Twin Treasures

[Year 2023–2024]

FIRST NAMES AND SURNAMES OF STUDENTS, GRADES: AFTENE BOGDAN GEORGE, ALBERT MATEI , MACSIM MIHAI, PASCARI TUDOR STEFAN, STIRBU ISACHE MARIO MIHAI, TANASA RAZVAN.

SCHOOL: Colegiul Naţional *C. Negruzzi*, Iaşi

TEACHER: ADRIAN ZANOSCHI, Colegiul Național *C. Negruzzi* and RAZVAN-ANDREI MORARIU, UNIVERSITY of Vienna

RESEARCHER: RAZVAN-ANDREI MORARIU, University assistant at University of Vienna

1. PRESENTATION OF THE RESEARCH TOPIC

Decision-making problems lie at the heart of many real-world challenges, from optimizing resource allocation to choosing the best strategy in uncertain situations. Probability theory provides a powerful framework for tackling these problems, offering tools to model uncertainty and quantify risks. By combining probabilistic reasoning with strategic thinking, we can search optimal solutions to complex scenarios, even with incomplete information. These problems combine mathematics, logic, and intuition, making them both fascinating and practical.

2. BRIEF PRESENTATION OF THE CONJECTURES AND RESULTS OBTAINED

Bob, a mathematician, is captivated by the myriad applications of mathematics in games. Consequently, he devised the following straightforward game, which appears to possess a far more intricate solution. The game is called ***The Twin Treasures***. During each turn, Bob has the option to open one treasure and get the reward that is inside. Once he closes the treasure, the treasure magically creates a new reward inside. Each treasure yields a reward according to the following criteria:

- It wins $(reward = 1)$ with a fixed (unknown) probability P_i.
- It loses $(reward = 0)$ with a fixed (unknown) probability $1 - P_i$.

There will be two type of games. In the first case, Bob does know the probabilities but not to which treasures correspond. In the second case, he does not know the probabilities at all. The final objective of the game is to determine, given the total number of times, T, Bob can open the treasures, how many times he should open each of the treasures in order to maximize the expected reward. In order to find the optimal strategy Bob tries to answer the following questions:

1) What is the expected reward if the first treasure is opened for T times?
2) What is the expected reward if the first treasure is opened for k times and the second for $(T-k)$ times?
3) For $\mathrm{T}=3$, $P_1=0.7$, $P_2=0.3$ is there a strategy for opening the treasures that has a higher expected reward than the expected reward of opening randomly one treasure for 3 times?
4) Prove that for any T, any $P_1>P_2$ the expected reward for a strategy of opening each treasure k times and the better one afterwards is better than randomly opening one treasure for T times for any $0<k<\frac{T}{2}$. In the case of an equal number of wins in the first $2k$ openings, the better treasure will be chosen randomly.
5) What is the expected reward if Bob opens each treasure k times and then opens only the one with the higher total reward for $\mathrm{T}-2k$ times? How can this reward be computed?
6) Can we make an informed decision for the current strategies about the optimal value of k when the probabilities P_1 and P_2 are known, but it is unclear which treasure corresponds to each probability? Furthermore, can we determine the best value of k when the probabilities themselves are entirely unknown?
7) Since the current strategy does not lead to finding the optimal value for k, how can we improve the current strategy?

3. SOLUTION

Let us enumerate the lemmas used in our solution.

Lemma 1: Let X and Y be two discrete random variables. Then: E[X+Y] = E[X]+E[Y].

Proof:
The expected value of the sum of two random variables is defined as:

$$E[X+Y] = \sum_{x,y} (x+y)P(X=x, Y=y).$$

One can get

$E[X+Y] = \sum_{x,y} xP(X=x, Y=y) + \sum_{x,y} yP(X=x, Y=y)$.

Since

$\sum_{y} P(X=x, Y=y) = P(X=x)$ and $\sum_{x} P(X=x, Y=y) = P(Y=y)$ one can obtain

$E[X+Y] = \sum_{x} x\,P(X=x) + \sum_{y} y\,P(Y=y)$, which leads to:

$$E[X + Y] = E[X] + E[Y].$$

Lemma 2: Let $X_1 \ldots X_n$ be n random variables. Then: $E[X_1 + \cdots + X_n] = E[X_1] + \cdots + E[X_n]$

Proof: It is enough to use Lemma 1 inductively.

