A Coin Flipping Game with Non-Transitive Odds

Stacy Jurgens

Abstract: Two players choose distinct sequences of Heads and Tails three entries long. A coin is then tossed and the outcome, Head or Tail, is recorded. The first player to observe his/her sequence occur is the victor. We will be examining the probabilities that one player wins over another depending on the sequences chosen. Fair and biased coins will be considered. The properties of the probabilities will then be examined. If the probability that sequence A occurs before sequence B is greater than 1/2, we say that sequence A "beats" sequence B. We will show that a biased coin (P(H)=2/5) yields the probabilistic result that the sequence TTT beats TTH beats THT, but TTT does not beat THT. Thus, the odds of the game are not necessarily transitive. We will examine this and other scenarios to find out why and when the odds of victory are non-transitive.

Preface

At the beginning of this project, my honors thesis, I was unsure of what was expected of me and what writing a math paper really involved. I started out looking at cellular automata, not coin tossing. I had been interested in probability and its calculations since my sophomore probability class, and wanted to do a project in that area. As I studied and experimented in the field of probabilistic cellular automata, I found that I was in over my head. The subject of cellular automata is vast, and I was having a hard time focusing my efforts on any one part. In particular, I found myself losing interest in the topic, thus slowing down my progress considerably. That is when I decided I had to change focus. I still wanted to work with probability because I enjoyed the calculations involved. But I wanted a subject that I could experiment in and understand and interpret my results. When my advisor gave me Arthur Engel's papers to read, I knew that I could use his method above and beyond what his papers narrated and perhaps find my own interesting results. That is exactly what I did. However, as I neared the end of my empirical research and began writing up my results, I found that other mathematicians had already done much of my work, [3], [4],[5]. Since I had developed my results independently, I continued preparing my thesis for submission, but also took the time to familiarize myself with the other research that had been done. In the end, I was satisfied with my work, but still in awe of the accomplishments of others and intrigued by the amount of questions still left unanswered. I present to you now my research, complete in its methods and purpose, results and questions.

Introduction

Coin tosses are as common today as the coins that are tossed. They are used to determine possession in football games, starters in competitions, and who gets to pick what movie to see on a Saturday night. Every probability textbook includes homework problems and examples dealing with the tossing of a coin and its outcome. We have examined a game that involves coin tossing, and have kept the rules simple. Anyone can play, and it does not take a Mathematics degree to understand the basic numbers involved. While this may sound simple, it is extremely interesting as we make the game more complex and the coin being tossed less fair. Let us begin.

Section 1: Playing the Game

To start we will look at the simplest of the game's variations. It requires two players but only one coin. For the purpose of this paper, we will assume the coin has a Head and a Tail side. Each player must choose a distinct sequence of Heads and Tails, each three entries long. An example would be Player 1 choosing HHH and Player 2 choosing HTH. The coin is then tossed and the outcome recorded until one of the sequences chosen occurs. Henceforth, the length of the game runs anywhere from 3 to an infinite number of tosses. Since there are only two players, the sequences must be distinct, and each must be exactly three entries long, the total number of scenarios is easy to compute. First of all, the total number of possible sequences from which each player can choose is:

$$(\text{number of possible outcomes})^{(\text{the number of entries in the sequence})} = 2^3 = 8.$$

Since the players must choose distinct sequences, that leaves the following number of total different scenarios: $8 \times 8 - 8 = 56$.

Now we will bring probabilities into the game. Using the method described in detail further on, the probability that sequence A occurs before sequence B is computed for each of the scenarios. **If the probability that sequence A occurs before sequence B is greater than 1/2, we say that sequence A "beats" sequence B.** In this way we can determine which sequences have better odds of winning. It is these odds that hold our interest.

Before describing the method used to determine the probabilities of each sequence winning, we must clarify a few definitions. **The term "win" in this paper refers to one sequence occurring before another**. This is different from the term "beats" which refers to the property that a sequence wins over another with probability greater than 1/2.

It is also important to point out some short cuts that can be taken due to the nature of the game. Probabilistically, one of the players wins every time. Thus out of two chosen sequences, **if one sequence wins with probability p then the other wins with**

probability 1-p. We will refer to this property as the "**One Victor Rule**". This relationship allows for the reduction of 56/2 = 28 computations, leaving only 28 to still be examined. Even more, 4 of these 28 remaining scenarios involve two sequences that differ only by the last entry. The probability of either sequence winning in these situations is simply the probability of the last entry occurring on each toss. For example, if a fair coin is being used (that is P(H)= P(T)=1/2) then the probability that HTH occurs before HTT is 1/2, the probability that the last toss will be a Head. A chart of the different sequences and their probability that the column wins over the row using a fair coin is given below. Note the symmetry occurring because if p= 1/2, then 1-p= 1/2 as well.

Chart 1.

	HHH	**HHT**	**HTH**	**HTT**	**THH**	**THT**	**TTH**	**TTT**
HHH		**1/2**	3/5	3/5	7/8	7/12	7/10	1/2
HHT	**1/2**		1/3	1/3	3/4	3/8	1/2	3/10
HTH	*2/5*	*2/3*		**1/2**	1/2	1/2	5/8	5/12
HTT	*2/5*	*2/3*	**1/2**		1/2	1/2	1/4	1/8
THH	*1/8*	*1/4*	*1/2*	*1/2*		**1/2**	2/3	2/5
THT	*5/12*	*5/8*	*1/2*	*1/2*	**1/2**		2/3	2/5
TTH	*3/10*	*1/2*	*3/8*	*3/4*	*1/3*	*1/3*		**1/2**
TTT	*1/2*	*7/10*	*7/12*	*7/8*	*3/5*	*3/5*	**1/2**	

The eight entries in bold print are those that involve sequences differing only by the last entry. Note that either the top or the bottom of the chart can be examined due to the "One Victor Rule".

