# ISYE 6644: SIMULATION:

# EVALUATING THE PERFORMANCE OF SCORING STRATEGIES IN YAHTZEE USING PROBABILISTIC MODELS AND SIMULATION

Matthew Cheung, Dylan Riggs

Group 28, Topic 12

Georgia Institute of Technology, 2024

ABSTRACT

Yahtzee is a game that incorporates strategy and probability in which players try to score as many points as possible over 13 rounds. This is done by choosing scoring categories based on their dice rolls. In this work, we compare three strategies: the Greedy Strategy, which focuses on maximizing points for each individual turn, the Balanced Strategy, which aims to spread scores evenly across categories using expected values of future rolls, attempting to ensure a higher final score, and lastly the Bonus First strategy, which has the goal of reaching the bonus points for scoring enough points in the upper section.

We employ probabilistic simulations to test these methods, incorporating essential Yahtzee scoring elements like reroll options and the 35-point bonus for achieving 63 or more points in the upper section. Each approach undergoes 1,000 simulated games to monitor dice outcomes, score allocations, and category selections. Using the data from these simulations allows us to determine the strengths of weaknesses of each strategy.

This analysis gives insight on how strategic choices inside the game can affect overall performance and offers intuition on decision-making processes in probabilistic games.

## BACKGROUND & DESCRIPTION OF PROBLEM

Yahtzee is a popular dice game that combines strategy and chance. Its captivating blend of score optimization and probabilistic decision-making has enthralled gamers for decades. Each turn, players can roll five dice up to three times and select from 13 scoring categories in an attempt to maximize their overall score over the course of 13 rounds. These categories are separated into a lower area that includes combinations like "Full House" or "Yahtzee" (five of a kind) and an upper section that involves sums of particular dice values (such as "Twos" or "Fives"). A 35-point bonus, which is often a key point to success for many players, is awarded for achieving a subtotal of 63 or higher in the top section. There is an inherent trade-off in the game's structure: should players seek high-scoring categories immediately, or disregard them for a more balanced scorecard?

There has been previous research on strategic decision-making in Yahtzee, like the study done by Cornell University on creating an ultimate gameplay algorithm. However, there is little simulation-based research available to discuss the relative advantages of other various gameplay tactics. Different approaches to managing these trade-offs in gameplay are represented by strategies like the Balanced Strategy, which strives for even category completion, and the Greedy Strategy, which emphasizes maximizing immediate scores. Knowledge about how these tactics function in various gameplay situations might help one make the best choices when faced with uncertainty.

Our goal in this study is to assess these tactics by simulating Yahtzee gameplay in great detail. Important aspects of the game, such as scoring systems, rerolling choices, and the top section bonus, are captured in our simulation. The fundamental probabilistic models frequently used in decision-making games are expanded upon in this paper and modified to account for the mechanics of Yahtzee.

The rest of this work is structured as follows: We start by outlining the process for modeling gameplay and putting the techniques into practice. The findings of our investigation are then shown, contrasting the Greedy, Balanced, and Bonus First techniques' performance on important metrics like overall scores and bonus attainment rates. Lastly, we talk about the ramifications of our results, how they apply to more general decision-making situations, and possible future developments for this work.

## THEORY

Developing successful strategy in Yahtzee requires an understanding of the associated probability. The game's decision-making process focuses on determining which dice to keep or reroll to maximize scoring potential and assessing the probability of reaching particular results from each roll. The underlying probability functions for Yahtzee are described in this section, beginning with the odds of the first roll's outcomes and moving on to the odds of following rolls.

### First Roll Probabilities

Each of the five dice has six possible results when all five are rolled, hence $6^5 = 7{,}776$ possible combinations. The basis for assessing the probability of any possible outcomes is the distribution of these combinations. On the initial roll, the likelihood of rolling any certain set of dice is:

$$P(Configuration) = \frac{1}{6^5}$$

For example, the probability of rolling a Yahtzee (all five dice showing the same number) is:

$$P(Yahtzee\ on\ First\ Roll) = \frac{6}{6^5} = \frac{1}{1{,}296}$$

## Probabilities for Subsequent Rolls

The reroll decisions depend on the probabilities of improving the current configuration. Given the dice kept after the first roll, the number of combinations for the remaining dice is:

$$P(Outcome \mid Reroll) = \frac{1}{6^n}$$

where *n* is the number of dice rerolled. For example, if two dice are kept and three dice are rerolled, there are $6^3 = 216$ possible outcomes for the reroll.

