Department of Economics and Finance

Chair of Gambling: Probability and Decision

Strategic Decision-Making Under Uncertainty: The Secretary Problem and The Prophet Inequality

Supervisor:
Prof.
Hlafo Alfie Mimun

Candidate:
Jie Kessy Lin
270001

Academic Year 2023/2024

Alla mia famiglia e a Michele.

Contents

Introduction

Decision-making problems are an everyday scenario and one of the most intriguing areas of study in the field of probability theory and behavioural science. The study involves analysing how individuals assess new data to make informed decisions. Each of us have faced this problem at least once and asked ourselves these kind of questions: How do we make the best decision? Is there an optimal strategy we can use? However simple the description of the decision-making scenarios may seem, it is a difficult task to find the most efficient and profitable solution to them. In this thesis, we will present a number of valid solutions using probability theories and mathematical computations.

To begin, the following examples will serve to illustrate the concept of decision-making problems. Imagine you are driving and looking for a parking spot: you are happiest if you are able to park the car closest to where you need to go. In fact, the longer you will have to walk once getting off the car the more you will feel discontent about your choice. Note that the parking spots available are visible to us one at a time and once we move past one, we cannot go back. Another example is the following game. You have n treasure chests in front of you, each one containing a cash prize that you can win. You can pick a chest and win the amount of cash that you find inside of it. However, at the very start of the game you have zero information about the prizes contained in each chest. It is the host's job (who is running the game) to open the chests for you, one at a time. Then, every time a chest is opened you must make a choice: you can either accept or reject the prize. if you reject the prize, it is lost to you forever and the game will continue; if instead you accept the prize, you end the game and earn the amount you found inside the chest. How do you maximise your expected winnings? How would you play this kind of games?

The irrevocable and time-critical decision under the uncertainty of possible future events is what characterizes the classical online selection problem. Moreover, in online selection problems, the random variables arrive in a sequential order, one at a time, and reveal their value each at their turn. We begin Chapter 1 with a theoretical background on the optimal stopping rule, or early stopping, to help the reader understand the theory of online selection and follow the resolution of the problem. In particular, in the first chapter, we are defining the optimal stopping theory as a tool used in decision-making scenarios to maximise an expected payoff or minimise an expected cost by choosing the timing of an action. What we focus on is the probability of picking a time N, at which the observation is stopped, that maximises the expected return.

We then move on to the analysis of two important problems in online selection, the Secretary Problem and the Prophet Inequality. In both problems, an agent is asked to make a decision and pick a value from a sequence of random variables. The rule is simple: the agent can only accept one value from the sequence and it cannot be one that has been previously rejected.
The Secretary Problem aims at maximising the probability of selecting the highest value while the Prophet Inequality attempts to compare, in terms of expectations, the performance of the gambler with the one of a prophet who knows the values of the variables in advance.

Right after section 1.1 on the optimal stopping rule, we start with our analysis on the Secretary Problem. The situation is exactly like the one we described for

online selection problems. The decision-maker is now a boss who has to hire a new worker from n applicants and our objective is to find the optimal strategy which makes the boss select the best applicant for the position. In this chapter, Ferguson's work is central to our investigation of this problem [4].
The strategy that he used was to reject the first $r-1$ applicants and choose the next one who was ranked the best relative to the past observations. We demonstrated that the optimal stopping occurs at about 37%. Then, we observed that for a large number of applicants, the probability that the strategy successfully picks the best applicant is also equal to 37%. To confirm this number, we made a simulation on MATLAB and derived Figure 1.1 from which we can see that the peak is achieved exactly at around 0.4 like we expected from our computation.

We dedicated the rest of the chapter to the analysis of a different version of the Secretary Problem. To be more precise, we used the concept of random walks to define what a random walk stock exchange market is and to show a connection with the Secretary Problem for the random walk. This would give the reader a more real-life application of what we examined so far on the decision-making problem. We heavily relied on both M. Hlynka and J.N. Sheahan [6] and Fama's works [3] for the last part of our examination on the Secretary Problem.
In Proposition 1.4.4 we were able to show that, differently from the classical view, in this new version of the problem the probability of success of the strategy would tend to zero as n tends to infinity. Therefore, as n tends to infinity, the optimal strategy would be to pick the first element of the sequence. An interesting note on our calculation of this probability is that it is surprisingly better than $1/n$. In fact, as we derived in section 1.4.1, the probability of picking the maximum value under strategy S_k is equal to $\sqrt{\frac{2}{\pi n}}$, where by S_k we mean the strategy of picking a value k for $k = 0, ..., n$.
Finally, we defined the random walk market as "*a market where successive price*

changes in individual securities are independent [3]" and proved the efficiency of the random walk model over the technical approaches.

In the second chapter we focused on the Prophet Inequality. We based this chapter on the research of Krengel, Sucheston and Garling in 1978 [7]. Our objective was to prove that the prophet can gain at most twice as much payoff as the gambler. So in this scenario we have two agents: a gambler and a prophet. While the prophet has knowledge on the values of the variables before they are revealed to him, the gambler makes his choice on whether picking or rejecting a value as it arrives to him. In other words, the gambler can gain information on their value only after he has observed them. In this chapter we proved the central theorem of the Prophet Inequality for independent random variables (Theorem 2.0.1):

$$\mathbb{E}[\max_i X_i] \leq 2\mathbb{E}[X_\tau]\,.$$

This is known as the first of many prophet inequalities in the optimal stopping theory and it aims at comparing the performance of online and offline algorithms for problems that involve selecting one or more elements from a random sequence.

On the following section of the same chapter we deal with a variation of the classic prophet inequality. In this case, the comparison is made between the performance of a prophet with complete foresight and that of a gambler who is charged a negative fixed cost for each observation. Our goal is the maximisation of the following difference:

$$\mathbb{E}[\max_{1\leq i\leq n} Y_i] - \mathbb{E}[Y_{sn}].$$

In conclusion, using the theory from [1] we introduce a prophet inequality in a difference form.

We hope that by the end of this work, the readers will be able to understand and apply the strategies to solve decision-making problems more efficiently. Additionally, our aim for this thesis is to provide a clearer and more comprehensive approach on how to make optimal choices regarding stopping times. We conclude this brief introduction by encouraging the readers to engage with this work with both curiosity and an open perspective.

Chapter 1

The Secretary Problem

The Secretary Problem or the Marriage Problem appeared in the late 1950's and early 1960's and gradually extended to now become a field within mathematics-probability-optimisation. Many probabilists and statisticians tried to develop and, at the same time, solve the problem. One of the firsts were Lindley (1961) Gilbert and Mosteller (1966), and later Freeman (1983), Samuels (1985), and Tamaki (1986). Simply put, the problem was described as one of picking the maximum value of a sequence of n independent variables when no recall is allowed. We can use Thomas S. Ferguson's list of features of the Secretary Problem to have a clearer view of the scenario [4]:

1. There is one secretarial position available.

2. The number n of applicants is known.

3. The applicants are interviewed sequentially in random order, each order being equally likely.

4. You can rank all applicants without ties. The decision to accept or reject must be based only on the relative ranks of those applicants interviewed so far.

5. An applicant once rejected cannot later be recalled.

6. Your payoff is 1 if you choose the best and 0 otherwise.

For now, we will focus on the simplest form of the Secretary Problem and disregard all its other variations. We can rewrite the features listed above and consider a more formal definition of the problem such as the one Ferguson used: "*A Secretary Problem is a sequential observation and selection problem in which the payoff depends on the observations only through their relative ranks and not otherwise on their actual values* [4]."

In the next section we will briefly explain the optimal stopping rule, a fundamental concept in the resolution of our problem.

1.1 Optimal Stopping Rule

The optimal stopping rule is often applied in probability theory and decision theory. Originally, it was studied for problems arising in the sequential analysis of statistical observations.

The theory is used in decision-making scenarios, exactly like the one we have seen in the Secretary Problem, and it has the objective of maximising an expected payoff or minimising an expected cost by making choices about the timing of an action, given sequentially observed random variables. In simple terms, it applies to situations where an individual has to make a decision or select an option from a sequence of choices, but can only evaluate each option sequentially and cannot return to a previously passed option.