Lemma 3: Let $X_1, X_2, X_3, \ldots, X_n$ be n random variables independent and identically distributed. Then

E[$X_1 + \ldots + X_n$] = nE[X_1].

Proof: Use Lemma 2

Solutions

Question 1

If the first treasure is opened for T times, the rewards from these trials are independent and identically distributed (i.i.d.) random variables with

$$E[X_1] = 1 \cdot P_1 + 0 \cdot (1 - P_1) = P_1,$$

where X_1 is the random variable that describes the reward in one opening of treasure 1.

Hence, by Lemma 3, the total expected reward from T openings of treasure 1 is $T \cdot P_1$.

Question 2

We will use the result derived in Question 1. If the first treasure is opened for k times and the second treasure for $(T-k)$ times, the total reward is:

$$k \cdot E[X_1] + (T-k) \cdot E[X_2],$$

where X_1 is the random variable that describes the reward in one opening of treasure 1 and X_2 is the random variable that describes the reward in one opening of treasure 2.

Hence, the expected total reward is:

$$k \cdot P_1 + (T-k) \cdot P_2.$$

Question 3

Recall that $P_1 = 0.7$ and $P_2 = 0.3$. (1)

The possible outcomes from the first two turns and their probabilities are:

1. Both treasures yield 0: $P(X_{11} = 0, X_{21} = 0) = (1 - P_1)(1 - P_2) = 0.3 \cdot 0.7 = 0.21$.

2. Treasure 1 yields 1 and Treasure 2 yields 0: $P(X_{11} = 1, X_{21} = 0) = P_1(1 - P_2) = 0.7 \cdot 0.7 = 0.49$.

3. Treasure 1 yields 0 and Treasure 2 yields 1: $P(X_{11} = 0, X_{21} = 1) = (1 - P_1)P_2 = 0.3 \cdot 0.3 = 0.09$.

4. Both treasures yield 1: $P(X_{11} = 1, X_{21} = 1) = P_1 P_2 = 0.7 \cdot 0.3 = 0.21$.

The expected reward is computed for each case:

1. Case 1: Both rewards are 0. Bob opens randomly for the third turn, so $E[\text{Reward}] = 0.5 \cdot P_1 + 0.5 \cdot P_2 = 0.5 \cdot 0.7 + 0.5 \cdot 0.3 = 0.5$

2. Case 2: Treasure 1 yields 1 and Treasure 2 yields 0. Bob opens Treasure 1 for the third turn, so $E[\text{Reward}] = 1 + P_1 = 1 + 0.7 = 1.7$.

3. Case 3: Treasure 1 yields 0 and Treasure 2 yields 1. Bob opens Treasure 2 for the third turn, so $E[\text{Reward}] = 1 + P_2 = 1 + 0.3 = 1.3$.

4. Case 4: Both rewards are 1. Bob opens randomly for the third turn, so $E[\text{Reward}] = 2 + 0.5 \cdot P_1 + 0.5 \cdot P_2 = 2 + 0.5 \cdot 0.7 + 0.5 \cdot 0.3 = 2.5$.

Now, the total expected reward is:

$E[\text{Total Reward}] = 0.21 \cdot 0.5 + 0.49 \cdot 1.7 + 0.09 \cdot 1.3 + 0.21 \cdot 2.5 = 0.105 + 0.833 + 0.117 + 0.525 = 1.58$.

This expected reward 1.58 is higher than the random opening strategy, which yields an expected reward of $3 \cdot (0.3 + 0.7) \cdot 0.5 = 1.5$. (2)

Question 4

Firstly, let us analyze the case of opening both treasures $\frac{T}{2}$ times. Using Lemma 2, the total expected reward will be:

$$E[\text{Total Reward case 1}] = \frac{T}{2}(E[X_1] + E[X_2]) = \frac{T}{2}(P_1 + P_2).$$

In the second case, the expected reward can be divided into 2 parts. Before analyzing this, let us denote by C_1 the random variable that describes the total win for the first k openings of treasure 1 and by C_2 the total win for the first k openings of treasure 2.

The expected reward for the first $2k$ openings is $k(P_1 + P_2)$.