## Expected Value Calculations

For category-specific decisions (e.g., aiming for a full house or a large straight), the expected value (EV) of keeping certain dice can be calculated by summing the probabilities of achieving the desired outcome multiplied by the associated score. For example, if aiming for a Yahtzee:

$$EV(Yahtzee) = \sum_{i=1}^{6} P(Yahtzee\ with\ i\ rerolled\ dice) \times 50$$

Where:

$$P(Yahtzee\ with\ i\ rerolled\ dice) = \binom{n}{i}\left(\frac{1}{6}\right)^i\left(\frac{5}{6}\right)^{n-i}$$

Here, $\binom{n}{i}$ represents the number of ways *i* successes can occur in *n* trials.

## Keeping Dice or Rerolling

The likelihood of getting a better result determines whether to keep or reroll a die. Keeping consecutive dice (e.g., 2, 3, 4) while aiming for a small straight, for example, increases the likelihood of completing the sequence in following rolls. The likelihood of finishing the sequence is:

$$P(Small\ Straight) = 1 - \left(\frac{5}{6}\right)^n$$

where *n* is the number of remaining dice to be rolled.

## Strategic Implications

These probabilities highlight how crucial it is to determine the expected value of outcomes and modify strategy in real time according to the status of the game. For instance, it is statistically beneficial to concentrate on numbers rolled with a higher frequency when striving for the upper section bonus (earning 63+ points).

We will use these theoretical probabilities in our simulation to assess the effectiveness of various tactics in the sections that follow.

APPLICATION

In order to get our results, we had to use the probabilities of any combination of dice as discussed above and create a simulation to determine the best strategy a person could use when playing Yahtzee to get the highest average score. Since there is luck and randomness involved, there is never any guarantee when it comes to this game. Instead, we can use the expected outcomes and the probabilities of each possible outcome to determine what dice a person should keep depending on how many rolls they have left, and which categories are available. By doing this, it should allow the person playing to maximize their expected points on each turn. Using this method would amount to the Balanced Strategy as discussed above, trying to find the best category to attempt to get points in based on the probabilities of the possible results for future rolls. In this way, it is factoring in both future probabilities and game context to determine how the player plays the game.

Breaking down how the algorithm itself works, dice are “rolled” using a random number generator and then we take the available categories, the current dice on the table, and the number of rolls remaining. Using these three variables, the program iteratively calculates the expected value of keeping specific dice combinations, simulating all possible outcomes of future rolls to choose the optimal dice to retain and either go on to the next roll or put on to the scoresheet. The expected value is derived by averaging the potential scores across all possible dice configurations weighted by their probabilities, since there is an assumed uniform distribution of dice rolls.  This ensures that the player makes decisions aligned with maximizing the average score in the long run.

Additionally, the algorithm also factors in the conditional probability of receiving bonus points as a possibility with the expected value, so if a certain combination increases the odds of hitting the bonus, it increases the expected value to account for that. This adjustment ensures that decisions about top-category scores – scores that affect receiving the bonus - are informed not only by immediate payoffs but also by their contribution toward the cumulative sum required for the bonus at the end of the game.

Example Code of Expected Value Calculations:

```
def expected_value(kept_dice, rolls_left, available_categories, top_category_sum):

if rolls_left == 0:

        max_score = max(calculate_score(kept_dice, category) for category in
available_categories)

        new_top_category_sum = top_category_sum + max_score

         if new_top_category_sum >= 63:

            max_score += 35

        return max_score

num_dice_to_roll = 5 - len(kept_dice)

possible_rolls = itertools.product(range(1, 7), repeat=num_dice_to_roll)

total_ev = 0

num_outcomes = 0

for roll in possible_rolls:

        new_dice = list(kept_dice) + list(roll)

        ev = expected_value(new_dice, rolls_left - 1, available_categories,
top_category_sum)

        total_ev += ev

        num_outcomes += 1
```

