How to Maximize Your Score in Solitaire Yahtzee*

Tom Verhoeff†

June 1999 INCOMPLETE DRAFT: Do not distribute!

Abstract

Yahtzee is a well-known game played with five dice. Players take turns at assembling and scoring dice patterns. The player with the highest score wins. Solitaire Yahtzee is a single-player version of Yahtzee aimed at maximizing one's score. A strategy for playing Yahtzee determines which choice to make in each situation of the game. We show that the maximum expected score over all Solitaire Yahtzee strategies is 254.5896. . . .

1 Introduction

Yahtzee is a dice game where players take turns at assembling and scoring dice patterns so as to obtain a higher score than their opponents. The precise rules of Yahtzee are explained in Section 2. At my home we have simplified the game for my three-year-old daughter, by restricting ourselves to the patterns of, what is called, the Upper Section. Even in this simplified game, there are situations where the innocent question "Daddy, what is my best choice?" cannot be answered directly. As a mathematical parent, I naturally wondered how to compute the best choice. That is how my obsession started.

My first stab at obtaining a better understanding of the simplified game, was triggered by the need to put together a 'design-directed' programming project for my first-year computing science students. Of course, I could not expect them to construct an optimal strategy. However, they should be able to capture their 'favorite' strategy in a programming language and to assess its performance by carrying out some simulations. From my own experiments in this direction, it became obvious that Yahtzee is indeed a game of skill and not just a game of chance. This made me even more eager to determine the optimal choices.

Yahtzee is a registered trademark of the Milton Bradley Company.

†Eindhoven University of Technology, Faculty of Mathematics and Computing Science, e-mail: `T.Verhoeff@tue.nl`

Overview

Section 2 defines the concepts and rules of Yahtzee. Even if you know the game, it is good to browse this section, because there are several variants that I distinghuish. My optimization effort is aimed at the official rules, which, it seems, are not that well known. In Section 3, I analyze the game and introduce such auxiliary notions as a strategy and its expected score. I explain how I have computed the maximum expected score taken over all strategies for Official Solitaire Yahtzee in Section 4. It should be noted that such an optimal strategy is very complicated and hard to learn by human players. Section 5 presents some characteristics of optimal strategies. These can help you improve your decisions. Remaining challenges appear in Section 6.

2 Defining Yahtzee

There are various ways to play *Yahtzee*. This article mainly deals with *Solitaire Yahtzee*, which is played alone. Unless otherwise stated, Yahtzee here refers to Solitaire Yahtzee.

First I define *Abstract Yahtzee* (§2.1), which captures the general rules of almost any kind of Yahtzee. Next I define two instances of Abstract Yahtzee, which I call *Junior Yahtzee* (§2.2) and *Basic Yahtzee* (§2.3). *Official Yahtzee* (§2.4) is the game I am interested in solving. It involves some complications and can be viewed as an extension of Abstract Yahtzee. In §2.5, I briefly discuss *Group Yahtzee*, which is played by two or more players. Finally, some other variants of Yahtzee are described in §2.6.

2.1 Abstract Yahtzee

A game of Yahtzee requires some **dice** (usually *five*) and a **score card**. On the score card is a list of **dice patterns**. Next to each pattern is a **box** in which to enter a **score**. Initially all boxes are **empty**. A game of Yahtzee consists of as many **turns** as there are patterns on the score card.

Each turn consists of an **assembly phase** followed by a **scoring phase**. The assembly phase is started by rolling all dice. Next come some **modification attempts** (usually *two*). In each modification attempt, the player chooses a subset of the dice (the **keepers**) and re-rolls the other dice. In the scoring phase, the player chooses a pattern on the score card with an *empty* box, and enters a score in it. The score is determined by the **pattern scoring rules** based on the final roll of the assembly phase, and the chosen pattern.

When all boxes have been filled with scores, the game ends. The **final score** is determined by the **final scoring rules**, based on the final state of the score card.

Here are some notes:

1. In a modification attempt, you may choose *any* set of keepers, ranging from

all dice to none. Keeping *all* dice boils down to skipping the modification attempt.

2. The two modification attempts are *independent*. That is, there is *no* restriction on the second set of keepers, e.g. it need not be a superset of the first set.

3. You may choose *any* dice pattern for scoring, as long as its box is empty. The resulting score can however be *zero*.

4. After scoring a dice pattern, its box is nonempty, possibly containing a *zero* score.

2.2 Junior Yahtzee

Junior Yahtzee is an instance of Abstract Yahtzee. It involves the six dice patterns and corresponding pattern scoring rules in Table 1. The score of a dice roll for a pattern equals the total value of all dice showing the pattern value.

Pattern	**Score**
Aces	total value of all ones
Twos	total value of all twos
Threes	total value of all threes
Fours	total value of all fours
Fives	total value of all fives
Sixes	total value of all sixes

Table 1: Dice patterns and their scoring rules for Junior Yahtzee

I distinghuish two variants of final scoring rules for Junior Yahtzee:

Without Bonus The final score is the total of all pattern scores.

With Bonus The final score is the total of all pattern scores, plus 35 (the **bonus**) if that total is at least 63 (the **threshold**).