Definition 1.1. *Optimal stopping problems always have two elements:*

(i) a sequence of random variables, $X_1, X_2, ...,$ whose joint distribution is assumed known, and

(ii) a sequence of real-valued reward functions,

$$y_0, y_1(x_1), y_2(x_1, x_2), ..., y_\infty(x_1, x_2, ...).$$

Having established Definition 1.1, we continue with the layout of the problem to identify the optimal stopping rule. Consider a situation where an individual is asked to observe a sequence of random variables, $X_1, X_2, ...$ coming one at a time until they wish to stop. After observing $X_1 = x_1, X_2 = x_2, ..., X_n = x_n$, for $n = 1, 2, ...$ they can either stop and earn a reward equal to $y_n(x_1, ..., x_n)$ or continue the observation of X_{n+1}. If the individual does not take any observation, their payoff will be equal to y_0. Instead, if they never stop, they will receive $y_\infty(x_1, x_2, ...)$.

Remark 1. *We assume the rewards are uniformly bounded above by a random variable with finite expectation. In this way, all expectations will make sense.*

The problem is to determine a randomised stopping rule or probability of stopping using randomised decisions consisting of the the following sequence of functions,

$$\phi = (\phi_0, \phi_1(x_1), \phi_2(x_1, x_2), ...),$$

where for all n and $x_1, ..., x_n$, $0 \leq \phi_n(x_1, ..., x_n) \leq 1$. If $\phi_n(x_1, ..., x_n)$ is either 0 or 1, the stopping rule is said to be non-randomised. Therefore, ϕ_0 is the probability that an individual does not take any observation, and $\phi_1(x_1)$ is the probability that an individual observes $X_1 = x_1$ and accepts it. To summarise, the random time N at which an individual stops the observation is defined by the stopping rule, ϕ, and the sequence of observations, $X = (X_1, X_2, ...)$, with $0 \leq N \leq \infty$. The probability mass function of N given $X = x = (x_1, x_2, ...)$ is denoted by $\psi = (\psi_0, \psi_1, ..., \psi_\infty)$,

$$\psi_n(x_1, ..., x_n) = \mathbb{P}(N = n | X = x) \qquad \text{for } n = 0, 1, 2, ...,$$

$$\psi_\infty(x_1, x_2, ...) = \mathbb{P}(N = \infty | X = x).$$

Finally, a stopping rule ϕ is chosen to maximise the expected return, $V(\phi)$

$$\phi_n(X_1, ..., X_n) = \mathbb{P}(N = n | N \geq n, X = x) \text{ for } n = 0, 1, ...$$
$$V(\phi) = \mathbb{E}_y N(X_1, ..., X_N) = \mathbb{E} \sum_{j=0}^{\infty} \psi_j(X_1, ..., X_j) y_j(X_1, ..., X_j) \,.$$

1.2 The Secretary Problem Solution

Going back to the initial problem, we will now find that there is a simple solution to it. Recall the decision-making scenario: we have a secretary position available and an applicant who has to be chosen from a group of n people who are interviewed in a sequential and random order; the applicants are valued based on their relative ranks and can either be rejected or accepted upon observation, with only one chance of being accepted. Once they are rejected, they cannot be later recalled.

The first step towards an optimal stopping requires us to restrict our attention on an employer, or boss, who rejects the first $r - 1$ applicants, for some integer $r \geq 1$. The employer will then continue the observation and choose the next applicant who has the best position in the ranking of the applicants rejected $(r-1)$. With this strategy, the probability of selecting the best applicant is $1/n$ for $r = 1$. For $n > 1$, we write the probability as

$$\phi_n(r) = \mathbb{P}(\textit{this strategy makes us select the best applicant}).$$

which we can further extend into

$$\begin{aligned} \phi_n(r) &= \sum_{j=r}^{n} \mathbb{P}(\textit{jth applicant is best and we select it}) \\ &= \sum_{j=r}^{n} \frac{1}{n} \cdot \frac{r-1}{j-1} = \frac{r-1}{n} \sum_{j=r}^{n} \frac{1}{j-1} \end{aligned} \tag{1.1}$$

In equation (1.1), $1/n$ is the probability that the j-th applicant is the best among n applicants, and $\frac{r-1}{j-1}$ is the probability that the rejected applicants $r - 1$ have values higher than the applicants before j.

Suppose the number of the applicants is large ($n \to \infty$) and $x = \frac{r}{n}$. Then the above equation becomes

$$\begin{aligned}\phi_n(r) &= \frac{r-1}{n}\sum_{j=r}^{n}\frac{1}{j-1} = \frac{r-1}{n}\sum_{j=r}^{n}\frac{n}{j-1}\cdot\frac{1}{n} = \\ &= \frac{r-1}{n}\sum_{j=r}^{n}\frac{1}{\frac{j-1}{n}}\cdot\frac{1}{n} = \frac{r-1}{n}\sum_{i=\frac{r}{n}}^{1}\frac{1}{i-\frac{1}{n}}\cdot\frac{1}{n} = \\ &= (x-\frac{1}{n})\sum_{i=x}^{1}\frac{1}{i-\frac{1}{n}}\cdot\frac{1}{n}\end{aligned}$$

As $n \to \infty$,

$$\phi_n(r) = x\cdot\int_x^1 \frac{1}{t}dt = x\cdot[\ln|t|]_{t=x>0}^{t=1} = x[0-\ln x] = -x\ln x$$

Lastly, to determine the value of x that maximises this quantity we compute the derivative with respect to x and set it higher than zero. Then, $g(x) = -x\ln x$ is maximised when

$$g'(x) = -\ln x - x\cdot\frac{1}{x} = -\ln x - 1 > 0$$

$$\ln x < -1$$

$$e^{\ln x} < e^{-1}$$

$$x < \frac{1}{e}$$

Therefore, if we let n tend to infinity, the optimal stopping (or probability of success of choosing the best applicant) occurs at about 37% which is the value of x that we were looking for.

$$x = 1/e \approx 0,37$$

In conclusion, this suggests us to wait until 37% of the candidates have been interviewed and then pick the next relatively best one. By substituting the value of x into equation (1.2) we find that the probability that the strategy picks the best applicant is also 37%. This solution is applied in the classical case of the Secretary Problem which is very interesting and even more so due to it being so common. In fact, many situations involve making decisions without the possibility of going back and changing them. Therefore, as we have pictured the simplest form of the problem we will now go on with additional strategies and the last insights to the problem before concluding the chapter.

1.3 Simulations of the Secretary Problem

Let us define S_T as the strategy that picks the first applicant $k \geq T$ such that the value of the k-th applicant is bigger than maximal value seen between the first $T-1$ applicants.

The code below, written with MATLAB, computes the probability

$$\mathbb{P}(S_T \text{ chooses the best applicant}) \qquad \text{for } T = 1, 2, \ldots, n\,.$$

Such a probability is computed by the code by approximating the fraction of successes over the total number of samples of the same experiment. This can be found in the code as `K(T)=c/samples` and the experiment we are focusing on is the one in which the Secretary Problem picks the maximum value in the sequence by using strategy S_T. Note that this approximation is justified by the Law of Large Numbers A.4.

```
samples=100;
n=1000;

K=zeros(n,1);

for T=1:n
```

```
  c=0;
  for i=1:samples
    correct=secretary(n,T);
    c=c+correct;
  end
  K(T)=c/samples;
end

plot(1:n,K(1:n))
```

The function `secretary(n,T)` gives 1 if the strategy S_T picks the maximum over the n secretaries, otherwise, it gives 0. We write a code for this function and consider secretaries samples value taken randomly from a discrete uniform distribution on the set $\{1, 2, \ldots, n\}$. The output of the code is Figure 1.1.

```
function correct=secretary(n,T)

v=zeros(n,1);
for i=1:n
    v(i)=randi(n);
end

M=v(1);
for i=2:n
    if v(i)>M
        M=v(i);
    end
end
```

```
max=v(1);
for i=2:T-1
    if v(i)>max
        max=v(i);
    end
end

if T<n
    j=T+1;
    while v(j)<max && j<n
        j=j+1;
    end
    S=v(j);
end

if T==n || j==n
    S=v(n);
end

if S==M
   correct=1;
else
   correct=0;
end
```

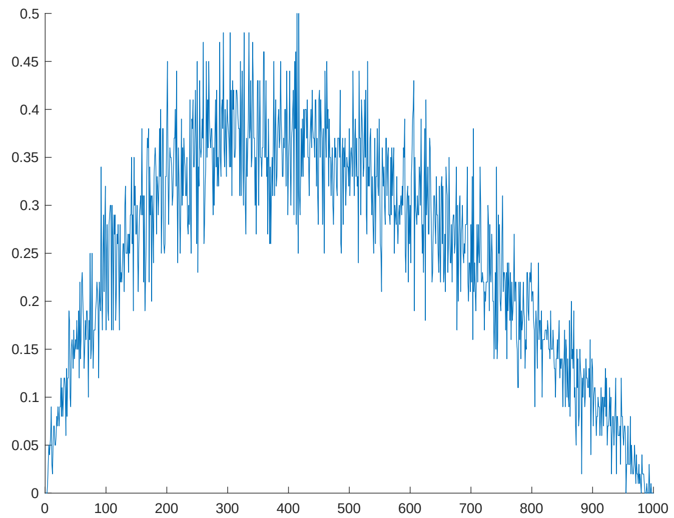

Fig. 1.1

1.4 The Secretary Problem for the Random Walk

In the next pages we will discuss in more detail the ideas and solutions presented in the paper authored by M. Hlynka and J.N. Sheahan [6]. The scholars solved the Secretary Problem in a situation characterised by dependent values and showed that the probability of success of their strategy tends to 0 as n tends to infinity, which differs from the result obtained in the classical case where the probability of success tends to $1/e$ when n tends to infinity.

