Investigating the 2-player stochastic duel

Justin Li

17 August 2025

1 Introduction

Suppose **Alice** and **Bob** (referred to **A** and **B** respectively) have guns and are able to shoot each other. They will eliminate the other player based on their probabilities of success, p and q. Additionally, they each have their own number of bullets and cannot eliminate the other player if their bullet count is reduced to 0. Each scenario can have the two players shooting simultaneously or in alternating order, with **A** going first every time. The objective of this investigation is to determine the probabilities of each outcome once the duel has ended and any significance within the data.

2 Methodology

For the purpose of analysis and reducing total data, starting bullet counts for the players are restricted between 1 and 10. Additionally, the probabilities of success range from 0.1 to 1, incrementing by 0.1 (0.1, 0.2, 0.3, ... 1.0). This results in 10000 distinct cases to be calculated. Some results are bound to be trivial and not interesting, such as A shooting first in an alternating duel with a 1.0 chance of success, with the probability matrix always resulting in A living and B dying. The goal is to find cases that may be more interesting.

One way of finding the probability matrix of outcomes in a particular duel can be found using a Monte Carlo simulation. With enough practice trials, the true probabilities can be estimated with a large number of trials. This method can be effective in presenting a rough estimate of the resulting probabilities and can be cross-checked with a recurrence to receive precise values for probabilities. Recurrence can be found with mathematical analysis or through a program, while a Monte Carlo simulation may only be implemented through code. Both methods were used to ensure that both methods were completed correctly.

2.1 Coding

The program was coded in Java and results would be written within a text file to reference. Using a Monte Carlo simulation and recurrence methods, probabilities would be determined for both simultaneous and alternating shooting scenarios. This resulted in a total of four different files needed for this analysis.

2.1.1 Output

The output would contain information about a certain situation (including number of bullets and the values of p and q for success probability. An example of the calculated probabilities can be seen below.

> Information:
>
> A:3 bullets 0.2 prob B:8 bullets 0.2 prob
>
> A lives, B Dies - 40.8752%
>
> A Dies, B Lives - 50.5936%
>
> A Lives, B Lives - 8.5312%
>
> A Dies, B Dies - 0.0%

The details and the resulting probabilities are listed, providing the probabilities for each end case. All probabilities add to 1 and correlate to that situation's calculated distribution.

2.1.2 Implementing the Monte Carlo Method

A large amount of trials will be conducted, with each trial being simulated with a random number generator to represent each shot, determining whether the shot was a success or a failure. The result is noted and complied within the program over many different trials for the situation. The relative frequency can be determined from the total of outcomes divided by the total amount of trials. This is done in the code by keeping track of the total instances of an outcome and once all trials of a situation are complete, the program writes the resulting information into the text file.

2.1.3 Implementing the Recurrence

Calculating a recurrence is less resource-intensive when running but more difficult to code into an algorithm. The solution for this is splitting the case into different segments. To understand the probability of both A and B surviving is

$$(1-p)^{bul_a} * (1-q)^{bul_b} \tag{1}$$

where $\mathbf{bul_A}$ and $\mathbf{bul_B}$ represent the bullets **A** and **B** have respectively. This is the case where all bullets fail and both **A** and **B** are both left alive.

The case where both **A** and **B** die only occurs during the simultaneous shot situation. The probability that both die during an alternating duel is 0 with this result being impossible. To calculate this, the total probability is,

$$\sum_{i=0}^{n} pq(1-p)^i(1-q)^i \tag{2}$$

with n representing the lesser of the two bullet amounts. It is simply the sum of both **A** and **B** both successfully eliminating each other while having perviously missed every single shot.