2.3 Basic Yahtzee

The next variant of Yahtzee, which I call Basic Yahtzee, has all the patterns of Official Yahtzee, but the scoring rules are somewhat simpler. Table 2 lists the *thirteen* patterns. These include the six patterns of Junior Yahtzee in, what is called, the **Upper Section**. The other seven patterns are in the **Lower Section**. The indicated pattern score is obtained only if the condition is met; otherwise, the score is zero.

Here is a brief explanation of the pattern conditions when scoring a dice roll in the Lower Section:

Three of a Kind At least three dice show the same value: a, a, a, b, c.

Pattern	Condition	Score
Aces	none	total value of all ones
Twos	none	total value of all twos
Threes	none	total value of all threes
Fours	none	total value of all fours
Fives	none	total value of all fives
Sixes	none	total value of all sixes
Three of a Kind	≥ 3 equals	total value of all dice
Four of a Kind	≥ 4 equals	total value of all dice
Full House	$2 + 3$ equals	25
Small Straight	sequence of 4	30
Large Straight	sequence of 5	40
Yahtzee	5 equals	50
Chance	none	total value of all dice

Table 2: Dice patterns and their scoring rules for Basic Yahtzee

Four of a Kind At least four dice show the same value: a, a, a, a, b.

Full House Three dice show one value and two another value: a, a, a, b, b with $a \neq b$.[1]

Small Straight At least four dice values are in sequence: a, $a{+}1$, $a{+}2$, $a{+}3$, b.

Large Straight All dice values are in sequence: a, $a{+}1$, $a{+}2$, $a{+}3$, $a{+}4$.

Yahtzee All dice values are the same: a, a, a, a, a.

The final scoring rules for Basic Yahtzee are similar to those of Junior Yahtzee. Again there are two variants depending on whether or not a bonus is given for scoring more than a treshold on the Upper Section (=Junior) patterns:

Without Bonus The final score is the total of all pattern scores.

With Bonus The final score is the total of all pattern scores, plus 35 (the **Upper Section Bonus**) if the total in the Upper Section is at least 63 (the **threshold**).

2.4 Official Yahtzee

Official Yahtzee is like Basic Yahtzee with Bonus, but extra Yahtzees (rolls of 5 equals) are treated specially in the pattern scoring rules:

Extra Yahtzee Bonus A Yahtzee earns a bonus of 100 (the **Extra Yahtzee Bonus**), no matter for what pattern it is scored, *provided* that the Yahtzee pattern has already been scored with 50. Note however that this bonus does *not* count toward reaching the threshold for the Upper Section Bonus.

[1] Some accept $a = b$ as well, but that is not according to the official rules.

Extra Yahtzee as Joker in Lower Section A Yahtzee, say of five y's, can be scored *in full* for the patterns *Full House* (score 25), *Small Straight* (score 30), or *Large Straight* (score 40), *provided* that both the Yahtzee pattern and the y-pattern in the Upper Section have already been scored (zeroes are OK). In that case, the Yahtzee is said to act as a **Joker**.

Note that the scoring rules for Official Yahtzee go beyond the rules of Abstract Yahtzee, because the pattern score not only depends on the roll and the chosen pattern, but also on (the current state of) the score card:

1. Whether or not the Extra Yahtzee Bonus is earned depends also on the state of the Yahtzee box.

2. Whether or not a Yahtzee (of value v) can act as Joker depends also on the state of the Yahtzee box and of the corresponding v-box in the Upper Section.

The final scoring rules for Official Yahtzee are similar to those of Basic Yahtzee with Bonus:

> The final score is the total of all pattern scores, plus 35 (the **Upper Section Bonus**) if the total of the Upper Section scores, each taken modulo the Extra Yahtzee Bonus (i.e. 100), is at least 63 (the **threshold**).

Note that regular scores (i.e. without Extra Yahtzee Bonus) are less than the Extra Yahtzee Bonus. Therefore, taking a score modulo the Extra Yahtzee Bonus means that this bonus is ignored. That is how the final scoring rules capture that Extra Yahtzee Bonuses do not count toward reaching the threshold for the Upper Section Bonus.

Rather than including the Extra Yahtzee Bonus in the pattern score, it is more practical to record Extra Yahtzee Bonuses in a separate box (as the versions sold in the U.S. and the Netherlands do) or to add them to the score in the Yahtzee box (as some web versions do). Scoring may then involve two boxes instead of just one, which also goes beyond the rules of Abstract Yahtzee.

2.5 Group Yahtzee

Every variant of Solitaire Yahtzee can be transformed into a corresponding variant of Group Yahtzee for two or more players. In Group Yahtzee, each player has a *personal* score card. Players take turns in *round-robin* fashion. A turn is played the same way as in Solitaire Yahtzee: first a dice roll is assembled and next it is scored on the personal score card. Players can see the current scores of each other and may use that knowledge when making their assembly and scoring choices. The **winner** is determined by the highest final score. There are no official rules for tie breaking.

Note that in Solitaire Yahtzee, the outcome of a game is a number, namely the final score. In Group Yahtzee, however, the outcome of a game is a player, namely the player with the highest score (assuming no tie).

2.6 Other Yahtzee Variants

There are many other variants of Yahtzee. The following list mentions a few of the more common ones:

Six-pack Yahtzee The rules of the version sold by Milton Bradley in the Netherlands state that the objective is to obtain the highest score in a total of *six* games played *sequentially*. [I coined the name 'Six-pack Yahtzee'.]