Below are some notations to keep in mind:

(C1) $\{X_i\}$ are i.i.d. random variables for $i = 1, ..., n.$,

(C2) $Y_0 = 0$,

(C3) $Y_i = Y_{i-1} + X_i$,

(C4) $Y^* = \max\{Y_i : 0 \leq i \leq n\}$,

(C5) $Y^{\#}$ = the value picked by using strategy S.

Our objective is to find a strategy such that $\mathbb{P}(Y^{\#} = Y^*)$ is maximised. In other words, we want a strategy which can give us a high probability of picking the optimal strategy. And by optimal strategy we mean one which can maximise the probability of picking the largest value among n values with no recall allowed.

Assume that we are in a generalised one-dimensional random walk.

Definition 1.2. *Let S_0 be the strategy which picks $Y_0 = 0$ for Y^*. And let S_k (for $0 < k \leq n$) be the strategy that picks Y_i after having looked at the values $Y_0, ..., Y_{k-1}$ such that $Y_i > max\{Y_0, ..., Y_{k-1}\}$. Know that $i \geq k$ and that if no such Y_i exists then the value picked will be Y_n.*

Assumption A1. *For $i = 1, ..., n$, let $\{X_i\}$ be i.i.d. continuous random variables with a distribution function F with density f symmetric about 0 and support equal to an interval.*

Assumption A2. *For $i = 1, ..., n$, and $m = 0, 1, ..., \infty$ let $\{X_i\}$ be i.i.d. discrete random variables such that $\mathbb{P}(X_i = m) = \mathbb{P}(X_i = -m) = p(m)$ (the distribution is symmetric about 0).*

We derive the following properties:

Proposition 1.4.1. *If assumption A1 holds then $\mathbb{P}(Y^{\#} = Y^*)$ is the same for all S_k and S_k is an optimal strategy.*

Proposition 1.4.2. *If assumption A2 holds then $\mathbb{P}(Y^{\#} = Y^*)$ is the same for all S_k.*

Proposition 1.4.3. *If assumption A2 hold then S_0 is an optimal strategy.*

We will focus on the discrete case and prove the properties following Gilbert and Mosteller (1966).

Proof. First, we define a candidate for $Y^{\#}$ from the sequence $\{Y_i\}$ which has a value greater than or equal to all the values observed until now. We can distinguish two different events: (i), $Y_i = Y^*$, 'win with Y_i' and (ii), $Y^{\#} = Y^*$, 'win'.
Let

$$g(i) = \mathbb{P}(\text{win with } Y_i | Y_i \geq Y_0, Y_1, ..., Y_{i-1})$$

and

$$h(i) = \mathbb{P}(\text{win with best strategy from } i+1 \text{ on} | Y_i \geq Y_0, Y_1, ..., Y_{i-1})$$

Then, given that $g(0) = \mathbb{P}(\text{win with } Y_0)$ and $h(n) = 1$, if Y_i is a candidate:

- An optimal strategy chooses Y_i as Y^* if $g(i) > h(i)$, in other words when $\mathbb{P}(\text{win with } Y_i) > \mathbb{P}(\text{win with } Y_{i+1} \text{ on})$,

- An optimal strategy can either choose Y_i as Y^* or wait for the next candidate i' with $g(i') > h(i')$ if $g(i) = h(i)$.

To show that S_0 is an optimal strategy for $\{X_i\}$ i.i.d. discrete random variables, we only need to prove that $g(0) \geq h(0)$. This means that the probability of winning by picking the candidate Y_0 is greater than or equal to the probability of winning by picking any other candidate from Y_1 on, given that the picked candidate has value higher than the values observed thus far.

Additionally, showing that S_0 is an optimal strategy is the same as demonstrating that $Y_0 = 0$ is truly the maximum we can achieve. In fact, by Property 1.4.1 we know that the probability that S_0 is an optimal strategy is the same for all $k = 0, ..., n$. And to make the proof easier, we can use $k = 0$ which is equivalent to using the strategy S_0. Hence using this strategy is the same as saying that we are choosing Y_0.

Set $u(r) = \mathbb{P}(Y_0$ is a maximum for the random walk sequence $Y_0, ..., Y_r)$, for

$0 \leq k \leq n$:

$$
\begin{aligned}
g(k) &= \mathbb{P}(\text{win with } Y_k | Y_k \geq Y_0, ..., Y_{k-1}) \\
&= \mathbb{P}(Y_k \text{ is a maximum for the random walk sequence } Y_k, ..., Y_n) \\
&= \mathbb{P}(W_0 \text{ is the maximum for the random walk sequence } W_0, ..., W_{n-k}) \\
&\text{(where } W_i = Y_{k+i} - Y_k) \\
&= u(n-k),
\end{aligned}
\tag{1.2}
$$

$$
\begin{aligned}
h(k) &= \mathbb{P}(\text{win with the best strategy from } k+1 \text{ on} | Y_k \geq Y_0, ..., Y_{k-1}) \\
&= \mathbb{P}(\text{win with the best strategy from step 1 on in a random walk sequence} \\
&\quad W_0, ..., W_{n-k}) \\
&\geq \mathbb{P}(\text{win by picking the last element of a random walk sequence} \\
&\quad W_0, ..., W_{n-k}) \\
&= \mathbb{P}(\text{win by picking the first element } W_0 \text{ of a random walk sequence} \\
&\quad W_0, ..., W_{n-k}) \\
&= u(n-k).
\end{aligned}
\tag{1.3}
$$

From equation (1.3) we obtain $h(k) \geq u(n-k)$ and with equation (1.2)we arrive to the conclusion that $h(k) \geq g(k)$ for $0 \leq k \leq n$. So, an optimal strategy would be to wait for the last value and choose it. However, what we wanted to prove was the opposite situation. To clarify, we wanted to show that S_0 is an optimal strategy by demonstrating that $g(0) \geq h(0)$. This can be proven by replacing '$\geq$' with '$=$' in the equation we computed above (1.3). Finally, we have that $g(k) = h(k)$ for all k (where $0 \leq k \leq n$), and consequently, the optimal strategy is to pick $Y_0 = 0$ for Y^*. This strategy is closely related to Feller's observation that the maximum value of a random walk is more likely to be at the start or at the end of the sequence than in the middle of it. □

Proposition 1.4.4. *Let n be a positive even integer and $\{X_i\}_{i=0}^n$ be a sequence of i.i.d. random variables with $\mathbb{P}(X_i = 1) = \mathbb{P}(X_i = -1) = 0.5$ and $Y_i = Y_0 + \sum_{j=0}^{i-1} X_j$ for $i = 0, ..., n$ with $Y_0 = 0$ Then, if assumption A2 holds, for all k=0,1,...,n, using strategy S_k,*

$$\mathbb{P}(Y^{\#} = Y^*) = \binom{n}{n/2} \frac{1}{2^n} \underset{n\to\infty}{\sim} \sqrt{\frac{2}{\pi n}}$$

1.4.1 Proof of Proposition 1.4.4

Recall that we are choosing an element only if it is the maximum of those seen up to that point (unless it is the last element). Together with proposition 1.4.4, this tells us that even though all of our strategies $S_0, S_1, ..., S_n$ are equivalent, the probability of choosing any element may not be the same. What interests us is $\mathbb{P}(Y_k^{\#} = Y_k^* | S_k)$. By proposition 1.4.2 we know that $\mathbb{P}(Y^{\#} = Y^*)$ is the same for all $k = 0, 1, ..., n$. Indeed, if we let $Y^{\#}$ be the picked value under strategy S_k then the probability is independent on the index k. Hence, since $Y^{\#} = Y_0 = 0$ when $k = 0$, we have

$$\mathbb{P}(Y^{\#} = Y^*) = \mathbb{P}(Y^* = 0)$$

Since Y^* is the maximum value of the random walk, if $Y^* = 0$ then

$$\mathbb{P}(Y^* = 0) = \mathbb{P}(Y_i \leq 0 \qquad \text{for all } i = 1, ..., n)$$

So to prove the proposition we have to show that the following equation is true

$$\mathbb{P}(Y_i \leq 0 \text{ for } i = 1, ..., n) = \binom{n}{n/2} \frac{1}{2^n} \underset{n\to\infty}{\sim} \sqrt{\frac{2}{\pi n}}$$

Consider the outcomes of X_i as prizes associated to a sequence of tossings of a fair coin. At each time i the coin is tossed: we win 1 unit in case of head and loose 1 unit in case of tail. Y_i is the total gain in i tossings.