For the simultaneous situations, the probability of one dying while the other lives is similar to that of both dying. The probability of A winning in a simultaneous duel is

$$\sum_{i=0}^{n} p(1-p)^i(1-q)^{i+1} \tag{3}$$

while the chance for **B** winning in such a duel is

$$\sum_{i=0}^{n} q(1-p)^{i+1}(1-q)^i \tag{4}$$

For the alternating duel, the probabilities of each winning are very similar. In fact, the probability of **B** winning in this scenario is also equal to the summation in 2.1.3(4). The probability of **A** winning is

$$\sum_{i=0}^{n} p(1-p)^i(1-q)^i \tag{5}$$

3 Data

With the large data set and the amount of variables in each scenario, a general analysis is difficult. There are also many trivial cases that are not very interesting. Overall, the situations where both p and q are .5 or above tend to end in a short amount of shots. On the contrary, the situations with .2 or.1 shots end often with a lot of shots taken or end up with both **A** and **B** surviving.

When comparing the Monte Carlo simulation results to the recurrence, the probabilities of each situation tend to match up and validate the results. At 100,000 repetitions of the simulation for each set of probabilities and bullet counts, variation is to be expected, but difference between the simulation probabilities and the recurrence probability distribution is generally under 0.5%.

For instance, the case where both **A** and **B** have 2 bullets and have success rates of 30% and 40% and **A** goes first is supposed to have the following probability distribution according to the recurrence:

A lives, B Dies - 42.6%

A Dies, B Lives - 39.76%

A Lives, B Lives - 17.639999999999997%

(Note slight variation in the total probability due to error with floating point values, the true value is 17.64%)

Meanwhile, the simulation at 100,000 trials yielded the following probabilities:

A lives, B Dies - 42.604%

A Dies, B Lives - 39.7986%

A Lives, B Lives - 17.5974%

For each outcome of the probability distribution, the probability differs from the recurrence probability by less than 0.05%.

A consistent issue throughout the data is the floating point error creating an erroneous probability distribution of less than $\mathbf{10^{-6}}$. These errors are more common in the recurrence formula compared to the simulations, with the increased use of float variables. Generally, the slight variations stemming from these errors is not significant enough to draw misled conclusions.

Each of the 4 results files (Alternating Shot Simulation, Alternating Shot Recurrence, Simultaneous Shot Simulation, Simultaneous Shot Recurrence) holds 10,000 different scenarios can be looked up within the text file using the unique detail code

A:$\mathbf{bul_A}$ bullets $\mathbf{p}$ prob B:$\mathbf{bul_B}$ bullets $\mathbf{q}$ prob

To identify certain patterns, the same generation of the text files can be modified to output only the desired outcome. For instance, the text file writer can be instructed to write only those outcomes where the probability of A winning the duel (only B dies) and B winning the duel (only A dies) are within 10% of each other. This method was used several times to determine certain unique properties within the larger data set.

4 Results

For the purposes of analysis, the probability distribution from the recurrences is used to draw conclusions with generally more accurate data. The data has been referenced with the simulation data and the similarities suggest that the recurrence distributions are accurate.

4.1

In a situation where bullets are not restricted (infinite bullets), the advantage of shooting first or a higher success rate is still not diminished. For example, the probability that both die in a simultaneous duel is similar to that of the equation in 2.1.3(2) with one adjustment to account for the infinite bullets

$$\sum_{i=0}^{\infty} pq(1-p)^i(1-q)^i \tag{6}$$

This also forces at least one player to die by the end of the duel as there is no set point where a duel can end other than a death. There is no situation anymore where both players live by the end of the duel. In fact, every single summation for the probabilities of who wins becomes an infinite sum, without a change to the terms being summed.

4.2

A consistent has the advantage during the alternating duel where **A** is able to take a shot before **B** gets a shot at all. In fact, **A** has the advantage and holds the winning percentage advantage 70.25% of the time. In the simultaneous duel, the advantage is reduced to only 49.04% of the time. Without the advantage of going first, the loss of 21.21% advantage is completely logical. In fact there is no advantage compared to the winning percentage of **B**, who also has a 49.04% chance of having the advantage the duel. This results in a 1.92% chance of neither being more likely to win than the other. These situations all stem from when both players have the same amount of bullets and probability of success at the same time (Neither player has a bullet or success probability advantage) or a 100% success chance regardless of their bullet count (Duel always ends in both dying at the first shot).