Triple Yahtzee Instead of playing multiple games sequentially, you can play them *concurrently*. In Triple Yahtzee, each player uses *three* score cards and, for each turn, the player has the extra freedom to choose a score card for scoring the turn.

With '1 pair' and '2 pair' patterns Of course, you can add, modify, or delete patterns and scoring rules as you like. One variant I have seen on the web includes two extra patterns: **one pair** (≥ 2 equals: a, a, b, c, d) and **two pairs** (≥ 2 equals of one value and ≥ 2 equals of another value: a, a, b, b, c with $a \neq b$. The score is the value in the pair(s).

3 Analyzing Yahtzee

In order to solve Yahtzee, it is necessary to analyze the game in terms of appropriate mathematical models. In particular, I will define such notions as game tree and playing strategy. These will also give us some insight into the complexity of the game.

3.1 Analyzing Micro Yahtzee

Let us start with the analysis of a much simpler instance of Abstract Yahtzee. **Micro Yahtzee** is played with *one* die and the score card has only *two* patterns: **Double** and **Square** (also see Table 3). Each of the two turns consists of a single roll with the die (*no* modification attempts), and the resulting value must be scored in a free pattern. Actually, only the first roll offers a choice, since the second roll must be scored in the remaining pattern. The objective is to obtain the highest final score, being the sum of the two pattern scores.

Pattern	Condition	Score
Double	none	the die value doubled
Square	none	the die value squared

Table 3: Dice patterns and their scoring rules for Micro Yahtzee

What would you do when playing Micro Yahtzee and on your first turn you roll a four: score 8 for *Double*, or 16 for *Square*? To answer that question, it is important to clarify what the objective is in mathematical terms, that is, to construct a mathematical model of the game.

The first model that I present is based on sequences of events occurring in a game of Yahtzee: dice rolls, and assembly and scoring choices. In Micro Yahtzee there are no assembly choices. At the start of Micro Yahtzee, in the first turn, one of six values comes up when the die is rolled, say 4. The player then chooses one of two patterns for scoring, based on the roll, say *Double* (scoring 8). Next, in the second turn, again one of six values comes up, say 3. Now there is no choice where to score it, since there is only one pattern left (*Square*, scoring 9). The final score for the example game is $2 \cdot 4 + 3^2 = 17$. Choosing *Square* after rolling 4, and *Double* after rolling 3 yields a final score of $4^2 + 2 \cdot 3 = 22$.

All event sequences can be collected in a **game tree**. Each *node* in the tree represents a game state, each *edge* an event. The game tree is grown as follows. The initial state is the *root* of the tree. For each possible event e in some state s, there is an edge labeled e from node s to a fresh node t, denoted by $s \xrightarrow{e} t$. The game ends in a *leaf*, that is, a node without outgoing edges.

Part of the game tree for Micro Yahtzee is shown in Figure 1. All edges are directed from left to right, and the nodes where the player needs to make a (scoring) choice are circled. The game shown in full corresponds to first rolling 4 and choosing pattern *Double*, then rolling 3 and finally "choosing" pattern *Square*.

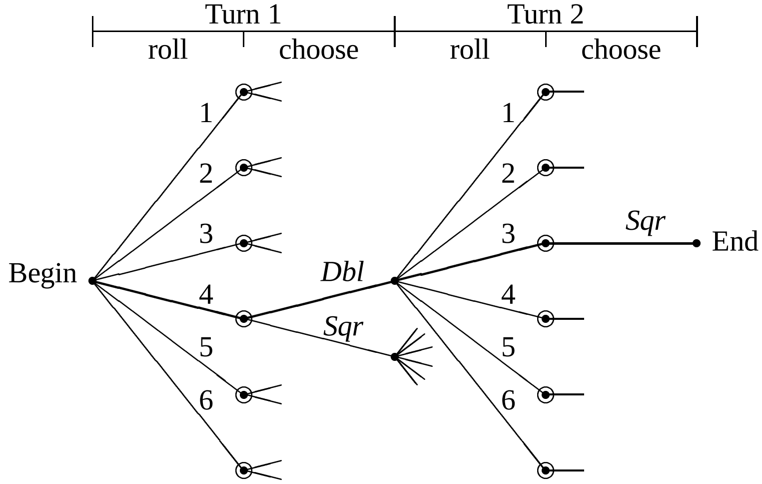

Figure 1: Part of Micro Yahtzee's game tree with one complete game

Each game corresponds to a path in the game tree from the root (Begin) to some leaf (End). The path is uniquely determined by the leaf. Thus, there are $6 \cdot 2 \cdot 6 = 72$ possible games of Micro Yahtzee.

In each node of the game tree, either all outgoing edges correspond to dice roll events or to choice events. A game of Yahtzee can also be considered to involve an extra player: the "*Dice Devil*", who randomly "chooses" the outcome of a dice roll. Each event in the game then consists of a player making a move. The Dice Devil starts the game and alternates its moves with those of the real player(s). Nodes where the Dice Devil moves are called **roll nodes**, and nodes where a "real" player moves are called **choice nodes** (circled in Figure 1).