All of these reductions result in 24 scenarios that must be more carefully examined. We will look at the bottom corner of the chart (the italicized entries) using the Probabilistic Abacus Method.

Section 2: The Probabilistic Method

The Probabilistic Abacus Method, developed by Arthur Engel [1], requires no more mathematical knowledge than basic algebra and the ability to count. It is after the probabilities have been computed that the more interesting results began to appear. The method is most easily described using a specific example, so we will introduce it by determining the probability that HHH occurs before HTH (that is, the probability that HHH wins over HTH) when a fair coin is tossed.

The first (and perhaps most difficult) step is to draw a diagram of all the possible outcomes of the flipping coin, keeping in mind that the coin will stop flipping only when one of the chosen sequences has occurred. For the case of HHH versus HTH, the diagram is as such:

Figure 1.

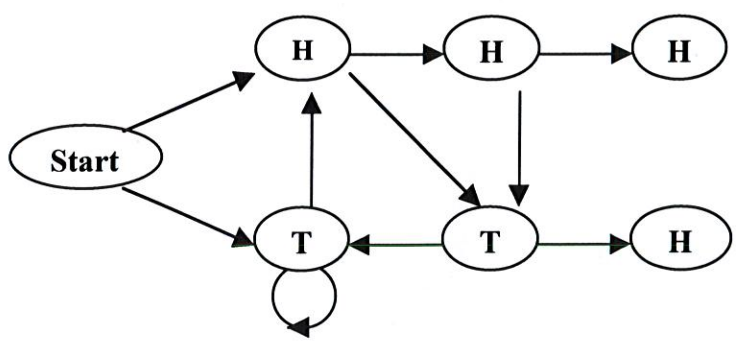

From the starting point the coin is tossed once, and its outcome is either a Head or a Tail, indicated by the two arrows going from "start" to "H" and "T". With each toss, there are two possible outcomes, thus each circle must have two arrows (indicating the two possible outcomes) originating from it. Of these two arrows, one must go to a "H" position and the other to a "T". The only exceptions to this rule are the circles representing the outcome of the last coin toss. At this point, one of the sequences has occurred and the game is over, the coin will not be tossed anymore, thus there will be no arrows originating from these circles. These rules are little more than basic observations, but they are very useful when constructing the diagrams. See Appendix A for a complete list of game diagrams involving sequences of length three.

Once the diagram is drawn, the problem is simply an algebraic one. Each arrow represents a coin toss while its destination represents a possible outcome. With a fair coin, the probability of each toss being a Head or Tail is 1/2, thus each arrow represents a path taken with probability 1/2.

Figure 2.

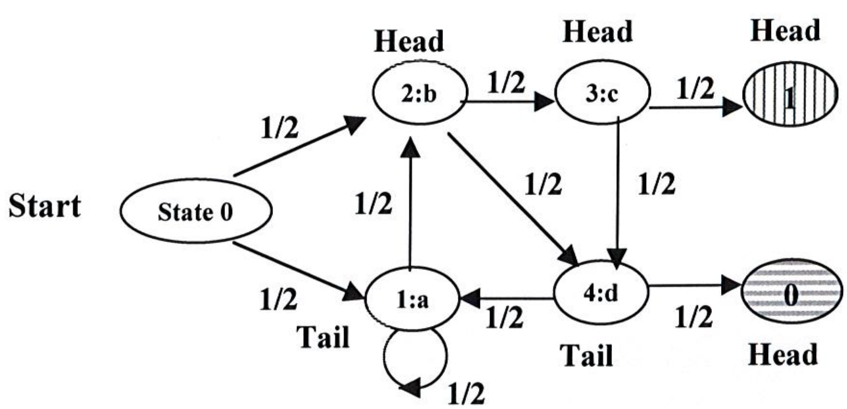

Each circle represents a state in the game. For example the start circle represents the game before any tosses have occurred, thus once a toss is made it is impossible to return to the start position. The expressions inside each of the intermediate circles represent the state's label (0,1,2,3, or 4) and the probability of HHH winning from that state (a,b,c, or d). (Note that the probability of winning from state 0 does not have a variable. The purpose of this exercise is to find that probability, and we will be defining it in terms of the other variables.) The two colored states are "absorbing" states, once you have reached one of them you cannot escape and the game is over. The circle with vertical stripes in Figure 2 represents one of the two possible results of the game, the situation where the HHH sequence occurs before the HTH (thus HHH wins). We will find the probability that HHH occurs before HTH, making the probability the HHH wins from the vertically striped state 1. The probability that HHH wins from the horizontally striped state is 0, since HTH has won at that state.

What we are looking for is the probability *from the start* that HHH occurs before HTH. To find this we must determine the intermediate probabilities in the remaining four states, that is determine a,b,c, and d.

Now, at any state on the diagram in Figure 2 (with the exception of the two absorbing states, whose probabilities are already known) there are two situations that can occur. Either a Head or a Tail will be tossed and the game will continue on appropriately. Thus there are two ways that victory can be obtained from each of these states, either by flipping a Head and carrying on or by flipping a Tail and continuing.

Recall that H= the event we toss a Head and T= the event we toss a Tail. Define W_n= the event that HHH wins from state n ($n \in \{0,1,2,3,4\}$, for the remainder of this paper the subscript n will refer to this indexing set). Recall also that a,b,c, and d are defined as the probability of HHH winning from a given state on the diagram, for example $P(W_1)$=a. Then generically the probability of HHH winning from any of the intermediate states can be thought of as follows:

(1) P(W)=P(winning and getting a Head or winning and getting a Tail).