(Figure 1) Shown above is psuedocode of our expected value algorithm that runs iteratively throughout our code

The next strategy that was used is the Greedy Strategy, which follows a more heuristic approach compared to the other strategies, as it does not implement the use of expected value and will instead attempt to go after the highest scoring category available at all times, then settling for whichever category gives the most points if the highest scoring category is not reached. Essentially, this means that this strategy will go for a Yahtzee every turn until it is reached and then attempt to get a large straight and so on. The only goal of this strategy is to attempt to get the highest scoring categories filled every time.

With the Greedy Strategy, the top categories are focused on less, but in attempting to get a Yahtzee, it means that the most common value dice will always be kept, so there are chances of some high scoring top categories. Overall, this strategy going into it was not expected to be the best, as there is a lot more risk involved, but could occasionally have very high scoring games, especially if you are playing with the bonus points for getting multiple Yahtzee's (typically 100 points) which was not factored into this simulation.

The final strategy that was implemented was the Bonus First strategy, which makes its only goal to get the bonus for receiving 63 or more points in the upper section of the scoresheet. When looking at the 6 categories that this has, one could get three of each value in this section (three

1’s through three 6’s) and receive the bonus. Knowing this, the algorithm ensures that one of the top categories is only used when there is at least a count of three in one of the categories.

The process in which the Bonus First Strategy goes through is a mix of heuristic and probabilistic approaches. The first step after a roll is to determine which dice had the highest count of the available categories in the upper section. The goal is then to continue this until the final roll, where the algorithm then selects a category if there is at least a count of 3 dice. If there is not three, then the highest scoring bottom category is selected. This means there is some risk of these categories getting 0 for their score. If every category has a score of 0, then the lowest probability of the remaining categories is selected to use for 0. So, if Yahtzee is available, it would be the first category with a 0 placed in it. Once all of the top categories are filled, the approach goes back to the algorithm that is used in the Balanced Strategy, using expected values based on probabilities to determine which dice to keep and which categories to select.

To run a simulation using any of the listed strategies, one would need to run the cells containing the code for each one of the strategies and then simply run the function that is named after the strategy that they want to select without any parameters. Print outs can be uncommented to analyze a specific run and look at individual choices that were made by the defined algorithm.

## RESULTS

In order to analyze which strategy is the best for a player to follow, we decided to run 1000 simulations of a game using each strategy. After each simulation, we saved the score, whether the bonus was reached, and the score in the upper section for comparison purposes. This resulted in the Balanced Strategy having the highest average score of 204.364, with the Bonus First Strategy having the next highest average score of 169.473, and the Greedy Strategy having the lowest average score of 145.165.