We study the event of having a reward equal to r at time n, which is the event of $\{Y_n = r\}$. We want to compute the following probability $p_{n,r} := \mathbb{P}(Y_n = r)$, where p denotes the number of heads and q, instead, denotes the number of tails obtained in the first n tossings. We can write the system

$$Y_n = r \leftrightarrow \begin{cases} p - q = r, \\ \\ p + q = r. \end{cases}$$

We get $p = \frac{n+r}{2}$ and $q = \frac{n-r}{2}$. Then we use the binomial coefficient to count the number of ways in which we can choose the tossings that give head (p) out of the total number of tossings (n). This number is equivalent to the number of paths starting with Y_0 and arriving with $Y_n = r$ since each path $Y_0, ..., Y_n$ is obtained through the sequence $X_1, ..., X_n$.

$$N_{n,r} := \binom{n}{\frac{n+r}{2}}$$

The probability of the event of a path occurring has probability $\frac{1}{2^n}$ since each X_i can be either 1 or -1 and the total number of possible paths is equal to 2^n. The probability of winning r units at time n is

$$p_{n,r} = N_{n,r} \cdot \frac{1}{2^n} = \binom{n}{\frac{n+r}{2}} \cdot \frac{1}{2^n}$$

Because what we wanted to compute was the probability that all Y_i are less than or equal to zero, we can rewrite the above formula

$$p_{n,0} = \binom{n}{\frac{n}{2}} \cdot \frac{1}{2^n}$$

Therefore

$$\mathbb{P}(Y_i \leq 0 \text{ for } i = 1, ..., n) = p_{n,o} = \binom{n}{\frac{n}{2}} \cdot \frac{1}{2^n}$$

This probability coincides to the probability of having a path with $Y_0 = 0$ and $Y_i \leq 0$ for all $i = 1, ..., n$. In conclusion, we prove our thesis by applying Stirling's formula

$$p_{n,0} \underset{n \to \infty}{\sim} \sqrt{\frac{2}{\pi n}}$$

To sum up, the optimal strategy for our problem, which maximises the probability of getting the maximum of a random walk, involves picking either the first or the last element of the sequence. A surprising fact to highlight is that this probability is considerably better than $1/n$ which is what is indicated in Property 1.4.4.

As we have seen, the problem can have many solutions as well as a number of versions and extensions. From a more practical perspective we can consider an alternative way to describe the problem. Instead of a secretary, we might think of a stock analyst as the protagonist of our scenario. He is in charge of picking the day in which a particular stock will have the highest price during a given month. And on the day when the analyst believes the stock to be the highest, he has to call his client to let him know of his view. At the end, after having observed the stock prices during the rest of the month, the analyst will receive a reward if he was correct or nothing otherwise.

Regarding this subject, Fama (1965) had the opinion that the analysis and, more specifically, the methods and solutions we have examined so far can be used to help the stock analyst to make a decision. Even more so if the stock prices behave like a random walk during the month, exactly like in the Secretary Problem.

1.5 The Random Walk Theory and the Stock Exchange Market

The stock exchange market is a market for shares in corporations. In particular, it is a secondary market where securities previously issued are bought and sold, serving as a trading platform for investors. Compared to bonds, stocks are riskier since they depend on the performance of the company. In fact, stocks actually represent ownership in a firm and can earn a return in two primary ways: through the increasing price of the stock over time and through dividends paid to stockholders. In this chapter we will focus on the first. In addition to this, there are two different types of stocks: the common stock and the preferred stock. The former type gives the stockholder the right to vote, receive dividends and residual claim, however in case of bankruptcy all bills (wages, supplier's bills, interest of bondholders) will be paid before the common stockholder. On the other hand, the preferred stock grants limited rights to the stockholder but provides a preferential treatment and prior claim: cash dividends are not paid to the common stockholders if the preferred stockholders have not yet been paid. Moreover, stocks are exchanged and traded over various markets: organized exchanges (e.g. NYSE); over the counter markets (e.g. NASDAQ); and electronic communication networks or ECNs (e.g. instinet, selectnet, NYSE arca). Lastly, stocks can also be bought and sold through exchange-traded funds (ETFs). ETFs are essentially baskets of securities such as stocks, bonds or commodities and track the performance and return of market indexes. They typically have low expense ratios and fewer broker commissions than buying the stocks individually.

Moving on, we will see a connection between the stock exchange market and the Secretary Problem for the random walk.

1.5.1 Techniques for Predicting Stock Prices

As we have concluded in the first sections of this chapter, the Secretary Problem reflects issues similar to those concerning the work of market analysts in the financial sector. Is there an efficient method for predicting stock market prices? How is an analyst going to maximise his return? How can he optimally select the highest price from the sequence of stock price changes? With these questions in mind, Fama examined models of stock price behaviour and allowed us to see the problem from a new and more practical perspective. He was interested in describing one model in particular, the theory of random walks, which has gained a significant importance in the field of finance and statistics.

In order to understand the relevance of random walks in stock market prices, there are two common techniques often adopted by market professionals that need to be discussed first: the "chartist" or "technical" theories, and the theory of fundamental or intrinsic value analysis. The former, in opposition to the random walk theory (which we will talk about next), are based on the assumption that history tends to repeat itself. In essence, this means that past patterns in individual securities are used to predict the future behaviour of the series. The idea that the series of successive price changes in individual securities are dependent is a key element and characteristic of such technical models. Thus the chartist theories assume that "*the sequence of price changes prior to any given day is important in predicting the price change for that day* [3]". An example of a chartist approach is the Dow Theory.

Next, fundamental analysis assumes that an individual security has an intrinsic value, or equilibrium price, at any given point in time. This value is based on the security's earning potential, which, in turn, can be influenced by the industry and economic conditions, and other financial and company management qualitative factors as well. This theory suggests to predict the security future price by carefully studying fundamental factors and determining whether the actual price is moving up or down in comparison to its intrinsic value. In a nutshell, this theory holds that determining the intrinsic value is the same as determining the future price of the security.

To proceed, whilst the technical analysis focuses on past market data to forecast future price movements, the random walk is a theory based on the concept that prices are independent and reflect all available (and important) information. This means that prices follow a random path, differently from the technical approach we mentioned first. Moreover, the independence of successive price changes in individual securities is a property of an efficient market also called by Fama as the 'instantaneous adjustment' property. This property has two important implications: first, the frequency with which actual prices overadjust to changes in intrinsic values will equal the frequency with which actual prices underadjust; second, the event of changes in intrinsic values due to new information will be anticipated by the market and preceded by the adjustment of actual prices. Furthermore, an efficient market is defined as one where the intrinsic value provides a reliable estimate of a security's actual price at any given time. This definition implies that current prices already reflect information from both past and anticipated future events. However, in an uncertain world, intrinsic values can never be determined exactly. Nevertheless, the collective actions of numerous market participants drive actual prices to fluctuate randomly around their intrinsic values.

All in all, by definition, a random walk market is: "*a market where successive price changes in individual securities are independent* [3]." Rephrased, it is a market where the sequence of price changes has no memory or record of past history. Hence, the random walk can be considered an adequate description of reality or otherwise, depending on the actual degree of dependence in the series of price changes. Precisely, in order to be adequate the degree of dependence must not be sufficient to make the expected profits of any more complicated and mechanical procedures or other techniques greater than the expected profits under the simple policy of buying and holding the security.

To prove the acceptability of the random walk model, empirical research has been conducted to test the hypothesis that successive price changes are indeed independent. With this intent, two approaches have been employed generating different conclusions:

(i) The first approach relies primarily on common statistical tools like serial correlation coefficients and analysis of runs of consecutive price changes of the same sign.

(ii) The second approach tests directly different mechanical trading rules to see if they provide profits greater than buy-and-hold.