A **strategy** for playing Yahtzee specifies for each choice node a way of choosing one of the outgoing edges. If every choice by strategy S depends only on the node (and not on other circumstances or chance), then S is called a **deterministic strategy**. In that case, strategy S is determined by a (partial) mapping C_S from choice nodes to their outgoing edges. If g is (the path to) a choice node, then $C_S(g)$ is the choice made by S. For example, the number of deterministic strategies for playing Micro Yahtzee equals $2^6 = 64$. Some of these strategies perform better than others.

For each ended game g, its final score

$$F(g) \tag{1}$$

can be computed by considering the events on the game's path. In Micro Yahtzee, a path with events i-*Double*-j-*Square* scores $2i + j^2$, whereas a path with events i-*Square*-j-*Double* scores $i^2 + 2j$.

Each playing strategy S induces a *probability distribution* on the the *sample space* consisting of all games. Let G_S be the *random variable* ranging over all games with probability determined by strategy S. The probability that game g occurs under strategy S is denoted by

$$P(G_S = g) \tag{2}$$

See [1] for an introduction to Probability Theory.

Given a deterministic strategy, the occurrence probability of every game can be computed as the product of all event probabilities on its path. Note that each die roll event is independent of all other events, and that choice events have probability one (if the strategy chooses it), or zero (otherwise). In Micro Yahtzee (with a fair die), every path has either probability $1/6^2$ (if the strategy chooses it), or zero (otherwise).

The final score is also a random variable whose probability distribution depends on the playing strategy S. This random variable is denoted by $F(G_S)$. The *mathematical expectation* of the final score or, briefly, the **expected final score** for strategy S is given by

$$E[F(G_S)] \;=\; \sum_g P(G_S = g)F(g) \tag{3}$$

The *Weak Law of Large Numbers* implies that, for each strategy S, its *mean* final score taken over N games approximates its *expected* final score $E[F(G_S)]$ with any accuracy level and confidence level provided that N is sufficiently large. More formally,

$$\forall_{\varepsilon>0}\,\forall_{\delta>0}\,\exists_N\,\forall_{n\geq N}\,P(|M_{S,n} - \mu_S| < \varepsilon) > 1 - \delta \tag{4}$$

where random variable $M_{S,n}$ is the mean final score of S over n games and μ_S is the expected final score $E[F(G_S)]$ of S. Thus, the expected final score is a reasonable performance measure for strategies. The best strategies are those achieving the highest expected final score.

Consider the following family of seven strategies $S.t$ $(0 \leq t \leq 6)$ for Micro Yahtzee:

$$S.t \text{ scores the first roll } i \text{ as } \textit{Double} \text{ if } i \leq t \text{, and as } \textit{Square} \text{ otherwise.} \tag{5}$$

Strategy $S.0$ always scores the first roll as *Square*, and strategy $S.6$ always scores the first roll as *Double*. For the expected final score of strategy $S.t$ one can derive:

$$\begin{aligned} E[F(G_{S.t})] &= \sum_{i \leq t} \sum_{j} (2i + j^2)/6^2 + \sum_{i>t} \sum_{j} (i^2 + 2j)/6^2 \\ &= \sum_{i \leq t} (2i + \sum_{j} j^2/6)/6 + \sum_{i>t} (i^2 + \sum_{j} 2j/6)/6 \\ &= \sum_{i \leq t} (2i + 15\frac{1}{6})/6 + \sum_{i>t} (i^2 + 7)/6 \end{aligned}$$

Table 4 lists the expected final scores for all strategies $S.t$. Note that the expected

t	0	1	2	3	4	5	6
$E[F(G_{S.t})]$	$22\frac{6}{36}$	$23\frac{25}{36}$	$25\frac{2}{36}$	$25\frac{33}{36}$	$25\frac{34}{36}$	$24\frac{29}{36}$	$22\frac{6}{36}$

Table 4: Expected final scores of strategies $S.t$ for Micro Yahtzee

final scores of strategies $S.0$ and $S.6$ are equal, because in those cases the strategy choice is 'blind' to the die roll. In fact, the expected final score is then just the sum of the expected scores for each pattern separately:

$$\begin{aligned} \sum_{i} \sum_{j} (2i + j^2)/6^2 &= \sum_{i} \sum_{j} (i^2 + 2j)/6^2 \\ &= \sum_{i} 2i/6 + \sum_{i} i^2/6 \\ &= E[2D] + E[D^2] \end{aligned}$$

where random variable D is the outcome of a single fair die roll. Keep in mind that $E[D^2] \neq E[D]^2$.

From Table 4 it is clear that strategy $S.4$ has the highest expected final score among all strategies $S.t$ (though it is only marginally better than $S.3$). However, Table 4 lists the expected final score of only 7 of the 64 strategies for Micro Yahtzee. Is there no better strategy than $S.4$? Intuitively this may seem obvious, but a computation would be more convincing.

For Micro Yahtzee it is straightforward to determine the optimal strategies (that obtain the highest expected final score). Consider the choice to be made in the first turn after rolling i. Since the next roll j must be scored at the other pattern, the expected score for that remaining pattern is either $E[2D] = 7$ or $E[D^2] = 15\frac{1}{6}$. If i is scored as *Double*, then the expected final score is $2i + 15\frac{1}{6}$, and if i is scored as *Square*, the the expected final score is $i^2 + 7$. These so-called *conditional expectations* are listed in Table 5. As can be seen from the rightmost column, $S.4$ is indeed the (only) optimal strategy for Micro Yahtzee.

