Examiner rated this a B

Methods of Maximization of Rewards with Predetermined Random Lists of Rewards

Given a hidden list of "n" items what methods of maximization of reward selections exist given that once an item is examined, it must either be selected or rejected immediately?

Mathematics

Word Count: 2923

Table of Contents

Introduction

The Sultan dowry problem is a quite simple problem, but its solution is rather complicated. The problem states: “A sultan has granted a commoner a chance to marry one of his *n* daughters. The commoner will be presented with the daughters one at a time and, when each daughter is presented, the commoner will be told the daughter's dowry” (Weisstein). The problem by itself can be answered quite simply, but its existence can provide insight into more generalized questions like given a hidden list of "n" items what methods of maximization of reward selections exist given that once an item is examined, it must either be selected or rejected immediately That can provide insight into how we make decisions. Upon being presented with a daughter, the commoner must immediately decide whether to accept or reject her. However, the sultan will allow the marriage to take place only if the commoner picks the daughter with the overall highest dowry. Then what is the commoner's best strategy, assuming he knows nothing about the distribution of dowries.” (Weisstein). The most important part of this problem as stated is that the commoner does not know anything about the distributions of the dowries. Although this makes the problem slightly less realistic, as it is obvious there would be some sort of maximum on the value that the dowry could be, for the problem we will assume that the commoner has zero way of knowing what the possible range of dowries could be.

The Sultan’s Dowry Problem

The best possible strategy that the commoner can use to maximize his probability of choosing the highest dowry involves waiting for *x* daughters and then afterwards choosing the next daughter with a dowry higher than one seen before (Weisstein). Throughout the paper, we will use many different variables for many different uses, although we will use many common variables. Some of these common variables will be as follows:

We will let X represent the number of daughters that will be passed before the commoner will pick the next highest daughter.

We will let N represent the total number of daughters that will be presented to the commoner.

We will let P represent the probability that the commoner will select the daughter with the highest dowry.

In this situation, the commoner will know nothing about the distribution of dowries, and as a result the only strategy he may take is one that compares dowries to the ones he has already seen. Furthermore, it is never implied in the question that the distribution will follow some sort of pattern, so while if we were dealing with a normal distribution there may be some ways we can use that knowledge to increase our odds, instead we must simply compare numbers to each other using operators like greater than or less than. As stated above, the commoner should have some value of X that after which he will select the next daughter with a dowry higher than any he had seen before. This value of X can be found by finding the lowest value of X such that the probability of the highest dowry has already been seen is just greater than the probability that the highest dowry has not been seen and it will be picked.

The value of this inequality, the probability that the highest dowry has already been seen, is rather simple to calculate. It can be found by dividing the number of dowries that have already been seen by the total number of daughters, or (X/N)

The second value of this inequality—The probability that the highest dowry has not been seen and it will be picked is much more difficult to calculate but is simpler when it is examined in chunks.

To start, we will calculate the probability that the 2nd highest daughter is the highest daughter that is seen in the initial set. This value can be found by finding the probability that the 2nd highest daughter is in the set of daughters first seen, and multiply it by the probability that given the 2nd highest daughter is in the set, no daughters higher than the 2nd is also present. This gives us

$$\frac{\left(\frac{X}{N}\right)(N-X)!\,(N-1)^{x-1}}{(N-1)!}$$

Finally, we must divide this value by 1, the probability that the highest daughter will be the first one high enough in the list after the turning point to be selected.

We will continue this with examining the situation where the 3rd highest daughter in the list is the highest that will be selected. This leaves us with the following probability.

$$\frac{\left(\frac{X}{N}\right)(N-X)!\,(N-2)^{x-1}(N-2)!}{2}$$

There many of these terms can be factored out of this equation, and luckily the rest of the equation simplifies very well. In the end, after iterating through all the potential highest daughter numbers, we are left with the following equation

$$\left(\frac{X}{N}\right) * \left(\left(\frac{1}{X+1}\right) + \left(\frac{1}{x+1}\right) + \cdots + \left(\frac{1}{N}\right)\right)$$

Finally, looking at the entire equation, we must find the smallest possible value such that

$$\left(\frac{X}{N}\right) \geq \left(\frac{X}{N}\right) * \left(\left(\frac{1}{X+1}\right) + \left(\frac{1}{x+1}\right) + \cdots + \left(\frac{1}{N}\right)\right)$$

Or, to further simplify it,

$$1 \geq \left(\left(\frac{1}{X+1}\right) + \left(\frac{1}{x+1}\right) + \cdots + \left(\frac{1}{N}\right)\right)$$

This can be solved easily by a calculator or a computer, and so we will generate a list of some of the possible values of N and their corresponding values of X, and analyze this for any trends that may become apparent

N	X	X/N
5	2	0.4
7	3	0.429
10	4	0.4
15	6	0.4
20	8	0.4
40	15	0.375
80	30	0.375
100	37	0.37
200	74	0.37
500	184	0.368
1000	368	0.368

This chart shows a few interesting insights. First, it shows that the correlation between the total number of daughters and the number of daughters the commoner should observe before being ready to make the decision to pick the next highest dowry he sees is a linear relationship, and a directly proportional one as well. This means that we can represent this relationship as a percentage. The commoner should wait until he has seen about 36.8%, or exactly 1/e of the daughters, rounded to the nearest daughter, until he should then pick the next highest dowry he sees.