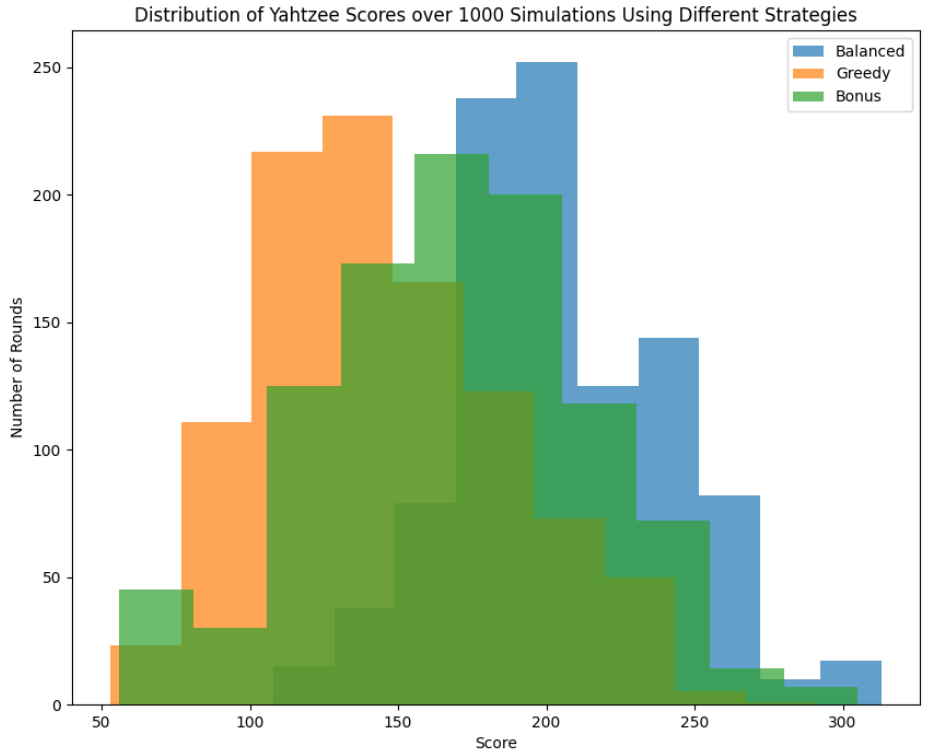


(Figure 2) Shown above is a histogram of the distributions of the scores using each of the three strategies defined in the paper

Looking into these results a bit further, we created 95% confidence intervals for each average score. The confidence interval from the Balanced Strategy ranges from 202.15 to 206.58. The confidence interval from the Bonus First Strategy ranges from 166.67 to 172.27. Lastly, the confidence interval from the Greedy Strategy ranges from 142.62 to 147.71. Due to the larger sample size of 1000, this allows the confidence interval to be narrower, so a greater understanding can be determined of which strategy is performing better. Since none of the confidence intervals have any overlap with one another, it is safe to assume that the Balanced Strategy would perform the best, then the Bonus First Strategy, and finally the Greedy Strategy at a 95% confidence level.

The next aspect that could be looked at is the rate at which the bonus is reached for hitting 63 points in the upper scoring section. Not surprisingly, the Bonus First Strategy performs at a significantly higher rate than the other two strategies, reaching the bonus 89.9% of the time in our simulations. The Balanced Strategy reaches the bonus 6.4% of the time and the Greedy Strategy reaches it just 2.5% of the time. This is a significant difference compared to the Bonus

First Strategy and shows that 35 bonus points are left on the table for the typically stronger performing Balanced Strategy.

Lastly, we looked at a similar statistic, the average score of the upper section for each strategy. The Bonus First Strategy had an average score of 69.828, the Balanced Strategy had an average score of 47.748, and the Greedy Strategy had an average score of 42.206. As noted previously, the Bonus First Strategy gets over 20 points more on average than the other two strategies, not including the 35 bonus points often received in those cases. This makes a near 55-point advantage in the upper section for the bonus first strategy compared to the other two strategies. Due to this, it means that the other strategies are also performing better in the lower section of the scorecard. Since the Balanced Strategy gets an average of 35 points more than the Bonus First Strategy, this means that the lower score is likely around 90 points higher to counteract the upper score difference.