Over the years, many studies have focused on the first approach to testing independence and showed that the sample serial correlation coefficients computed for successive price changes were extremely close to zero. This result is an empirical evidence against the idea of dependence in the series of changes.

Despite the evidence presented, technical theorists continued to view the random walk model as inadequate. It was in this context that Alexander's filter technique was introduced, aiming to apply more sophisticated criteria to analyze the independence of successive price changes in financial markets.

The filter technique systematically generates buy and sell signals based on specific percentage movements in the market, and incorporates the short-selling investment strategy. The strategy, also called short position, requires an investor to borrow shares of a stock (or another asset) and to sell them on the open market, planning to buy them back in the future at a lower price. Hence, the profit is generated from a decline in the asset's price during the holding period.

In our example, we will use a 5% filter: when the price of a security increases by at least 5%, buy and hold the security until its price moves down at least 5% from a subsequent high, at which simultaneously sell and go short. The short position is maintained until the daily closing price rises at least 5% above a subsequent low. This strategy can be divided into four steps:

- The initial buy signal allows the investor to profit from the continued rise in price,
- The holding period helps the investor to benefit from any further increase in price,
- The sell and short signal expects the downward trend to continue and allows the investor to profit from further decline in price,
- Lastly, the short position period is held until the price rises by 5% which signals to cover and buy the security, this signal assumes a continued raise in price and another opportunity to profit from the upward movement.

In brief, the strategy was tested on individual securities, demonstrating that the simple buy-and-hold method consistently yielded greater profits. This finding confirmed the efficiency of the random walk model and supported the theory of independence of price changes.

Chapter 2

The Prophet Inequality

In this chapter we are going to focus on the classic Prophet Inequality, a mathematical tool used in the field of economics and decision theory, which involves important mathematical theories such as the optimal stopping rules, pure online algorithms and also mechanism design. The importance of the Prophet Inequality lies in its help to address fundamental questions about decision-making under uncertainty.
The first to prove the famous Prophet Inequality were Krengel, Sucheston and Garling in 1978 [7]. In particular, their work explained the relative power of online and offline algorithms in Bayesian settings. We will analyze the classic framework in which it is set and try to reach to the final result through a mathematical proof.

The problem is the following: given a sequence of random numbers and a reward equal to the value chosen, can we measure the performance of a gambler compared to that of a “prophet” if the gambler is the person who is going to make a choice without knowledge of the values in the sequence whilst the prophet knows all the values in advance and can make the most rewarding decision? The Prophet Inequality affirms that a gambler who knows the distribution of each random variable can achieve at least half as much reward, in expectation, as a prophet who knows the sampled values of each random variable and can choose the biggest one. In other words, we are considering a game in which two people are playing: the gambler and the prophet. The gambler is a player who is observing a sequence of independent, non-negative random variables with finite expectations. From this sequence, a value is drawn in an indexed order, and after every draw the gambler can either terminate the game by accepting the value observed or continue the game and discard the value observed. His aim is to achieve the highest reward possible and maximise the expected value related to it. On the other hand, the prophet is someone who knows in advance the highest value in the sequence and can easily achieve the best result.

Now that we have a clear picture of the problem we are back to our first question, how is the gambler going to perform compared to the prophet? We will try to answer to this question in the rest of the chapter.

To begin with, the term of 'relative power' refers to a comparative effectiveness of using different methods and approaches in achieving a particular objective or, in our case, solving a decision-making problem. The different methods used in our problem are the online and offline algorithms in Bayesian settings. Offline algorithms typically have access to the whole dataset at once while online algorithms handle data as it arrives sequentially using the new data to change their estimation. In our case, the Bayesian setting is the one in which the online algorithm is linked to a gambler who knows the distribution from which the sequence will be sampled, whereas the offline algorithm is instead linked to a prophet who can foretell the entire sequence and stop at its maximum value.

What the three mathematicians showed is that, if we consider the game we have described in the previous pages, the prophet can gain at most twice as much payoff as the gambler, a player who must choose the stopping time based only on the current and past observations.

Suppose:

(A1) $X_1, X_2, ..., X_n$ is a sequence of independent, non-negative, real-valued random variables;

(A2) $\mathbb{E}[\max_i X_i] < \infty$ the maximum value in the sequence has finite expectation;

(A3) a threshold is defined $T = \mathbb{E}[\max_i X_i]/2$;

(A4) τ is a stopping time such that $X_\tau \geq T$.

Theorem 2.0.1 (Prophet inequality for independent random variables.)**.** *If the assumptions defined above holds, then*

$$\mathbb{E}[max_i X_i] \leq 2\mathbb{E}[X_\tau]\,.$$

The inequality compares the performance of online and offline algorithms for problems that involve selecting one or more elements from a random sequence. We are going to prove that there is an online algorithm whose expected payoff is at least half of the expected weight of the maximum weight and that the factor 2 is the optimal constant and cannot be improved.

Proof. First, we want to calculate the expectation of X_τ. By Lemma A.1

$$\mathbb{E}[X_\tau] = \int_0^{+\infty} \mathbb{P}(X_\tau > x)dx = \int_0^T \mathbb{P}(X_\tau > x)dx + \int_T^{+\infty} \mathbb{P}(X_\tau > x)dx\,.$$

We have two cases:

(i) $X_\tau > x$ for $x \in (T, +\infty)$;

(ii) $X_\tau > x$ for $x \in (0, T]$.

By assumptions 2, if $X_i < T$ for all $i = 1, ..., n$, we assume $\tau = \infty$ and define $X_\infty := X_n$. Moreover, $p = \mathbb{P}(\max_i X_i \geq T)$.

Let us start from point (i). Note that for any $x > T$ we have

$$\mathbb{P}(X_\tau > x) = \sum_{i=1}^{n} \mathbb{P}(X_\tau > x|\tau = i)\mathbb{P}(\tau = i) + \mathbb{P}(X_\tau > x|\tau = \infty)\mathbb{P}(\tau = \infty)\,. \quad (2.1)$$

We know that

$$\mathbb{P}(X_\tau > x|\tau = \infty)\mathbb{P}(\tau = \infty) = \mathbb{P}(X_\infty > x|\tau = \infty)\mathbb{P}(\tau = \infty) = \\ = \mathbb{P}(X_n > x|\tau = \infty)\mathbb{P}(\tau = \infty) = 0 \cdot \mathbb{P}(\tau = \infty) = 0\,.$$

So we have that the expression (2.1) becomes

$$\mathbb{P}(X_\tau > x) = \sum_{i=1}^{n} \mathbb{P}(X_i > x|\tau = i)\mathbb{P}(\tau = i) = \\ = \sum_{i=i-1}^{n} \mathbb{P}(X_i > x|X_1 < T, X_2 < T, ..., X_{i-1} < T, X_i > T) \cdot \mathbb{P}(\tau = i)\,. \quad (2.2)$$

Since the random variables $\{X_i\}_{i=1}^n$ are independent

$$\mathbb{P}(X_i > x|X_1 < T, X_2 < T, ..., X_{i-1} < T, X_i > T) = \mathbb{P}(X_i > x|X_i > T) = \\ = \frac{\mathbb{P}(X_i > x, X_i > T)}{\mathbb{P}(X_i > T)} = \frac{\mathbb{P}(X_i > x)}{\mathbb{P}(X_i > T)}\,. \quad (2.3)$$

Hence, by independence of the random variables and by (2.3), we can rewrite (2.2) as

$$\mathbb{P}(X_\tau > x) = \sum_{i=1}^{n} \frac{\mathbb{P}(X_i > x)}{\mathbb{P}(X_i > T)} \cdot \mathbb{P}(\tau = i) = \\ = \sum_{i=1}^{n} \frac{\mathbb{P}(X_i > x)}{\mathbb{P}(X_i > T)} \cdot \mathbb{P}(X_i > T) \cdot \mathbb{P}(X_j < T \text{ for } j = 1, ..., i-1) = \quad (2.4) \\ = \sum_{i=1}^{n} \mathbb{P}(X_i > x) \cdot \mathbb{P}(X_j < T \text{ for } j = 1, ..., i-1)\,.$$

Remark 1. *It is possible to show that the probability of the event* $\{max_i X_i < T\}$ *is less than or equal to the probability of the event* $\{X_j < T \text{ for } j = 1, ..., i-1\}$.