The expected rewards for the remaining $T - 2k$ openings can be obtained by considering the probabilities of getting more wins for treasure 1, more wins for treasure two, or the same number of wins. This leads to the following expected reward $(T - 2k)\left(P_1 \cdot P(C_1 > C_2) + P_2 \cdot P(C_2 > C_1) + \frac{P_1+P_2}{2} \cdot P(C_1 = C_2)\right)$

Hence, $E[\text{Total Reward case 2}] = k \cdot (P_1 + P_2) + (T - 2k)\left(P_1 \cdot P(C_1 > C_2) + P_2 \cdot P(C_2 > C_1) + \frac{P_1+P_2}{2} \cdot P(C_1 = C_2)\right)$.

Since $P_1 > P_2$ involves (3) $P(C_1 > C_2) > P(C_2 > C_1)$, it means that $P_1 \cdot P(C_1 > C_2) + P_2 \cdot P(C_2 > C_1) + \frac{P_1+P_2}{2} \cdot P(C_1 = C_2) > \frac{P_1+P_2}{2}$, which implies that

$E[\text{Total Reward case 2}] > E[\text{Total Reward case 1}]$.

Question 5

We will denote by $X_1, X_2, \dots, X_k$ the random variables that describe the wins in each of the first k openings of treasure 1 and by $Y, Y_2, \dots, Y_k$ the corresponding random variables of treasure 2.

Let us denote by $C_1 = X_1 + X_2 + \dots + X_k$ and $C_2 = Y_1 + Y_2 + \dots + Y_k$.

As before, we will denote by P_1 the probability to win when we open treasure 1 and P_2 is the probability to win when we open treasure 2.

We will now enumerate the steps for computing the expected reward for all possible values of k. After following these steps, we can find which value of k leads to the highest expected reward.

First step: We will compute the probability for getting i wins from the first k openings for C_1 or C_2, where i ∈{0, …, k}. A pseudo code that does this is the following:

$$for\ (int\ i\ =\ 0;\ i\ \le k;\ i++)\ \{$$
$$C_1[i] = \binom{k}{i} \cdot P_1^i \cdot (1 - P_1)^{k-i}$$
$$C_2[i] = \binom{k}{i} \cdot P_2^i \cdot (1 - P_2)^{k-i}\}$$

Second step: We compute the probability of getting $C_1 = C_2$

$$P(C_1 = C_2) = \sum_{i=0}^{k} C_1[i]C_2[i] \coloneqq P_{12}$$

In this case we will randomly choose which treasure we consider to be better and open that for the remaining number of times.

Third step: We compute the probability of getting $C_1 > C_2$

$$\sum_{i=1}^{k} C_1[i]\left(\sum_{j=0}^{i} C_2[j]\right) \coloneqq P_{11}$$

One can observe that the term inside the parenthesis can be seen as a cumulative distributive function evaluated at different values. However, this cdf does not have an elementary closed form, so numerical computation is needed.

Fourth step: We compute the probability of getting $C_2 > C_1$ for the first k openings of each:

$$P_{22} \coloneqq 1 - P_{11} - P_{12}$$

Fifth step: We compute the expected reward for the first k openings of each.

$$E[X_1 + X_2 + \cdots + X_k] = kE[X_1] = kP_1$$
$$E[Y_1 + Y_2 + \cdots + Y_k] = kE[Y_1] = kP_2$$

Sixth step: We'll compute the expected reward for $T - 2k$ openings of the better one.

Let Z be the random variable that describes the wins from the treasure that is observed as being the better after k openings of both of them. So, Z will be X_1 with probability $\left(P_{11} + \frac{P_{12}}{2}\right)$ and Y_1 with probability $\left(P_{22} + \frac{P_{12}}{2}\right)$.

Consequently, the expected win for the last $T - 2k$ openings is

$$(T - 2k)E[Z] = (T - 2k)\left[\left(P_{11} + \frac{P_{12}}{2}\right)E[X_1] + \left(P_{22} + \frac{P_{12}}{2}\right)E[X_2]\right]$$

Seventh step: We compute the total expected reward:

$$kP_1 + kP_2 + (T - 2k)\left[\left(P_{11} + \frac{P_{12}}{2}\right)P_1 + \left(P_{22} + \frac{P_{12}}{2}\right)P_2\right] = E[W]$$

We implemented this code in C++ and used it on three different set of parameters to check which is the optimum number of openings for different values of T, P_1 and P_2. The plots (crated in python) show the expected value for all possible choices of k:

Plot for $T = 100$, $P_1 = 0.55$, $P_2 = 0.5$

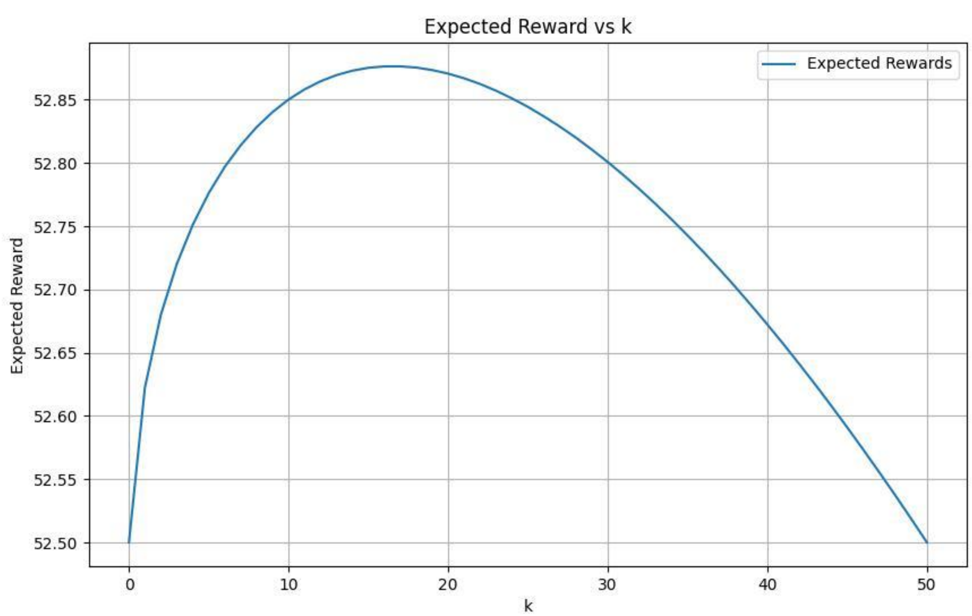

Plot for $T = 100, P_1 = 0.7, P_2 = 0.5$

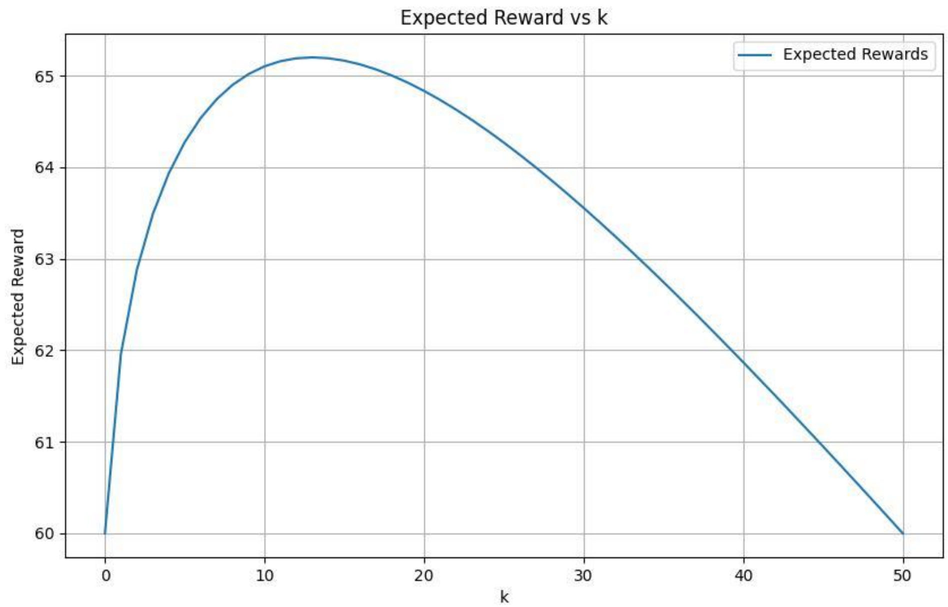

Plot for $T = 100, P_1 = 0.7, \ P_2 = 0.4$

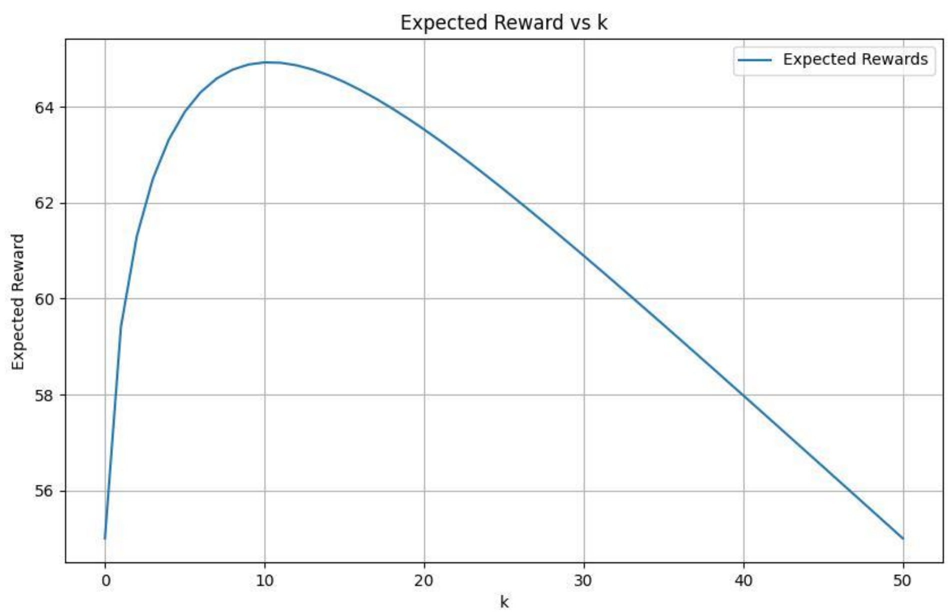

The code can be found at https://github.com/Razvan8/Twin-Treasures.git.

Question 6

Since we showed in Question 5 that we can compute the expected reward when T, P_1, P_2 and k are given, we can choose the best k if we know T, P_1, P_2 . However, if P_1 and P_2 are not known, we cannot find the best k for that specific case. However, it is clear that any k would lead to a higher expected reward than opening a treasure randomly due to the analysis done in Question 4.

Question 7

Firstly, let us state again that in the case where the probabilities of winning for the two treasures are known (even without knowing which treasure is better), the strategy described in the previous point can compute the optimal number of openings. Therefore, this method leads to finding the best k in this scenario. (4)

However, when P_1 and P_2 are unknown, the method from the previous point cannot be applied directly. It is evident that exploring both treasures is necessary to estimate which one is better (as discussed in Question 4), but determining the optimal value for k is not straightforward. A more sophisticated approach might therefore yield better results.