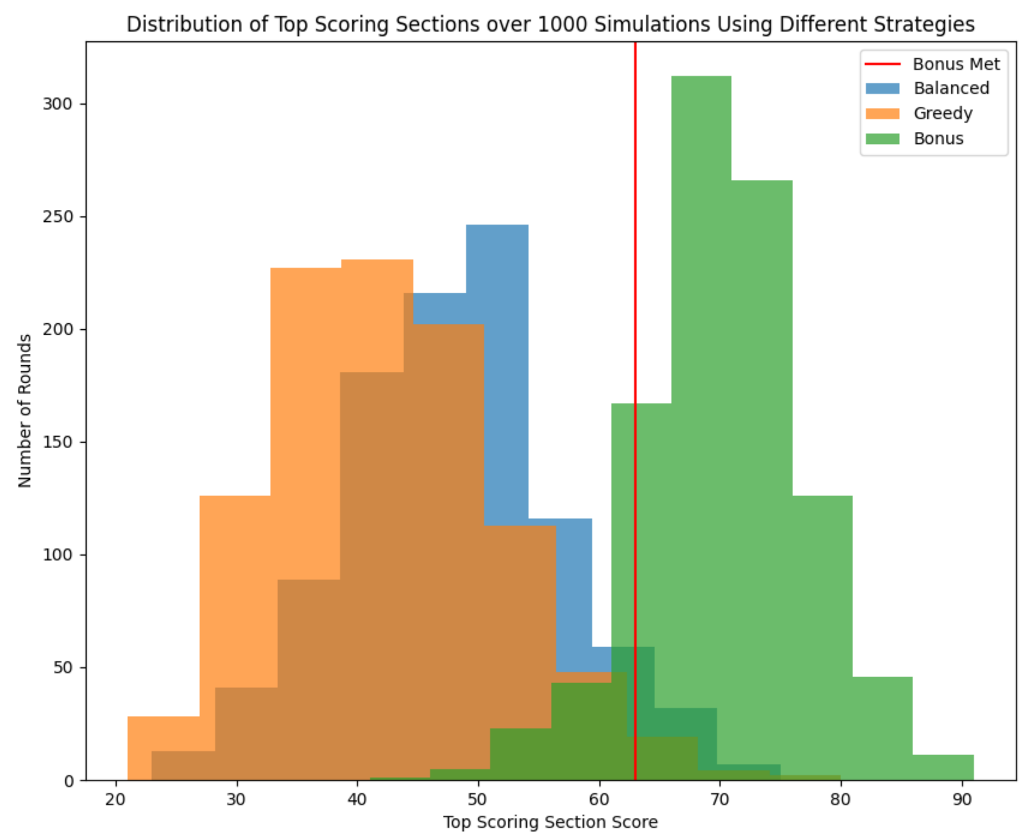


(Figure 3) This figure shows a histogram of the score distributions from the upper section of the scoresheet using our three strategies. The vertical line represents the cutoff point for 35 bonus points.

CONCLUSIONS

Overall, the simulations that were made determined that the best strategy to use was the Balanced Strategy, followed by the Bonus First Strategy, then the Greedy Strategy. Using this system won't always work however, as all three had games that performed better than some of the lower scores for other strategies. To maximize the likelihood of winning however, the Balanced Strategy would give you the strongest chance to win in any game that was played.

There are improvements that I believe could be made to these strategies down the road. After analyzing the results, it was clear that the Bonus First Strategy performed the best in the upper section of the scoresheet while the Balanced Strategy performed the best in the lower section of the scoresheet. If the Balanced Strategy had some modifications made to it in the form of heuristics that punished or rewarded the categories in the upper section depending on how many dice for that number were selected, it could possibly be an improvement on the current algorithm. This way, instead of simply taking the highest expected value or the highest points, it would make smarter selections of the upper categories and when to go after them or use them when the turn ends.

With the Greedy Strategy, I believe an improvement that could be made is to implement expected values into its algorithm and add weights to different categories that modify the expected values based on how important each category is. With this, one could determine which categories are the most important and make them more or less likely to be selected instead of just going after the highest available score in each round, reducing the overall risk by introducing an algorithmic approach to the heuristic approach that was used.

When someone is physically playing the game of Yahtzee themselves, it would be hard to implement the Balanced Strategy without knowing all of the probabilities and expected values of future combinations. Thus, it may be more realistic to look at specific combinations of dice and use this algorithm to determine for which combinations which dice should be kept, and which ones should be rerolled. It is likely that most people use a mix of the Greedy Strategy and the Bonus First Strategy when they are playing the game, as many will want to attempt to reach the bonus while also attempting to fill out the bottom section of the scoresheet as much as possible. This strategy would be the optimal end goal strategy for further work while also creating a sort of "cheat sheet" that allows players to know which dice to keep in which scenarios that maximizes their ability to win, which is the overall goal of this project.

REFERENCES

**Yahtzee Manifesto.** *Yahtzee Manifesto*, www.yahtzeemanifesto.com/. Accessed 1 Dec. 2024.

**Official Game Rules.** "Yahtzee." *Official Game Rules*, www.officialgamerules.org/board-games/yahtzee. Accessed 1 Dec. 2024.

**Cornell University Department of Mathematics.** "Yahtzee." *Mathematics Exploration Center: Probability*, pi.math.cornell.edu/~mec/2006-2007/Probability/Yahtzee.html. Accessed 1 Dec. 2024.