Since the events of getting a Head or a Tail are mutually exclusive, the following holds:

(2) $$P(W) = P((W \cap H) \cup (W \cap T)) = P(W \cap H) + P(W \cap T) = P(W|H)P(H) + P(W|T)P(T).$$

$P(W|H)$ is the probability of winning from the state *traveled to* upon flipping a Head, likewise $P(W|T)$ is the probability of winning from the state *traveled to* upon flipping a Tail. Also, since we are using a fair coin in this example, P(H)=P(T)=1/2. Thus the probability of HHH winning from the starting state (and thus the probability that HHH wins) in Figure 2 is:

(3) $P(W_0) = a \times 1/2 + b \times 1/2 = a/2 + b/2.$

Repeating this process to find $P(W_n)$ for each of the intermediate states yields:

(4) $P(W_1) = a = a \times 1/2 + b \times 1/2$

(5) $P(W_2) = b = c \times 1/2 + d \times 1/2$

(6) $P(W_3) = c = d \times 1/2 + 1 \times 1/2$

(7) $P(W_4) = d = a \times 1/2 + 0 \times 1/2$

Solving this system of linear equations algebraically yields a=2/5, b=2/5, c=7/10, and d=1/5. Thus, by substituting these values back into (3) we get $P(W_0)$=2/5. Thus the probability that HHH wins over HTH is 2/5. Thus, since 2/5 < 1/2, HHH does not "beat" HTH. The technique is quite simple once you get the hang of it, and the remaining 23 situations are quickly analyzed, resulting in Chart 1.

Section 3: The Biased Coin

Looking at the same example, P(HHH wins over HTH), only with a biased coin, requires only a small adjustment in our diagram and calculations. The shape of the diagram remains the same, what has changed is the probability that one travels from one state to another. Looking at a generic coin, P(H)=p and P(T)=1-p, thus the diagram is as follows:

Figure 3.

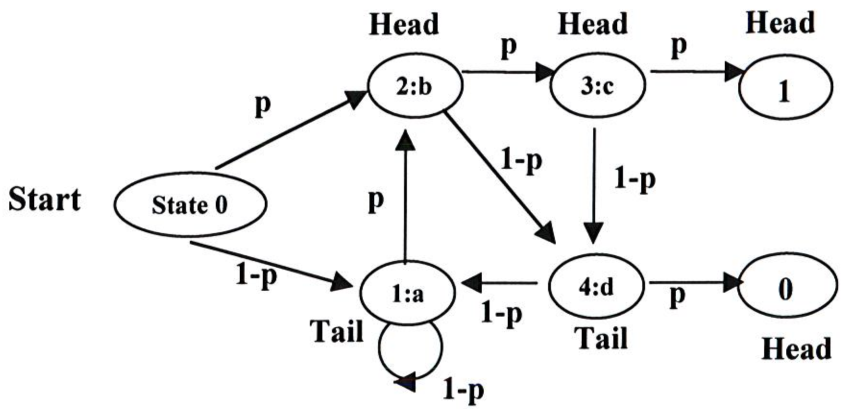

The same technique used with the fair coin is now used to determine the values of a,b,c, and d.

(8) $P(W_0) = a \times (1-p) + b \times p$

(9) $P(W_1) = a = a \times (1-p) + b \times p$

(10) $P(W_2) = b = c \times p + d \times (1-p)$

(11) $P(W_3) = c = 1 \times p + d \times (1-p)$

(12) $P(W_4) = d = 0 \times p + a \times (1-p)$

Solving the system of equations in (9)-(12) yields, $a = b = -p/(p^2 - p - 1)$.

Plugging these values back into (8) gives the following result:

(13) P(HHH wins over HTH)= $p/(1 + p - p^2)$.

The preceding was a step by step example of how to determine the probability that one sequence occurs before another, using both a fair and a biased coin. Note that in this specific example $P(W_0)=P(W_1)=P(W_2)$. This is the case when both sequences begin with the same outcome, H or T. Thus neither can began its path to victory until that first outcome is obtained.

Now that we have found the probability that HHH occurs before HTH, finding the probability that HTH wins over HHH is easy due to "One Victor Rule".

(14) P(HTH wins over HHH)=1-P(HHH wins over HTH)= $(1-p^2)/(1+p-p^2)$.

Repeating this method for the remaining 23 scenarios and then applying the "One Victor Rule" results in Chart 2. Recall that the values in the chart are the probability that the column wins over (occurs before) the row.

Chart 2.