First roll i	Expected final score when first choice is *Double* $2i + 15\frac{1}{6}$	*Square* $i^2 + 7$	Optimal choice
1	$17\frac{1}{6}$	8	*Double*
2	$19\frac{1}{6}$	11	*Double*
3	$21\frac{1}{6}$	16	*Double*
4	$23\frac{1}{6}$	23	*Double*
5	$25\frac{1}{6}$	32	*Square*
6	$27\frac{1}{6}$	43	*Square*

Table 5: Conditional expected final scores and optimal choices for Micro Yahtzee

For certain *zero-sum games*, it is possible that nondeterministic strategies do better than deterministic strategies. *[Reference to Game Theory]* The preceding analysis may help convince you that, for Solitaire Yahtzee, nondeterministic strategies cannot do better than deterministic strategies. The reason is that die roll events are independent of all other events, both other die roll events and choice events. The Dice Devil is not demonic. Its ‘strategy’ is fixed in advance, namely to provide fair die rolls.

3.2 Analyzing Basic Yahtzee Without Bonus

Let us now consider Basic Yahtzee Without Bonus. Its game tree is enormous. There are thirteen turns, each consisting of three rolls, and after each roll there is a choice, either for keeping some dice or for scoring one of the free patterns. Thus, each game consists of $13*3*2 = 78$ events: $13*3 = 39$ roll events and $13*3 = 39$ choice events (of which $13*2 = 26$ keep events and 13 score events).

An estimate of the size of the game tree (i.e. the number of leaves = the number of games) can be obtained from the following model. The first roll event of each turn has $6^5 = 7776$ possible outcomes (five dice with six values each). Each modification attempt starts by picking a subset of the dice ($2^5 = 32$ possibilities) and then rerolling the remaining dice (6^k possibilities when rerolling k dice). Together, there are $\sum_{k=0}^{5}\binom{5}{k}6^k = 7^5 = 16807$ possibilities[2] for each of the two modification

[2]This number can be obtained directly (without binomial coefficients): when keeping and rerolling, you choose for each of the five dice whether to keep it (one possibility) or to reroll it

attempts. After the last modification attempt of turn t ($1 \leq t \leq 13$), there are $14-t$ possible choices for scoring in one of the remaining empty pattern boxes. The number of possibilities for turn t are thus

$$6^5 \left(7^5\right)^2 (14-t) \quad = \quad 2\,196\,527\,536\,224\,(14-t) \tag{6}$$

Consequently, the total number of possible Basic Yahtzee games in this model equals

$$\left(6^5 7^{10}\right)^{13} 13! \quad \approx \quad 1.725 \times 10^{170} \tag{7}$$

The nonzero probabilities of games in this model range from

$$\left(6^{-5}\right)^{3\times 13} \approx 5.5 \times 10^{-151} \quad \text{to} \quad \left(6^{-5}\right)^{13} \approx 3.8 \times 10^{-50} \tag{8}$$

corresponding to always rerolling all dice and never rerolling any dice.

The number of deterministic strategies in this model for Basic Yahtzee is truly horrendous. The number c of choice nodes in the game tree is a number like that given by (7). The number of options per choice ranges from 1 (pattern choice in the last turn before the end) to 32 (for the subset of keepers). So, the number of deterministic strategies is something like 10^c. Note, however, that these are *completely specified* strategies, that is, they determine a choice in every imaginable game state, even if that state is not reachable from the initial state under the strategy.

3.2.1 Refining the model with bags of dice values

The preceding model assumes that the dice are distinghuishable. However, when you consider the rules of Basic Yahtzee, it is clear that the dice can be assumed *in*distinghuishable.[3] Only the *bag of values* matters for applying the (scoring) rules; the identity of the die that has rolled a value is irrelevant. Hence, many of the games treated separately above, can be considered equivalent. For instance, the roll 4, 1, 2, 5, 1 is equivalent to 1, 1, 2, 4, 5 (and to 58 other rolls; also see Table 6).

Before remodeling the game with indistinghuishable dice, let me classify and count the various ways of rolling five dice. This is not strictly necessary, but the results can be of help when testing Yahtzee software. Let $V = \{1, \ldots, 6\}$ be the set of six possible dice values. A **roll of five distinghuishable dice** is a *list* of five values from V, that is, a mapping from the set of five dice to V. As we have seen earlier, there are $6^5 = 7776$ such rolls. Each of these rolls is equally likely (given fair dice). Therefore, the preceding model is called the **List Model**.

A **roll of five indistinghuishable dice** is a *bag* of five values from V, that is, a mapping β from V to $\{0, \ldots, 5\}$ such that $\sum_{v=1}^{6} \beta(v) = 5$, where each $\beta(v)$ is value v's *occurrence count* or *multiplicity* in β. There are 'only' $\binom{5+6-1}{5} = 252$ such rolls. These rolls are *not* equally likely. Note that each list λ corresponds to

(six possibilities), giving altogether $(1+6)^5$ possibilities.

