{1}\right) = \frac{1}{n\left(\mathbf{1}^\intercal \mathbf{M}_\mathcal{T}^{-1} \mathbf{1}\right)} \mathbf{1}^\intercal \mathbf{M}_\mathcal{T} \mathbf{M}_\mathcal{T}^{-1} \mathbf{1} = \frac{\mathbf{1}^\intercal \mathbf{1}}{n\left(\mathbf{1}^\intercal \mathbf{M}_\mathcal{T}^{-1} \mathbf{1}\right)}.$$

Since $\mathbf{1}^\intercal \mathbf{1} = n$, we obtain our desired result. □

Let us look at some examples.

Example 28. Let $\mathcal{T} = \{THT, HTT\}$. Our correlation matrix is given by

	T	H	T	H	T	T
THT	2^3		2^1			2^1
HTT		2^2		2^3		

The inverse of $\mathbf{M}_\mathcal{T}$ is

$$\mathbf{M}_\mathcal{T}^{-1} = \frac{1}{36} \begin{pmatrix} 4 & -1 \\ -2 & 5 \end{pmatrix}.$$

Using the above result, we have

$$\mathbb{E}[L_\mathcal{T}] = \frac{1}{\frac{1}{36}(4 - 1 - 2 + 5)} = 6.$$

Example 29. Let $\mathcal{T} = \{TT, THT\}$. Our correlation matrix is given by

	T	T	T	H	T
TT	2^2	2^1			2^1
THT		2^1	2^3		2^1

One can calculate the inverse of $\mathbf{M}_\mathcal{T}$ to be

$$\mathbf{M}_\mathcal{T}^{-1} = \frac{1}{28} \begin{pmatrix} 5 & -1 \\ -1 & 3 \end{pmatrix}.$$

Thus,

$$\mathbb{E}[L_\mathcal{T}] = \frac{1}{\frac{1}{28}(5 - 1 - 1 + 3)} = \frac{14}{3}.$$

5 Further Questions

In this workshop, we managed to derive closed forms for $\mathbb{E}[L_\mathcal{T}]$ and $\mathbb{P}[W_\mathcal{T} \in [t]]$. There are, however, many more questions we can ask about this game:

- Is there a closed form for $\mathbb{E}[L_{\mathcal{T}} \mid W_{\mathcal{T}} \in [t]]$?
- If $W_{\mathcal{T}} \in [t]$, what is the distribution of $L_{\mathcal{T}}$? How many words in $[t]$ have length n? Equivalently, given that $L_{\mathcal{T}} = n$, what is the probability that $W_{\mathcal{T}} \in [t]$?
- Given n terminators, each at most length k, what is the minimum and maximum value of $\mathbb{E}[L_{\mathcal{T}}]$?
- What is the significance of $\mathbf{M}_{\mathcal{T}}^{-1}$? What does it mean to "invert" a correlation matrix?

Slightly modifying our original coin-flip problem also opens up a whole can of worms:

- What if we stopped flipping the coin once we see all terminators?
- What if we allowed up to k appearances of a single terminator?
- If we flip a fair coin n times, what is the probability that a terminator t appears?

References

[1] L.J Guibas and A.M Odlyzko. String overlaps, pattern matching, and nontransitive games. *Journal of Combinatorial Theory, Series A*, 30(2):183–208, 1981.

[2] Vikas Jain. Why is the expected number coin tosses to get HTH is 10? Mathematics Stack Exchange.

[3] Aaron Montgomery. Slick way to solve P(pattern A appears before pattern B) type of problems in a coin toss process. Mathematics Stack Exchange.

[4] Scott M. LaLonde. The Martingale Stopping Theorem. https://math.dartmouth.edu/~pw/math100w13/lalonde.pdf, 2013.

[5] Adam Lelkes. Martingales and the Optional Stopping Theorem. https://doi.org/10.59350/hgcyv-5x530, 2014.

[6] Mihai Nica. A Coin Flip Paradox and the ABRACADABRA Theorem for infinite monkeys: How long does it take? https://www.youtube.com/watch?v=t8xqMxlZz9Y, 2022.

[7] John Sylvester, Nicolás Rivera, Luca Zanetti, and Thomas Sauerwald. Lecture 8: The Optional Stopping Theorem. https://www.cl.cam.ac.uk/teaching/1819/Probablty/materials/Lecture8.pdf, 2019.