	HHH	HHT	HTH	HTT	THH	THT	TTH	TTT
HHH		$1-p$	$\frac{1-p^2}{1+p-p^2}$	$\frac{(p-1)^2(p+1)}{p^3-p+1}$	$1-p^3$	$\frac{(p-1)^2(p^2+p++1)}{p^2-p+1}$	$\frac{(p-1)^2(p^2+p+1)}{p^3-p+1}$	$\frac{(p-1)^3(p^2+p+1)}{p^4-2p^3-p^2+2p-1}$
HHT	p		$\frac{p-1}{p-2}$	$\frac{p-1}{p-2}$	$1-p^2$	$(p-1)^2(p+1)$	$\frac{(p-1)^2(p+1)}{p^2-p+1}$	$\frac{(p-1)^3(p+1)}{p^3-3p^2+2p-1}$
HTH	$\frac{p}{1+p-p^2}$	$\frac{-1}{p-2}$		$1-p$	$\frac{p^2-1}{p-2}$	$\frac{(p-1)^2(p+1)}{p^2-p+1}$	$1-2p^2+p^3$	$\frac{-p^4+3p^3-2p^2-p+1}{p^2-p+1}$
HTT	$\frac{p^2}{p^3-p+1}$	$\frac{-1}{p-2}$	p		p	$\frac{p^2-p+1}{p+1}$	$(p-1)^2$	$-(p-1)^3$
THH	p^3	p^2	$\frac{p-p^2-1}{p-2}$	$1-p$		$1-p$	$\frac{1-p}{p^2-p+1}$	$\frac{-(p-1)^2}{p^3-3p^2+2p-1}$
THT	$\frac{-p^2(p^2-p-1)}{p^2-p+1}$	$p(1+p-p^2)$	$\frac{p^2(2-p)}{p^2-p+1}$	$\frac{-p(p-2)}{p+1}$	p		$\frac{1}{p+1}$	$\frac{p-1}{p^2-p-1}$
TTH	$\frac{-p^3(p-2)}{p^3-p+1}$	$\frac{-p^2(p-2)}{p^2-p+1}$	$p^2(2-p)$	$p(2-p)$	$\frac{p^2}{p^2-p+1}$	$\frac{p}{p+1}$		$1-p$
TTT	$\frac{-p^3(p^2-3p+3)}{p^4-2p^3-p^2+2p-1}$	$\frac{-p^2(p^2-3p+3)}{p^3-3p^2+2p-1}$	$\frac{p^2(p^2-3p+3)}{p^2-p+1}$	$p(p^2-3p+3)$	$\frac{p^2(p-2)}{p^3-3p^2+2p-1}$	$\frac{p(p-2)}{p^2-p-1}$	p	

Section 4: Results

Now that all of the work is finished we look at the interesting implications of it. By entering Chart 2 into a Microsoft Excel program we were able to look at many different results using different biased coins. In trying different values for p=P(H), we came upon the following lemma:

Lemma 1. If P(H)=2/5, then the odds of a sequence beating another are not transitive.

Proof: Let P(H)=2/5.

Then the probabilities of column winning over the row are as follows:

Chart 3.

	HHH	HHT	HTH	HTT	THH	THT	TTH	TTT
HHH		3/5	21/31	63/83	117/125	351/475	351/415	137/188
HHT	2/5		3/8	3/8	21/25	63/125	63/95	27/55
HTH	10/31	5/8		3/5	21/40	63/95	93/125	279/475
HTT	20/83	5/8	2/5		2/5	19/35	9/25	27/125
THH	8/125	4/25	19/40	3/5		3/5	15/19	45/77
THT	124/475	62/125	32/95	16/35	2/5		5/7	15/31
TTH	64/415	32/95	32/125	16/25	4/19	2/7		3/5
TTT	51/188	28/55	196/475	98/125	32/77	16/31	2/5	

Note that:

(15) P(TTT wins over TTH)=3/5>1/2$\Rightarrow$TTT beats TTH

(16) P(TTH wins over THT)=5/7>1/2$\Rightarrow$TTH beats THT

(17) P(TTT wins over THT)=15/31<1/2$\Rightarrow$TTT does not beat THT

Thus although TTT beats TTH beats THT, TTT does not beat THT. The odds of a sequence beating another are not transitive for p=P(H)=2/5.

A reasonable question to ask at this point is: For what values of p=P(H) are the odds of victory transitive? In order to determine this, all possible trios of sequences must be examined, a daunting task. We will start by looking only at the trio TTT, TTH, THT.

Example 1.

First we add the stipulation that TTT beats TTH beats THT. Now we must find a value of p=P(H) such that TTT does not beat THT, or P(TTT wins over THT)<1/2.

Using the results in Chart 2:

(18) TTT beats TTH$\Rightarrow$P(TTT wins over TTH)>1/2$\Rightarrow 1-p>1/2\Rightarrow p<1/2$

(19) TTH beats THT$\Rightarrow$P(TTH wins over THT)>1/2$\Rightarrow 1/(p+1)>1/2\Rightarrow$p<1

Assuming TTT does not beat THT$\Rightarrow$P(TTT wins over THT)<1/2,

(20) $$\frac{p-1}{p^2-p-1}<\frac{1}{2}\Rightarrow 2p-2>p^2-p-1\Rightarrow 0>p^2-3p+1\Rightarrow p>\frac{3-\sqrt{5}}{2}$$

Therefore $\frac{3-\sqrt{5}}{2}<p<\frac{1}{2}$ $\Rightarrow$TTT beats TTH beats THT and TTT does not beat THT.

Example 2.

We examine the case where HHT beats HTT beats THH but HHT does not beat THH.

(21) HHT beats HTT$\Rightarrow$P(HHT wins over HTT)>1/2$\Rightarrow p>0$

(22) HTT beats THH$\Rightarrow$P(HTT wins over THH)>1/2$\Rightarrow p<1/2$

(23) HHT does not beat THH$\Rightarrow$P(HHT wins over THH)<1/2$\Rightarrow p^2<1/2\Rightarrow p<\sqrt{2}/2$

Since $1/2<\sqrt{2}/2$, $0<p<1/2$ satisfies the three equations, making this case non-transitive. It is of interest to note here that if p=1/2, then the P(HTT wins over THH)=1/2, which means that neither sequence has higher odds of winning than the other.

Example 3.

Looking at a longer chain of sequences, we examine the case where HHT beats HTT beats TTH beats THH but HHT does not beat THH.