[3] If the dice are *un*fair, each in its own way, then this assumption is no longer be valid.

one bag β with $\beta(v) = \#\{\, i \mid \lambda(i) = v \,\}$, but each bag corresponds to many lists (see column A in Table 6). The probabilities range from $1/7776$ for any roll of 5 equal values, to $120/7776 = 1/64.8$ for any roll of 5 distinct values. Thus, the nonzero probabilities of games in this model range from

$$\left(6^{-5}\right)^{3\times 13} \approx 5.5 \times 10^{-151} \quad \text{to} \quad \left(\frac{120}{7776}\right)^{13} \approx 2.8 \times 10^{-24} \tag{9}$$

corresponding, on one hand, to a game where all dice are always rerolled and each roll results in 5 equals (yes, highly unlikely), and, on the other hand, to a game where no dice are ever rerolled and every roll results in 5 distinct values (still quite unlikely). This new model is called the **Bag Model**.

	Bag class		A	B	C
values	multiplicities m_i	n_i	#lists/bag	#bags	#lists
a, a, a, a, a	5, 0, 0, 0, 0, 0	5, 0, 0, 0, 0, 1	1	6	6
a, a, a, a, b	4, 1, 0, 0, 0, 0	4, 1, 0, 0, 1, 0	5	30	150
a, a, a, b, b	3, 2, 0, 0, 0, 0	4, 0, 1, 1, 0, 0	10	30	300
a, a, a, b, c	3, 1, 1, 0, 0, 0	3, 2, 0, 1, 0, 0	20	60	1200
a, a, b, b, c	2, 2, 1, 0, 0, 0	3, 1, 2, 0, 0, 0	30	60	1800
a, a, b, c, d	2, 1, 1, 1, 0, 0	2, 3, 1, 0, 0, 0	60	60	3600
a, b, c, d, e	1, 1, 1, 1, 1, 0	1, 5, 0, 0, 0, 0	120	6	720
		Total		252	7776

Table 6: Classifying and counting rolls of 5 dice (a, b, c, d, e all distinct)

Table 6 classifies and counts the bags and corresponding lists of five values from V. The classification is based on the bag of the value bag's multiplicities. Two value bags that have the same bag of multiplicities are in the same bag class. The bag of multiplicities is given here by sorting the multiplicities in descending order: $m_1, m_2, m_3, m_4, m_5, m_6$. For instance, the rolls 1, 1, 1, 4, 5 and 1, 3, 3, 3, 6 are in the same bag class, because their multiplicities are 3, 0, 0, 1, 1, 0 and 1, 0, 3, 0, 0, 1 respectively, both reducing to 3, 1, 1, 0, 0, 0 after sorting.

Column A of Table 6 gives the number of lists corresponding to one bag in the class, that is, the number of ways to permute the values. This number is obtained as the multinomial coefficient

$$\binom{5}{m_1\, m_2\, m_3\, m_4\, m_5\, m_6} = \frac{5!}{m_1!m_2!m_3!m_4!m_5!m_6!} \tag{10}$$

where the m_i are the multiplicities with $\sum_{i=1}^{6} m_i = 5$. Column B gives the number of bags in the class, that is, the number of ways to permute the multiplicities, thereby inducing an assignment of distinct values from V to the relevant a, b, c, d, e. This number is obtained as the multinomial coefficient

$$\binom{6}{n_0\, n_1\, n_2\, n_3\, n_4\, n_5} = \frac{6!}{n_0!n_1!n_2!n_3!n_4!n_5!} \tag{11}$$

where $n_i = \#\{ j \mid m_j = i \}$; note that $\sum_{i=0}^{5} n_i = 6$ and $\sum_{i=0}^{5} i n_i = 5$. Column C gives the number of lists in each class, that is, $C = AB$.

Pattern	**Probability**	**Expected Score**
Aces	1	0.83
Twos	1	1.67
Threes	1	2.50
Fours	1	3.33
Fives	1	4.17
Sixes	1	5.00
Three of a Kind	1656/7776	3.73
Four of a Kind	156/7776	0.35
Full House	300/7776	0.96
Small Straight	1200/7776	4.63
Large Straight	240/7776	1.23
Yahtzee	6/7776	0.04
Chance	1	17.50
Grand Total		45.95

Table 7: Probabilities and expected scores for the Basic patterns in a single roll

Using Table 6, it is straightforward to compute the probabilities for rolling, in a single roll, each of the Basic Yahtzee patterns, that is, a roll satisfying the pattern's condition. These probabilities are listed in Table 7. Here follows some further explanation, in the order of increasing complexity.

Yahtzee is obtained by any bag in the class of 5 equals. This corresponds to 6 lists.

Full House is obtained by any bag in the class of 3+2 equals. This corresponds to 300 lists.

Four of a Kind is obtained by a *Yahtzee* or any bag in the class of 4+1 equals. This corresponds to $6 + 150 = 156$ lists.

Three of a Kind is obtained by a *Four of a Kind*, a *Full House*, or any bag in the classes of 3+1+1 equals. This corresponds to $156 + 300 + 1200 = 1656$ lists.

Large Straight is obtained by 2 of the 6 bags in the class of 5 distinct values (1+1+1+1+1 equals): 1, 2, 3, 4, 5 and 2, 3, 4, 5, 6. This corresponds to $2 \times 120 = 240$ lists.