In simple words, the event of finding the highest value in the sequence $X_1, X_2, ..., X_n$ *before the threshold value implies that all the other values in the sequence* $X_1, X_2, ..., X_{i-1}$ *are also found before the threshold value since there is no value higher than* T*.*
As a result,

$$i - 1 < n\,,$$

The probability of having the maximum value, one variable in the sequence, before T *is lower compared to the probability of having all the other values except for* X_n *before* T*.*
Indeed,

$$\{max_i X_i < T\} =$$
$$= \{X_1 < T, X_2 < T, ..., X_n < T\} \subseteq \{X_1 < T, X_2 < T, ..., X_{i-1} < T\} =$$
$$= \{X_j < T \ for\ j = 1, ..., i-1\}\,,$$

So the following statement holds true

$$\mathbb{P}(max_i X_i < T) \leq \mathbb{P}(X_j < T \ for\ j = 1, ..., i-1)\,,$$

And by definition of p

$$\mathbb{P}(X_j < T \ for\ j = 1, ..., i-1) \geq 1 - p\,.$$

Then, (2.4) becomes

$$\mathbb{P}(X_\tau > x) \geq \sum_{i=1}^{n} \mathbb{P}(X_i > x) \cdot (1-p) = (1-p) \sum_{i=1}^{n} \mathbb{P}(X_i > x)\,. \tag{2.5}$$

Next, we can easily see that by the union bound,

$$\mathbb{P}(\mathrm{max}_i X_i > x) = \mathbb{P}(\exists\, i \text{ such that } X_i > x) = \mathbb{P}(\cup_{i=1}^{n} \{X_i > x\} \leq \sum_{i=1}^{n} \mathbb{P}(X_i > x)\,.$$

Hence, (2.5) becomes

$$\mathbb{P}(X_\tau > x) \geq (1-p)\mathbb{P}(\mathrm{max}_i > x)\,. \tag{2.6}$$

Let us now move to point (ii). Note that if $x \le T$, we have

$$\begin{aligned}\mathbb{P}(X_\tau > x) &= \sum_{i=1}^{n} \mathbb{P}(X_\tau > x | \tau = i)\mathbb{P}(\tau = i) + \mathbb{P}(X_\tau > x | \tau = \infty)\mathbb{P}(\tau = \infty) = \\ &= \sum_{i=1}^{n} 1 \cdot \mathbb{P}(\tau = i) + \mathbb{P}(X_\tau > x | \tau = \infty)\mathbb{P}(\tau = \infty) = \\ &\ge \sum_{i=1}^{n} \mathbb{P}(\tau = i) = \mathbb{P}(\tau < \infty) = \mathbb{P}(\max_i X_i > T) = p\,.\end{aligned} \tag{2.7}$$

Finally, by definition of T and by Lemma A.1

$$\begin{aligned}2T = \mathbb{E}[\max_i X_i] &= \int_0^T \mathbb{P}(\max_i X_i > x)dx \, + \int_T^{+\infty} \mathbb{P}(\max_i X_i > x)dx = \\ &\le \int_0^T 1dx \, + \int_T^{+\infty} \mathbb{P}(\max_i X_i > x)dx = T \, + \int_T^{+\infty} \mathbb{P}(\max_i X_i > x)dx\,,\end{aligned}$$

so

$$2T \le T \, + \int_T^{+\infty} \mathbb{P}(\max_i X_i > x)dx\,,$$

and

$$T \le \int_T^{+\infty} \mathbb{P}(\max_i X_i > x)dx\,.$$

The above result, together with (2.6) and (2.7) help us prove Theorem 2.0.1. We can show that

$$\begin{aligned}\mathbb{E}[X_\tau] = \int_0^{+\infty} \mathbb{P}(X_\tau > x)dx &= \int_0^T \mathbb{P}(X_\tau > x)dx \, + \int_T^{+\infty} \mathbb{P}(X_\tau > x)dx = \\ &\ge \int_0^T p\,dx \, + \int_T^{+\infty} (1-p)\mathbb{P}(\max_i X_i > x)dx \ge pT \, + (1-p)T = \\ &= T = \frac{1}{2}\,\mathbb{E}[\max_i X_i]\,.\end{aligned}$$

□

The last thing we want to show is that the constant 2 appearing in Theorem 2.0.1 is the optimal constant, it is not possible to find a value $C \in (0, 2)$ for which $\mathbb{E}[\max_i X_i] \le C\,\mathbb{E}[X_\tau]$ for any stopping rule τ.

Assume $n = 2$, in other words, we have two random variables in the sequence X_1

and X_2. Given $\varepsilon \in (0,1)$, we define $X_1 := 1$ and

$$X_2 = \begin{cases} \frac{1}{\varepsilon}, & \text{with probability } \varepsilon, \\ 0, & \text{with probability } 1-\varepsilon. \end{cases}$$

We can observe that $\frac{1}{\varepsilon} > 1$ and $\mathbb{E}[X_2] = \mathbb{E}[X_1] = 1$. Therefore, under any stopping rule τ, we have $\mathbb{E}[X_\tau] = 1$. So,

$$\max\{X_1, X_2\} = \begin{cases} \frac{1}{\varepsilon}, & \text{with probability } \varepsilon, \\ 1, & \text{with probability } 1-\varepsilon. \end{cases}$$

From this we can compute

$$\mathbb{E}[\max\{X_1, X_2\}] = 2 - \varepsilon\,.$$

When $\varepsilon \to 0$ we have the following result

$$\mathbb{E}[\max\{X_1, X_2\}] = 2 = 2 \cdot 1 = 2 \cdot \mathbb{E}[X_\tau]\,.$$

2.1 The Prophet Inequality with Cost for Observations

The Prophet Inequality we have analysed in the previous section aimed at comparing the performance of online and offline algorithms in the selection of a value drawn from a random sequence. Now, we will consider a different problem scenario where the comparison is made between the performance of a prophet with complete foresight and the performance of a gambler who is observing a sequence of i.i.d., real-valued random variables, exactly as in the classical Prophet Inequality that we have just observed at the beginning of this chapter. What's new is a non-negative fixed cost charged for each observation. This case was described by E.

Samuel-Cahn in 1992 [1], who introduced the main theorem for this new prophet problem.

Before we begin with the mathematical analysis, there are some hypothesis which need to be taken into account.

Assume:

(B1) X_i is independent, $i = 1, 2, \dots, 0 \leq X_i \leq 1$;

(B2) $c \geq 0$ is a fixed constant;

(B3) the optimal stopping problem for the sequence $Y_i = X_i - ic$, $i = 1, 2, ...,$ corresponds to a reward X_i minus a fixed cost c of sampling, for each observation;

(B4) $V(Y_1, ..., Y_n) = sup\,\{\mathbb{E}[Y_\tau] : \tau \leq n,\ \tau$ is a stopping rule$\}$;

(B5) $[x]$ denotes the largest integer smaller than x.

Theorem 2.1.1 (Difference Prophet Inequality with cost.)**.** *Let* $\{X_i\}_{i=1}^n$ *be a sequence of i.i.d. with* $0 \leq X_i \leq 1$.

(a) For $0 \leq c \leq 1$ *fixed and a positive integer* n*,*

$$E[\max_{1\leq i\leq n} Y_i] - V(Y_1, ..., Y_n) \leq \left[\frac{1}{c}\right]\ c\,(1-c)^{[1/c]+1}, \tag{2.8}$$

(b) For $n \geq 1$ *fixed and all* $c \geq 0$*,*

$$E[\max_{1\leq i\leq n} Y_i] - V(Y_1, ..., Y_n) \leq (1 - 1/n)^{n+1}, \tag{2.9}$$

(c) For all $c \geq 0$ *and all finite or infinite sequences* $(n = +\infty)$*,*

$$E[\max_{1\leq i\leq n} Y_i] - V(Y_1, ..., Y_n) \leq e^{-1}. \tag{2.10}$$

The theorem presents three different situations and defines the best possible bound for each one of them based on the fixed cost and length of the sequence. We will focus on the first inequality of the theorem (2.8).