(24) HHT beats HTT$\Rightarrow$P(HHT wins over HTT) >1/2$\Rightarrow p>0$

(25) HTT beats TTH$\Rightarrow$P(HTT wins over TTH)>1/2$\Rightarrow \frac{2-\sqrt{2}}{2}<p$

(26) TTH beats THH$\Rightarrow$P(TTH wins over THH)>1/2$\Rightarrow p<\frac{\sqrt{5}-1}{2}$

(27) HHT does not beat THH$\Rightarrow$P(HHT wins over THH)<1/2$\Rightarrow p<\sqrt{2}/2$

Since $\frac{\sqrt{5}-1}{2} < \frac{\sqrt{2}}{2}$, $\frac{2-\sqrt{2}}{2} < p < \frac{\sqrt{5}-1}{2}$ satisfies the four equations, making this case non-transitive.

Section 5: Questions

When the writer began her research, the Probabilistic Abacus method was the main focus. The method is extremely simple yet incredibly powerful, and appeals to the visual person with its game diagrams. The mathematics involved in running the method are simple enough to be handled by grade-schoolers as they begin to study probability and coins, making it an attractive teaching tool. As the method was used repeatedly, numbers were acquired that led to the above observations and examples. At this point in the research, many questions still remain. Why does the method itself work? It is known that the Mathematics functioning behind the Probabilistic Abacus involve chip-firing and absorbing Markov Chains, see Engel [2]. Then there are questions specific to the game itself. Is there a value range for p such that the odds of winning in a game are always transitive (regardless of the sequences chosen)? And why does the non-transitivity occur?

Section 6: Extensions

Now that we have examined the Probabilistic Abacus method and its application to the coin-tossing game involving sequences of length three, it is worth noting that there are many other games that can be studied using this method. The main requirement a game must have in order to fit into the Probabilistic Abacus is that states in the game can be visited more than once, resulting in infinitely many outcomes (or "tosses") yet still comprehendible because of the recurrence of states. Some examples of other games that can be studied follow.

Example 1.

Our example of a coin being tossed and players choosing distinct sequences of three possible outcomes can be extended to games where the sequences are more than three entries long. Of course, the longer the sequence the more involved the game diagrams and computations. Still, analyzing the game remains simple with the use of the Probabilistic Abacus Method. A specific example would be the sequence HTHT versus HTTH, which has the following game diagram.

Figure 4.

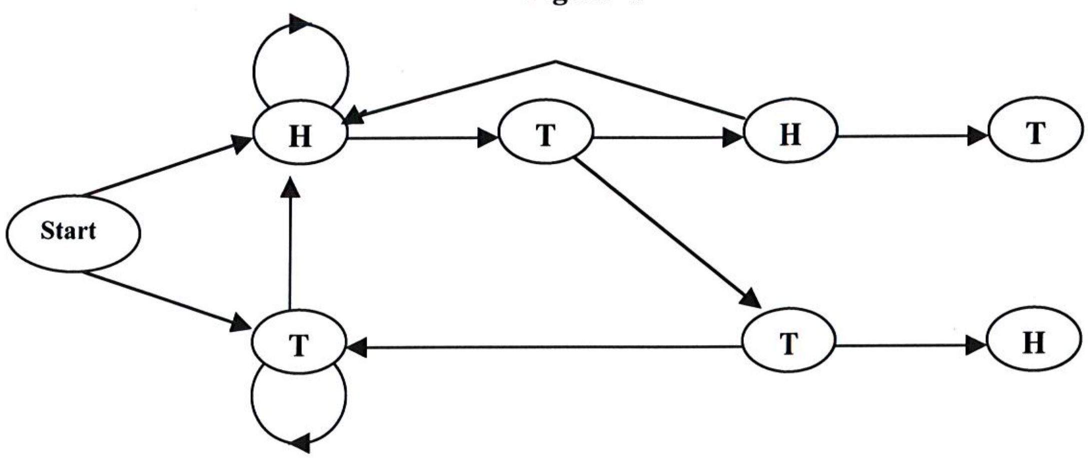

Example 2.

Another prop that yields interesting and exciting games is a die. Games played with a die (as in a pair of dice) are more complex because of the six-possible outcomes. To begin looking at the probabilities involved in rolling a die it may help to categorize the possible outcomes, such as evens and odds. Then sequences can be chosen regarding which category the outcome falls in, such as the three-entry sequence Even-Even-Odd. Other possible games include rolling a pair of dice and observing the sum, difference, product, or quotient of their faces. To add even more complexity consider dice with more than six faces. For example, the board game Scattegories© contains a 26-sided die with the letters of the alphabet on its faces. While time consuming, it is possible that the reader may find the probability that his/her name will be rolled before that of Elvis Presley's.

There are countless other games that can be played using a prop that yields a finite number of outcomes, we have only discussed the more common ones, coins and dice. And with each prop there are numerous ways to choose sequences and patterns when determining the rules for a game.

Conclusion

The Probabilistic Abacus Method is a very nice tool in examining the probabilities associated with infinite patterns of a finite number of outcomes. While the method itself is powerful both in its simplicity and reasoning, the probabilities extracted from it also yield interesting results. And what's the probability of that?

References

[1] Arthur Engel, *The probabilistic abacus*, Educational Studies in Mathematics **6** (1975), 1-22.

[2] Arthur Engel, *Why does the probabilistic abacus work?*, Educational Studies in Mathematics **7** (1976), 59-69.

[3] Shuo-Yen Robert Li, *A martingale approach to the study of occurrence of sequence patterns in repeated experiments*. Ann. Probab. 8 (1980), no. 6, 1171-1176.