Small Straight is obtained by a *Large Straight*, or 2 additional bags in the class of 5 distinct (viz. 1, 3, 4, 5, 6 and 1, 2, 3, 4, 6), or 12 of the 60 bags in the class 2+1+1+1 equals: 1, 2, 3, 4 or 2, 3, 4, 5 or 3, 4, 5, 6 with any of the four values doubled. This corresponds to $4 \times 120 + 12 \times 60 = 1200$ lists.

It is also easy to compute the expected score of a single roll for each Basic pattern. In the Upper Section, the expected score for the v-pattern ($1 \leq v \leq 6$) is $5vP(D = v) = 5v/6$, where random variable D is the outcome of a fair die roll. The expected scores for *Three of a Kind*, *Four of a Kind*, and *Chance* equal their probability times $5E[D] = 17.5$, because each die value is equally likely in these patterns. The expected scores for *Full House, Small Straight, Large Straight*, and *Yahtzee* equal their probability times 25, 30, 40, and 50 respectively. These expected pattern scores are also listed in Table 7.

The numbers in Table 7 are related to strategies that reroll arbitrarily chosen dice and that score each roll in an arbitrary free pattern, where 'arbitrary' means 'regardless of the roll'. There are numerous ways to make such arbitrary choices. For example, never reroll any dice, always reroll all dice, choose the topmost free pattern, etc. These correspond to distinct strategies with the same expected final score of just less than a measly 46, which most human players will easily beat.

3.2.2 The number of games in the Bag Model

To compute the number of games in the Bag Model is a little harder than in the List Model. The reason is that the various rolls (value bags) do not all have the same number of subbags to choose for keeping. In the List Model, each roll allows precisely $2^5 = 32$ choices for keeping. In the Bag Model, a roll of five equal values, for instance, allows only 6 choices (zero to five of that value), whereas a roll of five distinct values allows 32 choices (each value is either chosen or not). Column T of Table 8 gives the number of subbags for each bag class[4]. This number is obtained as $\prod_{i=1}^{6}(m_i+1)$, where the m_i are the multiplicities. The remainder of Table 8 is explained below.

Bag class	5	4 1	3 2	3 1^2	2^2 1	2 1^3	1^5	4	3 1	2^2	2 1^2	1^4	3	2 1	1^3	2	1^2	1	0	T
5	1							1					1			1		1	1	6
4 1		1						1	1				1	1		1	1	2	1	10
3 2			1						1	1			1	2		2	1	2	1	12
3 1^2				1					2		1		1	2	1	1	3	3	1	16
2^2 1					1					1	2			4	1	2	3	3	1	18
2 1^3						1					3	1		3	4	1	6	4	1	24
1^5							1					5			10		10	5	1	32

Table 8: Number of subbags for bags of 5 dice, arranged by class

The unkept dice, say k in number, are rerolled and combined with the kept dice to give again a bag of size 5. The number of bags over V with size k equals $\binom{k+\#V-1}{k}$. The total number of such bags for $0 \leq k \leq 5$ equals $\binom{5+\#V}{5}$. Here is the

[4] Bag class identifiers are shown in abbreviated form, e.g. 3 1^2 denotes 3, 1, 1, 0, 0, 0.

whole list:

k	0	1	2	3	4	5	total
# bags over V of size k	1	6	21	56	126	252	462

(12)

The roll after the first modification attempt is then subjected to a second modification attempt. In order to compute the number of possibilities, it is convenient to keep track of how many bags of each class occur at every stage of the turn. Table 8 lists, for each class of bags over V of size 5, how many subbags it has in each of the classes of bags of size at most 5. For instance, a bag in the class of 2+2+1 equals has 4 subbags of size 3 in the class of 2+1 equals: there are 4 ways to leave out two values, namely 2 ways of leaving out 2 equals, and two ways of leaving out the single value and one of the double values.

Table 9 lists, for each class of bags over V of size at most 5, in how many ways it can be completed to a bag of size 5 in a given class. For instance, a bag of size 3 in the class of 2+1 equals can be completed in 8 ways to a bag of size 5 in the class of 2+2+1 equals: there are 4 ways to add two equal new values, and 4 ways to add one new value and one value present once. Column T shows the total number of completions. These totals are the same as given by (12).

Bag class	#	5	4 1	3 2	3 1^2	2^2 1	2 1^3	1^5	T
5	6	1							1
4 1	30		1						1
3 2	30			1					1
3 1^2	60				1				1
2^2 1	60					1			1
2 1^3	60						1		1
1^5	6							1	1
4	6	1	5						6
3 1	30		1	1	4				6
2^2	15			2		4			6
2 1^2	60				1	2	3		6
1^4	15						4	2	6
3	6	1	5	5	10				21
2 1	30		1	2	4	8	6		21
1^3	20				3	3	12	3	21
2	6	1	5	10	10	20	10		56
1^2	15		2	2	12	12	24	4	56
1	6	1	10	10	30	30	40	5	126
0	1	6	30	30	60	60	60	6	252

Table 9: Number of completions to bag of 5 dice, arranged by class

Let

- e be the 1×19-matrix $[0 \ldots 0\,1]$,
- A the 7×19-matrix of Table 8,

- B the 19×7-matrix of Table 9, and
- u the 7×1-matrix $[1 \ldots 1]^{\mathrm{T}}$.