For example, one could initially open each treasure a small number of times to gather data and decide which treasure appears more profitable. Subsequently, instead of exclusively opening the treasure with the higher observed reward, it may be advantageous to open it more frequently while still occasionally opening the other treasure. This strategy would reduce the risk of mistakenly favoring the wrong treasure due to early uncertainty. The greater the observed advantage of one treasure in terms of wins, the more frequently it should be opened compared to the other.

However, the theoretical development of these ideas is beyond the scope of this discussion and is left for future work.

4 CONCLUSION

In the Twin Treasures problem, we employed basic probability theory, logical reasoning, and intuition to navigate the complexities of uncertainty and decision-making. By modeling the problem mathematically, we analyzed strategies to maximize rewards despite incomplete information and proved that there are strategies that lead to higher expected reward than random strategies.

Notes d'édition

(1) The proposed strategy consists in opening once each of the treasures during the first two turns, and for the third turn following the case study presented at last.

(2) Observe that the strategy of opening three times treasure 1 yields an expected rewards of 3*0.7=2.1.

(3) For k=1 we have four possible outcomes, with $P(C_1>C_2) = P_1*(1-P_2)$ and $P(C_2>C_1) = P_2*(1-P_1)$. From $P_1>P_2$ we deduce $(1-P_2)>(1-P_1)$ hence $P(C_1>C_2) > P(C_2>C_1)$. For k=2 we have 16 possible outcomes, with $P(C_1>C_2) = P_1^2*(1-P_2)^2+P_1^2*(1-P_2)+P_1*(1-P_2)^2$ and symmetrically for $P(C_2>C_1)$. From $P_1>P_2$ and the symmetry we deduce again $P(C_1>C_2) > P(C_2>C_1)$. The same reasoning applies for all k.

(4) These optimal k are of course the abscissa corresponding to the maxima of the bell-shaped curves.