[4] Martin Gardner, *Martin Gardner's Sixth Book of Mathematical Games from Scientific American*, W. H. Freeman and Company, San Francisco, 1971.

[5] Martin Gardner, *Wheels, Life, and Other Mathematical Amusements*, W. H. Freeman and Company, New York, 1983.

Appendix A: Game Diagrams for Sequences of Length Three

HHH versus HHT

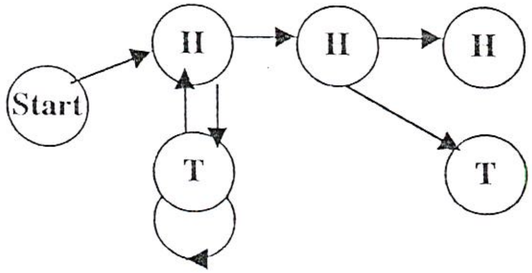

HHH versus THH

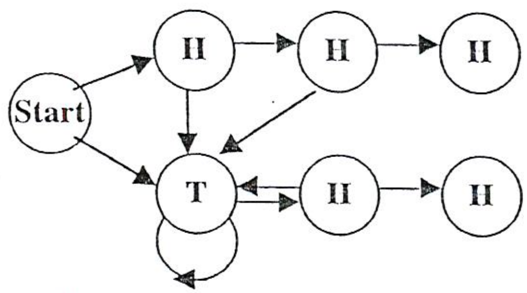

HHH versus TTT

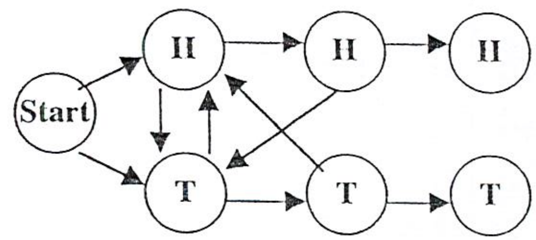

HHH versus HTH

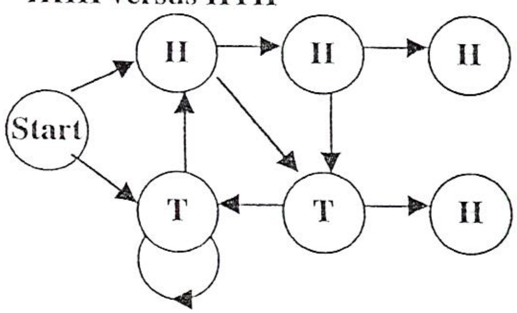

HHH versus THT

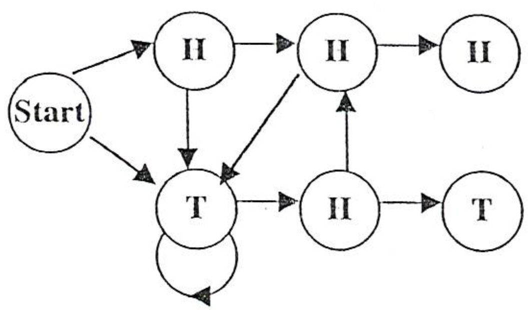

HHT versus HTH