4.3

The probability that **A** or **B** just survives rather than wins is simple as well. The probability is the sum of likelihood of one player winning and the likelihood of both surviving the duel. High likelihoods that both survive the duel result in high total survivability percentages, as would be expected.

4.4

Sometimes **A** can have every single advantage (more bullets, higher success chance and the advantage of going first) and yet still win a significant amount of time.

There are 838 different scenarios in the alternating duel where **B** has at least a 10% chance of winning the duel outright (0.838% of all situations). These generally tend to be in situations where **A** has a very low probability of success (30% or less) and the disadvantage is very small (10% difference in success or 1 bullet difference).

When the threshold of **B** winning is raised to 20%, the amount of situations is reduced to 311 (0.313 of all situations). Generally, the same general patterns of **B** being able to reach this threshold is similar to that of the 10% threshold. That would make sense considering that this is a subset of the 10% condition. There is much less leniency of being just slightly worse off compared to **A** in probability and bullet count, with very few of the results having differences in success probability greater than 10%. Difference in bullet count is less restrictive than that of success probability but generally follows the guideline for reaching this threshold.

Once the threshold is raised to 30% win probability, only 36 situations satisfy this condition. Most of these situations rely on high bullet counts and moderately low success rates to minimize the probability that both survive so that is more likely that **B** is able to win in some of the later rounds of fire. The probability is so low that the highest probability of both surviving is 1.89%, in the case where **A** has 8 bullets and a 0.3 probability of success while **B** has 5 bullets and a 0.2 probability of success. The ability of **B** to win beyond 30% of the time while being completely disadvantaged is minimal. The highest probability of **B** winning with the disadvantage is 31.645% where **A** and **B** have 10 and 9 bullets, with 0.3 and 0.2 probabilities of success respectively.

4.5

There are certain conditions to achieve a roughly equal probability distribution across all outcomes within a situation. Using selective text file output, only situations where all outcome probabilities were between 28% and 38% were found within the alternating duel. The range gives leniency with out equal the outcomes are. There are a total of 16 situations that satisfy this criteria. These situations fall within 3 categories.

The first category encompasses all situations where **A** has 1 bullet with a 0.3 success probability while having **B** win around 50% of the time assuming **A** misses, creating a distribution of roughly

30% - A wins

35% - B wins

35% - Both survive

The second category composes of all situations where **A** has 2 bullets with 0.2 success probability. The outcome distribution is balanced through **B** having low bullet counts and success rates.

The last category of is when both **A** and **B** both have 0.1 chances of success with high bullet counts. This minimizes the starting advantage for **A** as much as possible while keeping the probability that both survive roughly equal to the probabilities of each **A** and **B** winning.

4.5.1

There are 756 situations where both A and B win roughly 50% of the time each (Probability of winning each between 44% to 56%). There is less of a pattern for these cases, with the only notable one being that **A** has a success probability between 0.2 and 0.5 to ensure that the chance that **A** wins on the first shot is not too great.

4.5.2

There are only 7 cases in the simultaneous case where all four outcomes are roughly equal. Probabilities were restricted between 0.15 and 0.35) This case is generally difficult to achieve as so many different variables need to be roughly in balance. Generally, these situations needed low amount of bullets and roughly equal probabilities of success to ensure each outcome had roughly similar probabilities.

5 Future Work

The most logical progression to this 2-player duel is to extend the situation to 3 players. In this circumstance, the decision of how the 3 players interact needs to be considered. The players could shoot in a circle, where **A** shoots **B**, **B** shoots **C** and **C** shoots **A**. They players could also randomly shoot between the other two players. If at any point the duel gets reduced to 2 players, the duel continues in a way identical to the 2 person duel until the amount of remaining bullets is reduced to 0 or another player is eliminated. The extra player creates much more depth with 8 different outcomes in the simultaneous duel where all players shoot at once. The amount of situations in the 3-player duel also increases to 1 million. Drawing coherent conclusion from such a massive data set is likely to be even more difficult than the 2-player duel.