This also gives us further insight into how probable it is that the commoner is able to select the daughter with the highest dowry. Because the ratio of daughters seen before the commoner has the optimal amount of information to make the decision is 1/e, we can now plug this value in for the minimum possible value for the first part of the inequality used to determine when the commoner should make the decision. The inequality used is the following:

$$\left(\frac{X}{N}\right) \geq \left(\frac{X}{N}\right) * \left(\left(\frac{1}{X+1}\right) + \left(\frac{1}{x+1}\right) + \cdots + \left(\frac{1}{N}\right)\right)$$

Specifically, we are interested in the expression on the left half of this inequality. This expression can be simplified significantly. The (X / N) expression can be simplified to 1/e, as we just established that 1/e will be approximately the ratio of daughters needed to be seen.

Furthermore, another inequality already used is useful to understanding this. The following is a simplified version of the inequality above.

$$1 \geq \left(\left(\frac{1}{X+1}\right) + \left(\frac{1}{x+1}\right) + \cdots + \left(\frac{1}{N}\right)\right)$$

The term on the right is the same as the remaining term we wish to simplify. The final simplified expression, displayed as an inequality, provides insights into the commoner's situation.

$$P \geq \left(\frac{1}{e}\right)$$

Shockingly, the commoner has a large probability of selecting the highest daughter—at 1/e or about 36.8% at the minimum, regardless of how many total daughters he will be presented with.

Variation One—Known Distribution

The sultan's dowry problem is an interesting mathematical problem, however the lack of information supplied to the commoner hinders its ability to apply to real life as well. To continue our examination of this topic, we will analyze variations to the problem that may have more real world applications. Because we will no longer be discussing the actual sultan's dowry problem itself, but rather some variations of the problem, it will be helpful to view the problem differently, although fundamentally the same. Throughout the two variations to the problem that we will discuss, we will examine the problem as a game where a guesser attempts to pick a card out of a stack with a predetermined number of cards in it, each with a number written on it.

For the first of the variations to the sultan's dowry problem, we will change the amount of information available to the guesser about the distribution of the numbers. Instead of the guesser knowing nothing about the distribution of the numbers, he will instead have full knowledge of the possible values that the cards could contain, and the probability of each of those numbers. We will also assume that the numbers are infinitely precise, so that two numbers could never be the exact same. While it matters that the guesser has complete knowledge of how distribution that the cards belong to, as for our analysis all that is important is the ranks of the cards, so we can use the percentile of the number on the card. It does not matter if the distribution is normal, skewed, uniform, or some other irregular shape, as for the purpose of this variation all distributions can be condensed to a uniform distribution between 0 and 1 representing the percentile of the number on the card.

As a result of us knowing the distribution of the outcomes at the start, it is no longer necessary to wait for a period of time to gather information about the distribution. Instead, the optimal strategy will be to wait until we see an outcome with a high enough percentile, and then select it. The probability that the highest outcome is selected can be found by multiplying the probability that various numbers of the set could get picked, multiplied by, for each, the probability that the highest of them would get picked. Once again, we will assume that two items in this list cannot have the same exact number associated with it.

To find what the appropriate percentile would be to select the value, we will use a similar principle to before. Like last time, we will evaluate the probability that, for any given percentile of value given, the odds of the value being the highest in the set, and the probability that in the rest of the set, there will be a higher value. Because these events are mutually exclusive and collectively exhaustive, we only need to solve for what value makes these odds 50% or greater. We will solve for the following:

$$0.5 \geq X^n$$

Where X is the percentile of the item selected, and n is the number of remaining options to pick. Interestingly, this also simplifies incredibly well, to give us the result that

$$X = \left(1 - \left(\frac{1}{n}\right)\right)$$

The following chart shows what the threshold looks like for many possible values of n, and what the probability of obtaining the highest value is.

N	Threshold	P
2	0.5	0.625
3	0.667	0.568
5	0.8	0.531
10	0.9	0.507
20	0.95	0.494
50	0.98	0.490
100	0.99	0.487
200	0.995	0.485
500	0.998	0.483

This chart and graphs show some interesting trends also. First of all, they show that, as expected based off of the earlier calculation, the value of a number has to be greater than its percentile for it to be selected.