Then the matrix product

$$e \times A \times (B \times A)^m \times u \tag{13}$$

gives the number of assembly possibilities for a single turn consisting of m modification attempts. Matrix e represents the initial situation of a turn, when no dice are kept, i.e. when a single bag of size 0 is to be completed to a bag of size 5. Multiplication by A gives a 7×1-matrix with the number of bags in each of the 7 classes of bags of size 5. Subsequent multiplication by B gives a 1×19-matrix with the number of subbags in each of the 19 classes of bags of size at most 5. Multiplication by matrix u adds all the entries. Substituting $m := 2$ in (13) yields 272 212 248 possibilities after two modification attempts. Consequently, the number of games in the Bag Model equals

$$272\,212\,248^{13}\,13! \quad \approx \quad 2.806 \times 10^{119} \tag{14}$$

This is a huge reduction compared to the number of games in the List Model given by (7), but still an enormous number.

It is instructive to understand the relationship between strategies in the List Model and strategies in the Bag Model. Every deterministic strategy S in the Bag Model can be translated into an equivalent deterministic strategy S' in the List Model, by translating lists into bags and then applying S. The resulting strategy S' is equivalent to S in the sense that it has the same characteristics, such as expected final score. The reverse translation does not always work, because the translation from bags into lists is not uniquely determined. Choices by a strategy in the List Model can be based on features of a roll that are ‘invisible’ in the Bag Model.

For instance, consider the first rolls 4, 1, 2, 5, 1 and 1, 1, 2, 4, 5, which differ as lists but are equal as bags. The dice kept after these two rolls can differ for a deterministic strategy in the List Model, but must be the same for a deterministic strategy in the Bag Model. A deterministic strategy S in the List Model can only be translated into an equivalent deterministic strategy S' in the Bag Model when choices by S are constant on bag classes. Such a deterministic strategy in the List Model is called a **bag-deterministic strategy**. Note that the translation of a deterministic strategy from the Bag Model to the List Model yields a bag-deterministic strategy.

3.2.3 Optimal strategies for Basic Yahtzee Without Bonus

The approach for optimizing Micro Yahtzee can be generalized. It yields a recursive evaluation scheme for the expected final score of a strategy for a Yahtzee game in general. Moreover, it leads to a means of determining strategies that maximize the expected final strategy.

Define the **additional score** by strategy S when continuing from game state g via g' to the end by

$$A(g' \mid g) \tag{15}$$

It has the following properties

$$\begin{aligned} A(g \mid \langle\rangle) &= F(g) \text{ if } g \text{ ended} \\ A(\langle\rangle \mid g) &= 0 \text{ if } g \text{ ended} \\ A(eg' \mid g) &= f(e \mid g) + A(g' \mid ge) \end{aligned}$$

Maximum score under official rules: All rolls are (appropriate) Yahtzees, yielding $5+10+15+20+25+30+35+30+30+25+30+40+50+30+12*100 = 5*21+35+4*30+25+40+50+1200 = 140+120+115+1200 = 1575$

4 Solving Official Yahtzee

Combinatorial and probabilistic issues have been dealt with. Computational and numeric issues. Double, self-organizing, dynamic programming.

5 Investigating Optimal Strategies

What is the worst opening roll? What is the best? What to do if you open with 4 sixes and 1 five? When to zero out Yahtzee or Large Straight? What is expected number of zeroed out categories? What is the minimum score? What is the median score?

	# Modification attempts		
Pattern	0	1	2
Aces	0.83333	1.52778	2.10648
Twos	1.66667	3.05556	4.21296
Threes	2.50000	4.58333	6.31944
Fours	3.33333	6.11111	8.42593
Fives	4.16667	7.63889	10.53241
Sixes	5.00000	9.16667	12.63889
Three of a Kind	3.72685	10.20870	15.19466
Four of a Kind	0.35108	2.42151	5.61126
Full House	0.98380	4.85254	9.15362
Small Straight	4.65278	12.37972	18.48075
Large Straight	1.26543	5.31739	10.61274
Yahtzee	0.03858	0.63157	2.30143
Chance	17.50000	21.25000	23.33333

Table 10: Expected optimal scores for a single turn

Compare the column for no modification attempts to Table 7.

Pattern	**Avg. score**	**% 0 scores**
Aces	1.88	10.87
Twos	5.28	1.83
Threes	8.56	0.95
Fours	12.16	0.61
Fives	15.69	0.51
Sixes	19.18	0.54
U. S. Bonus	23.80	32.00
Three of a Kind	21.66	3.26
Four of a Kind	13.09	36.37
Full House	22.59	9.36
Small Straight	29.46	1.80
Large Straight	32.70	18.24
Yahtzee	16.87	66.25
Chance	22.01	0.00
Extra Yahtzee	9.58	91.76
Grand Total	254.51	0.00

Table 11: Scoring statistics of an optimal strategy for Official Yahtzee, based on one million simulated games

6 Conclusion

6.1 Acknowledgments

I would like to thank Marga Daniëls, Jurjen Bos, . . . for their helpful comments and suggestions.

References

[1] W. Feller. *An Introduction to Probability Theory and Its Applications: Volume I*. Third Edition, Wiley, 1968.