Remark 2. *We will consider only $c \leq 1$, otherwise both the maximum and the optimal stopping value will be obtained for $n = 1$ and the difference will be equal to 0. The purpose of this note is to help with the explanation of Theorem 2.1.1.*

2.2 Proof of the Difference Prophet Inequality

We assume $n \geq 2$, otherwise the theorem is trivial. For the proof of the theorem Chow and Robbins in 1961 have considered the optimal stopping problem (for infinite horizon) for the payoff sequence $Y_i^* = \max_{1 \leq j \leq i} X_j - ic$, $i = 1, 2, ...,$ where $X_1, X_2, ...,$ are i.i.d. with finite expectation and $c > 0$. They have shown that there exists an optimal rule s given by:

$$s = \inf\{i : X_i \geq \beta\},$$

where β is the unique value for which $\mathbb{E}[\max\{X_1 - \beta, 0\}] = c$. This result suggests that since $Y_i \leq Y_i^*$, the rule is also optimal for the payoff $Y_i = X_i - ic$ and as a consequence $\mathbb{E}[Y_s] = \beta$. Moreover, we define $s_n = \min\{s, n\}$ for Y_i and all $t \leq n$. Then, we can set our goal to be the maximisation of the following difference:

$$\mathbb{E}[\max_{1 \leq i \leq n} Y_i] - \mathbb{E}[Y_{sn}].$$

As the theorem suggests, this difference is always less than or equal to the right hand-side of (2.8) and (2.9), and its maximum is achieved at the equality for special Bernoulli random variables taking values of 0 and 1 with positive probability. Moreover, let $V(Y_1, ..., Y_n) = \mathbb{E}[Y_{sn}]$. Lastly, for the purpose of the thesis, there are two additional identities take note of:

- The expectation of X_1 given that the value of X_1 is bigger or equal than beta,

$$\mathbb{E}[X_1|X_1 \geq \beta] - \beta = \frac{c}{u}, \tag{2.11}$$

- And the probability of X_i assuming the value of 1,

$$p = c/(1-\beta). \tag{2.12}$$

Lemma 2.2.1.

$$\mathbb{E}[Y_{sn}] = \beta - (1-u)^n(\beta - \mathbb{E}[X_1|X_1 < \beta]), \tag{2.13}$$

Let $u = \mathbb{P}(X \geq \beta)$ *and the value in the equation equal to* β *when* $u = 1$.

Proof. We start from the left side of the equation in (2.13)

$$\mathbb{E}[Y_{sn}] = \sum_{i=1}^{n} \mathbb{E}[Y_{sn}|sn = i]\,\mathbb{P}(s_n = i). \tag{2.14}$$

Then we decompose the equation and study $\mathbb{E}[Y_{sn}|sn = i]$ and $\mathbb{P}(s_n = i)$ separately. For $i = 1, ..., n-1$:

$$\begin{aligned}\mathbb{E}[Y_{sn}|s_n = i] &= \mathbb{E}[Y_i|X_1 < \beta, ..., X_{i-1} < \beta, X_i \geq \beta] = \\ &= \mathbb{E}[X_i - ic|X_1 < \beta, ..., X_{i-1} < \beta, X_i \geq \beta] = \\ &= \mathbb{E}[X_i - ic|X_i \geq \beta] = \mathbb{E}[X_i|X_i \geq \beta] - ic = \mathbb{E}[X_1|X_i \geq \beta] - ic.\end{aligned} \tag{2.15}$$

Recall the Law of Total Probability (Theorem 2.2). For $i = n$:

$$\begin{aligned}\mathbb{E}[Y_{sn}|s_n = n] &= \mathbb{E}[X_n - nc|s_n = n] = \\ &= \mathbb{E}[X_n - nc|\{X_1 < \beta, ..., X_{n-1} < \beta, X_n \geq \beta\} \cup \\ &\cup \{X < \beta, ..., X_{n-1} < \beta, X_n < \beta\}] = \\ &= \mathbb{E}[X_n - nc|X_1 < \beta, ..., X_{n-1} < \beta] = \mathbb{E}[X_n - nc] = \mathbb{E}[X_n] - nc = \\ &= \mathbb{E}[X_1] - nc.\end{aligned} \tag{2.16}$$

For $i = 1, ..., n-1$:

$$\begin{aligned}\mathbb{P}(S_n = i) &= \mathbb{P}(X_1 < \beta, X_2 < \beta, ..., X_{i-1} < \beta, X_i \geq \beta) = \\ &= \mathbb{P}(X_1 < \beta) \cdot \mathbb{P}(X_2 < \beta) ... \mathbb{P}(X_{i-1} < \beta) \cdot \mathbb{P}(X_i \geq \beta) = \\ &= (1-u)^{i-1} \cdot u.\end{aligned} \tag{2.17}$$

Again, keep in mind Theorem 2.2. For $i = n$:

$$\begin{aligned}\mathbb{P}(S_n = n) &= \mathbb{P}(X_1 < \beta, ..., X_n < \beta) + \mathbb{P}(X_1 < \beta, ..., X_{n-1} < \beta, X_n \geq \beta) = \\ &= \mathbb{P}(X_1 < \beta, ..., X_{n-1} < \beta) = \mathbb{P}(X_1 < \beta)\,\mathbb{P}(X_2 < \beta) ... \mathbb{P}(X_{n-1} < \beta) = \\ &= (1-u)^{n-1}.\end{aligned} \tag{2.18}$$

To simplify, set $\alpha = (\mathbb{E}[X_1] - nc)(1-u)^{n-1}$. Then, given the results in (2.15), (2.16), (2.17), and (2.18), (2.14) becomes

$$\begin{aligned}\mathbb{E}[Y_{sn}] &= \sum_{i=1}^{n-1} \mathbb{E}[Y_{sn}|sn = i] \cdot \mathbb{P}(s_n = i) + \mathbb{E}[Y_{sn}|sn = n] \cdot \mathbb{P}(s_n = n) = \\ &= \sum_{i=1}^{n-1} (\mathbb{E}[X_1|X_i \geq \beta] - ic)(1-u)^{i-1}u + \alpha = \\ &= \sum_{i=1}^{n-1} \mathbb{E}[X_1|X_i \geq \beta](1-u)^{i-1}u - \sum_{i=1}^{n-1} ic(1-u)^{i-1}u + \alpha = \\ &= \mathbb{E}[X_1|X_i \geq \beta]u \sum_{i=1}^{n-1}(1-u)^{i-1} + cu \sum_{i=1}^{n-1} i(1-u)^{i-1}(-1) + \alpha.\end{aligned}$$

Remark 3. *Note that*

$$i(1-u)^{i-1} \cdot (-1) = \frac{d}{du}(1-u)^i.$$

Therefore by the above mentioned Remark and by Lemma A.3, we can write

$$\begin{aligned}
\mathbb{E}[Y_{sn}] &= \mathbb{E}[X_1|X_1 \geq \beta]u\sum_{i=1}^{n-1}(1-u)^{i-1} + cu\sum_{i=1}^{n-1}\frac{d}{du}(1-u)^i + \alpha = \\
&= \mathbb{E}[X_1|X_1 \geq \beta]\, u \cdot \frac{1-(1-u)^{n-1}}{u} + cu\frac{d}{du}\sum_{i=1}^{n-1}(1-u)^i + \alpha = \\
&= \mathbb{E}[X_1|X_1 \geq \beta]u \cdot \frac{1-(1-u)^{n-1}}{u} + cu\frac{d}{du}(\frac{1-(1-u)^n}{u} - 1) + \alpha = \\
&= \mathbb{E}[X_1|X_1 \geq \beta](1-(1-u)^{n-1}) + c \cdot \frac{[n(1-u)^{n-1}u - 1 + (1-u)^n]}{u} + \alpha.
\end{aligned}$$

We put the original value of α back in the equation and set $\gamma = \mathbb{E}[X_1|X_1 \geq \beta]$.

$$\begin{aligned}
\mathbb{E}[Y_{sn}] &= \gamma - \gamma(1-u)^{n-1} + cn(1-u)^{n-1} - \frac{c}{u} + \frac{c(1-u)^n}{u} + \mathbb{E}[X_1](1-u)^{n-1} + \\
&- nc(1-u)^{n-1} = \\
&= \gamma - \gamma(1-u)^{n-1} - \frac{c}{u} + \frac{c(1-u)^n}{u} + \mathbb{E}[X_1](1-u)^{n-1} = \\
&= \gamma - \frac{c}{u} - \gamma(1-u)^{n-1} + \frac{c(1-u)^n}{u} + \mathbb{E}[X_1](1-u)^{n-1}.
\end{aligned}$$

Remember (2.11), then

$$\begin{aligned}
\mathbb{E}[Y_{sn}] &= \beta - \gamma(1-u)^{n-1} + (\gamma-\beta)(1-u)^n + \mathbb{E}[X_1](1-u)^{n-1} = \\
&= \beta - \gamma(1-u)^{n-1} + \gamma(1-u)^n - \beta(1-u)^n + \mathbb{E}[X_1](1-u)^{n-1} = \\
&= \beta - \beta(1-u)^n - \gamma(1-u)^{n-1}(1-(1-u)) + \mathbb{E}[X_1](1-u)^{n-1} = \qquad (2.19) \\
&= \beta - \beta(1-u)^n - \gamma(1-u)^{n-1}u + \mathbb{E}[X_1](1-u)^{n-1} = \\
&= \beta - \beta(1-u)^n + (1-u)^{n-1}[\mathbb{E}[X_1] - \gamma u].
\end{aligned}$$