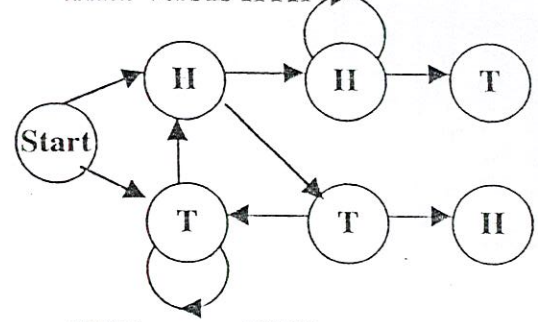

HHH versus HTT

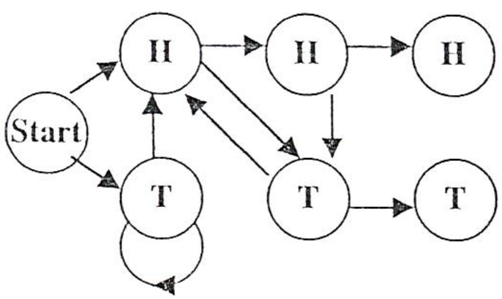

HHH versus TTH

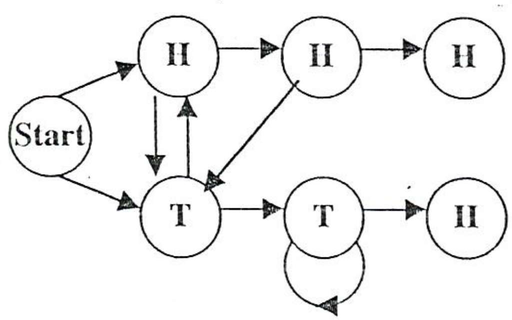

HHT versus HTT

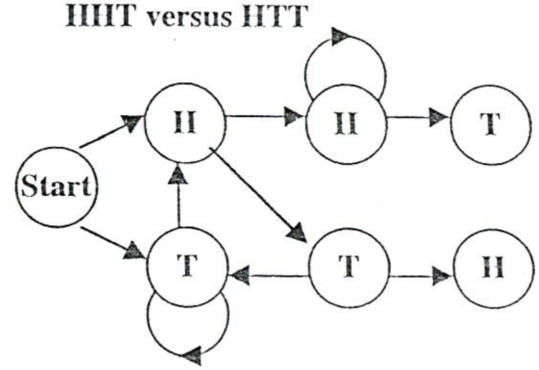

HHT versus THH

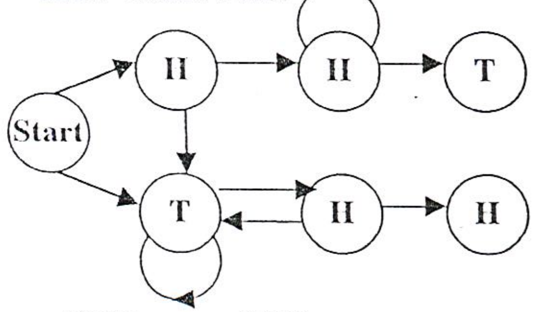

HHT versus TTT

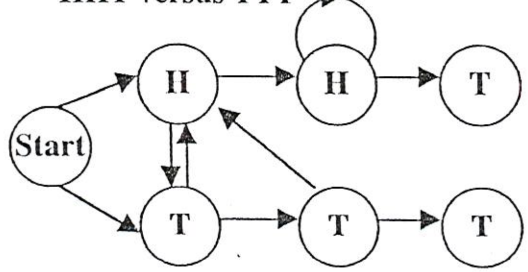

HTH versus THT

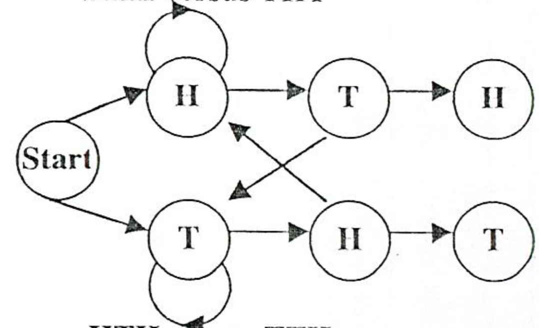

HHT versus THT

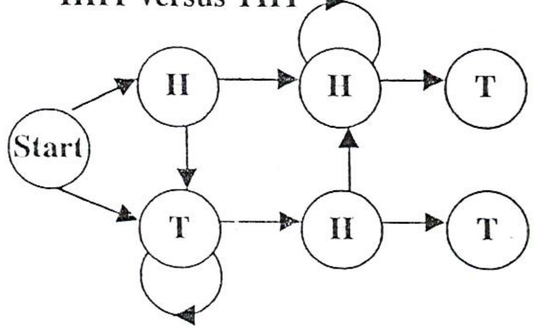

HTH versus HTT

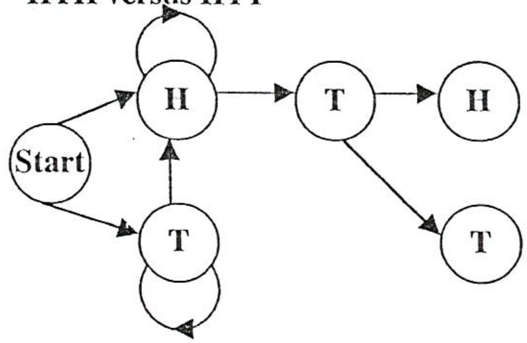

HTH versus TTH

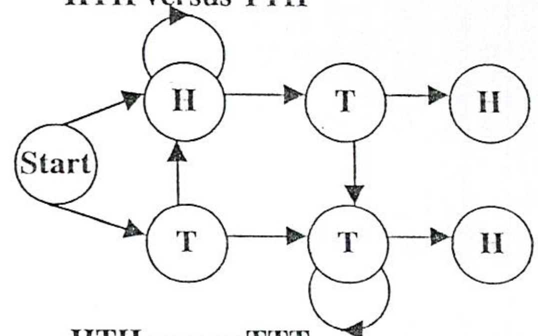

HHT versus TTH

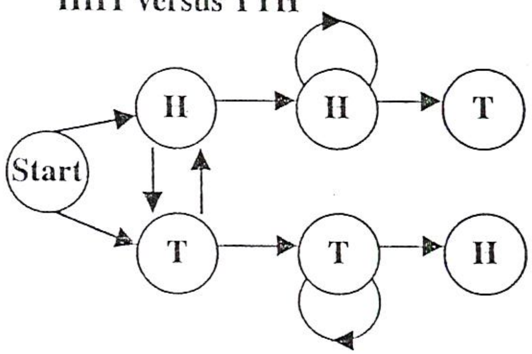

HTH versus THH