This method is improvable, however, when understanding that an initial group of 10 possible values where the guesser has already seen the first of the possible values is almost no different than a group of 9 possible values where the guesser has not seen any of the possible values. The only difference here is that for the group of 10 initial values, one of them is already known. Therefore, we can change this into a 2-step process to determine whether the guesser should pick the number presented to them.

The guesser should select the value if

1. The number presented is greater than all of the numbers already seen.

And

2. The number presented is greater than (1 - (1 / k)) where k is the total number of values left to seen

The difference this change makes is small, but the change is significant. This change can be seen in the chart below

N	P
2	0.625
3	0.590
5	0.586
10	0.583
20	0.579
50	0.575
100	0.573
200	0.571
500	0.569

As the charts show, the second method of selecting the highest value is much better than the first, with the probability of an N value of 500 is about 8% more than the probability using the second method. It is also extremely better than the situation in the sultan's dowry problem, where given the information about the possible distribution of the values/dowries.

Variation Two—Maximizing Averages

One final aspect of this situation is to discuss how to maximize the average outcome. This has many real worlds uses, including obtaining, on average, the best deals when shopping.

Unlike the last situation, however, this situation relies on the distribution being one specific way so we will only analyze the normal distribution. Like in the other problems, however, all normal distributions have essentially the same features, so we will specifically be analyzing a uniform distribution between 0 and 1. This allows us to use many of the same methods as in the part where we analyzed the probability of picking the greatest value given a known distribution.

This situation is distinct from the other variations in one major way. It does not matter what has already been seen, but all that matters is the current value of the item and how many items left.

For any situation, the decision to either pass an item or keep it can be determined by comparing the expected value of keeping the current item to the expected value of playing out the rest of the game. Calculating the expected value is easy for the current item—just whatever the number is—but calculating the expected value for the whole game proves to be a more difficult part. Calculating this for a game with one is easy—as it is simply taking the average of all of the possible values for that one value, giving us an expected value of 0.5.

For a game with 2 cards, the player will either take the first card if its expected value is greater than 0.5 or take the second card otherwise. The first card will be selected 50% of the time with an average expected value of 0.75 when it is selected, or the second card will be selected—again with a probability of 50%--and will have an expected value of 0.5. Combined, this gives us an expected value of 0.625 for a game with 2 cards

Finally, the last game we will examine will be a game with three cards. In this situation, the player can either decide to take the first card presented to them, with a known value, or instead play the game with 2 cards, with a known expected value of 0.625. In a similar way as in the last two times, we can determine that the expected value of the game with three cards will be 0.6953125.

Based on these first three games, a recursive formula becomes apparent. To find the expected value for a game with N cards played by a perfect player, all one needs to know is the expected value of a game with N-1 cards in it. The formula to calculate this is as follows, where N represents the total number of cards in the game and E represents the expected value of a game.

$$E(N) = \frac{\big(1 + E(N-1)\big)\big(1 - E(N-1)\big)}{2} + E(N-1)^2$$

Or, in a more simplified form,

$$E(N) = 0.5(1 + E(N-1)^2)$$

The values for these games can be seen in the following chart.

N	Expected Value
1	0.5
2	0.625
3	0.6953125
5	0.7750815
10	0.8610982
20	0.9198874
50	0.9641452
100	0.9812084
1000	0.9980172
10000	0.9998002

There appears to be at least somewhat of an inverse relationship between the total number of cards in a game and the expected value of playing the game. While other relationships appear to exist at smaller values, these relationships seem to fall apart at larger values of N because, according to the chart, the end behavior of this relationship is to approach 1 as N approaches infinity, which does not share in end behavior with the many patterns suggested by the early data.

A similar method could be applied to other types of distributions as well. By comparing the value of the current card to the expected value of playing out the rest of the game, one can determine whether the decision to keep the current card or continue with the game would lead to a higher expected value.

Conclusion

This essay explores the sultan's dowry problem, a problem where a commoner is presented with "N" daughters one at a time, each with an associated dowry, and upon being presented with a daughter, must immediately decide whether to marry that daughter or not, knowing that he will only be allowed to marry if he has selected the daughter with the highest dowry. The essay further discussed the details of some variations to this problem, one where the commoner has the information about the possible distribution of such dowries, and another where the commoner aims to maximize the average dowry received without the restriction that he may only marry if he selects the daughter with the highest dowry. For further research, it would be interesting to extend the analysis on variation two to include more types of distributions and analyze the differences between them. Further, it would be interesting to modify the original problem to allow the commoner to marry any of the top, for example, three daughters, and observe if any new patterns become apparent.

References

Weisstein, Eric W. "Sultan's Dowry Problem." From *MathWorld*--A Wolfram Web Resource. https://mathworld.wolfram.com/SultansDowryProblem.html