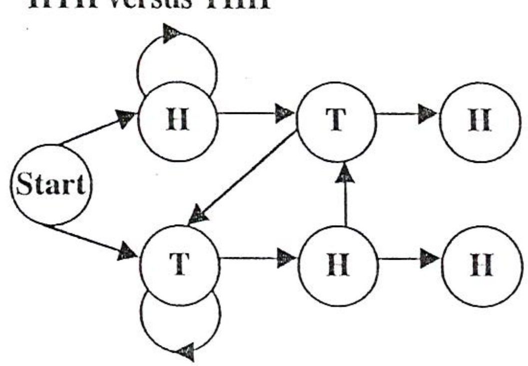

HTH versus TTT

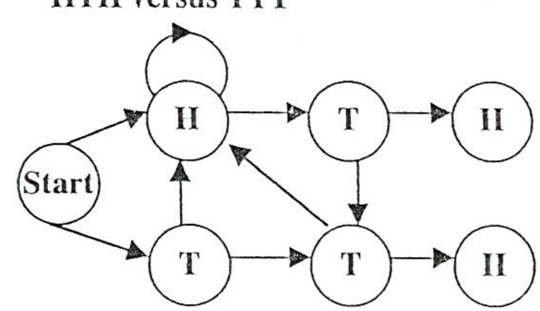

HTT versus THH

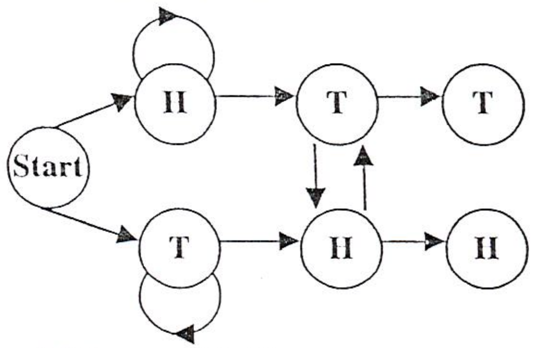

HTT versus TTT

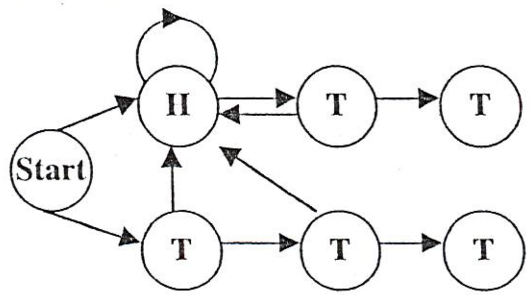

THH versus TTT

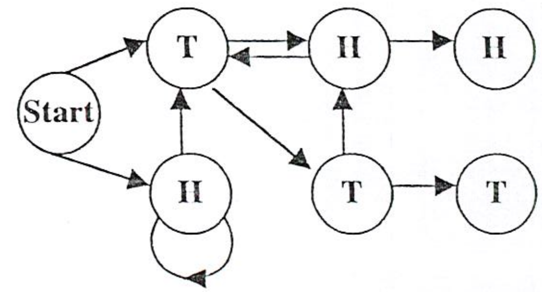

HTT versus THT

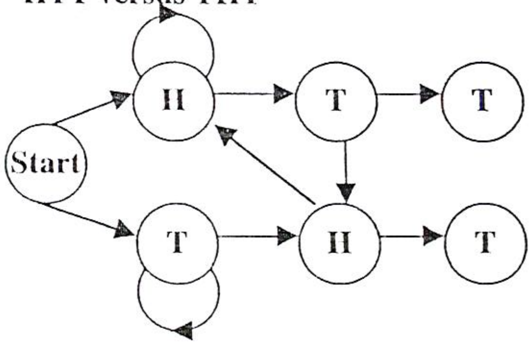

THH versus THT

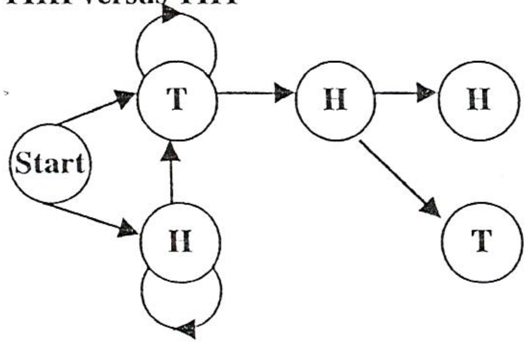

THT versus TTH

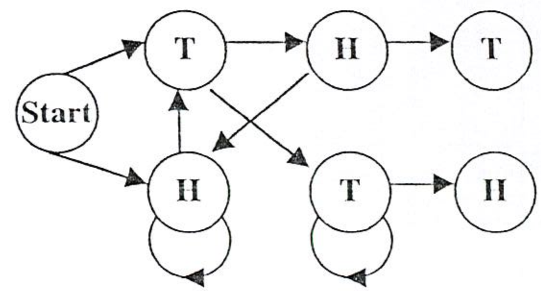

HTT versus TTH

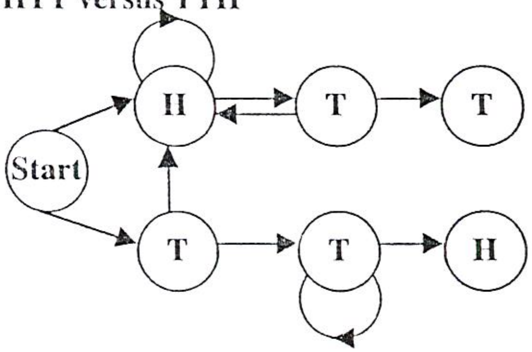

THH versus TTH

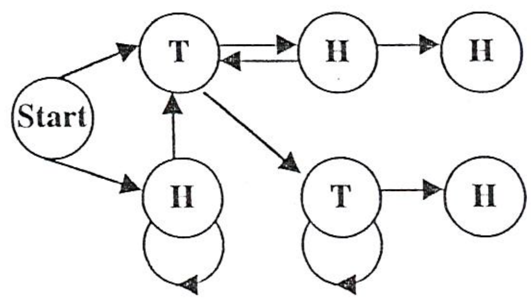

THT versus TTT

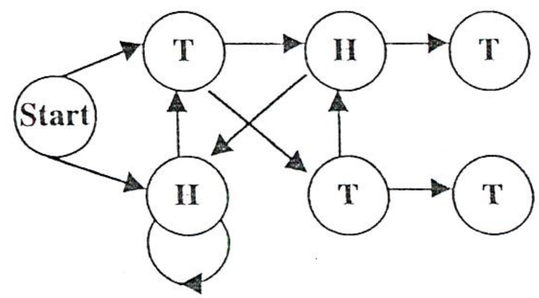

TTH versus TTT