Remark 4. *Note that*

$$\mathbb{E}[X_1] - \gamma u = \mathbb{E}[X_1] - \mathbb{E}[X_1|X_1 \geq \beta]\mathbb{P}(X_1 \geq \beta),$$

and

$$\mathbb{E}[X_1] = \mathbb{E}[X_1|X_1 \geq \beta]\mathbb{P}(X_1 \geq \beta) + \mathbb{E}[X_1|X_1 < \beta]\mathbb{P}(X_1 < \beta),$$

therefore the equation becomes

$$\mathbb{E}[X_1] - \gamma u = \mathbb{E}[X_1|X_1 < \beta]\mathbb{P}(X_1 < \beta) = (1-u) \cdot \mathbb{E}[X_1|X_1 < \beta].$$

Finally, we can rewrite (2.19) as (2.13)

$$\begin{aligned}\mathbb{E}[Y_{sn}] &= \beta - \beta(1-u)^n + (1-u)^{n-1}(1-u) \cdot \mathbb{E}[X_1|X_1 < \beta] = \\ &= \beta - \beta(1-u)^n + (1-u)^n \cdot \mathbb{E}[X_1|X_1 < \beta] = \\ &= \beta - (1-u)^n(\beta - \mathbb{E}[X_1|X_1 < \beta]).\end{aligned}$$

□

Having proved Lemma 2.2.1, we will now continue with the explanation of the following proposition in order to prove the theorem. For this intent, consider:

- $X_i \sim Ber(p)$ of parameter $p = \frac{c}{1-\beta}$,
- $r = \sup\{i : 1 - ic > -c\}$ for fixed values of c,
- $Y_i \leq 0$ for all $i > r$,
- $D_n = \mathbb{E}[\max_{1 \leq i \leq n} Y_i] - \mathbb{E}[Y_{sn}]$.

Remark 5. *By definition, r is the highest possible i value for which $1 - ic > -c$, in other words, it is the least upper bound for the set $(-c, 1 - ic)$. We can affirm that $Y_1 > Y_{r+1}$ since when taking $r+1$ the worst value assumed by payoff $Y_1 (= -c)$ is bigger than the best value assumed by payoff $Y_{r+1} (= 1 - (r+1)c)$. Thus, Y_j cannot be the maximum payoff for values of $j > r + 1$.*

Proposition 2.2.2. *For all $n \geq r$, and c and p fixed:*

$$D_n \leq D_r = (1-p)^r (r-1)c. \tag{2.20}$$

Proof. We know D_n to be the difference between the expected payoff of a prophet with perfect foresight and the expected payoff of a gambler who is observing the sequence of random variables and stopping at time s_n. As we have successfully proved $\mathbb{E}[Y_{sn}]$, all is left to do now is to compute $\mathbb{E}[\max_{1 \leq i \leq n} Y_i]$.
Firstly:

- $\max\{Y_1, ..., Y_n\} = Y_i$ for $2 \leq i \leq r$,

- $\max\{Y_1, ..., Y_n\} = Y_1$.

As a result we have,

$$\begin{aligned}\mathbb{E}[\max_{1\leq i\leq n} Y_i] = \mathbb{E}[\max_{1\leq i\leq r} Y_i] = \sum_{j=2}^{r} \mathbb{E}[Y_j | \max_{1\leq i\leq r} Y_i = Y_j]\mathbb{P}(\max_{1\leq i\leq r} Y_i = Y_j) + \\ + \mathbb{E}[Y_1 | \max_{1\leq i\leq r} Y_i = Y_1]\mathbb{P}(\max_{1\leq i\leq r} Y_i = Y_1).\end{aligned} \tag{2.21}$$

To make the computation clearer we decompose equation (2.21) in three parts:

(i) For $2 \leq j \leq r$,

$$\mathbb{E}[Y_j | \max_{1\leq i\leq r} Y_i = Y_j] = 1 - jc,$$

(ii) For $2 \leq j \leq r$,

$$\begin{aligned}\mathbb{P}(\max_{1\leq i\leq r} Y_i = Y_j) &= \mathbb{P}(Y_1 = -c, ..., Y_{j-1} = -(j-1)c, Y_j = 1 - jc) = \\ &= \prod_{k=1}^{j-1} \mathbb{P}(Y_k = -kc) \cdot \mathbb{P}(Y_j = 1 - jc) = (1-p)^{j-1}p,\end{aligned}$$

(iii)

$$\begin{aligned}&\mathbb{E}[Y_1 | \max_{1\leq i\leq r} Y_i = Y_1]\mathbb{P}(\max_{1\leq i\leq r} Y_i = Y_1) = \\ &= (1-c)\mathbb{P}(Y_1 = 1-c) + (-c)\mathbb{P}(Y_1 = -c, Y_2 = -2c, ..., Y_r = -rc) = \\ &= (1-c)p - c(1-p)^r.\end{aligned}$$

We use our results to solve for $\mathbb{E}[\max_{1\leq i\leq n} Y_i]$:

$$\begin{aligned}\mathbb{E}[\max_{1\leq i\leq n} Y_i] &= \sum_{j=2}^{r}(1-jc)(1-p)^{j-1}p + (1-c)p - c(1-p)^r = \\ &= \sum_{j=2}^{r}(1-p)^{j-1}p - \sum_{j=2}^{r} jc(1-p)^{j-1}p + (1-c)p - c(1-p)^r.\end{aligned} \tag{2.22}$$

By applying A.3

$$\sum_{j=2}^{r}(1-p)^{j-1}p = p\sum_{j=2}^{r}(1-p)^{j-1} = p\sum_{k=1}^{r-1}(1-p)^{k} = p[\frac{1-(1-p)^r}{1-(1-p)} - (1-p)^0] =$$
$$= p[\frac{1-(1-p)^r}{p} - 1] = 1-(1-p)^r - p. \tag{2.23}$$

And

$$\begin{aligned}
\sum_{j=2}^{r} jc(1-p)^{j-1}p &= cp\sum_{j=2}^{r} j(1-p)^{j-1} = cp\sum_{j=2}^{r}[-\frac{d}{dp}(1-p)^j] = \\
&= -cp\frac{d}{dp}\cdot(\sum_{j=2}^{r}(1-p)^j) = \\
&= -cp\frac{d}{dp}\cdot(\frac{1-(1-p)^{r+1}}{1-(1-p)} - (1-p)^0 - (1-p)^1) = \\
&= -cp\frac{d}{dp}\cdot(\frac{1-(1-p)^{r+1}}{p} - 1 - 1 + p) = \\
&= -cp\cdot[\frac{-(r+1)(1-p)^r(-p) - (1-(1-p)^{r+1})}{p^2} + 1] = \\
&= -cp\cdot[\frac{(r+1)(1-p)^r}{p} + \frac{-1+(1-p)^{r+1}}{p^2} + 1] = \\
&= -c(r+1)(1-p)^r - \frac{(1-p)^{r+1}}{p}c + \frac{c}{p} - cp.
\end{aligned} \tag{2.24}$$

We use the solutions in (2.23) and (2.24) and rewrite equation (2.22).

$$\begin{aligned}
\mathbb{E}[\max_{1\leq i\leq n} Y_i] &= c(r+1)(1-p)^r + \frac{(1-p)^{r+1}}{p}c - \frac{c}{p} + cp + 1 - (1-p)^r - p + \\
&+ (1-c)p - c(1-p)^r = \\
&= (1-p)^r[c(r+1) - 1 - c + \frac{(1-p)}{p}c] - \frac{c}{p} + cp + 1 - p + (1-c)p.
\end{aligned} \tag{2.25}$$

Recall the definition of p from (2.12). Equation (2.2) becomes

$$\begin{aligned}\mathbb{E}[\max_{1\leq i\leq n} Y_i] &= (1-p)^r[cr+c-1-c+(1-p)(1-\beta)]-1+\beta+cp+1-p+\\ &+p-cp=\\ &=\beta+(1-p)^r[cr-1+(1-p)-(1-p)\beta]=\\ &=\beta+(1-p)^r[cr-1+1-p-c]=\beta+(1-p)^r[c(r-1)-\beta].\end{aligned} \tag{2.26}$$

Then, by Lemma 2.2.1 and (2.26)

$$\begin{aligned}D_n &= \mathbb{E}[\max_{1\leq i\leq n} Y_i]-\mathbb{E}[Y_{sn}]=\\ &=\beta+(1-p)^r[c(r-1)-\beta]-\beta+(1-u)^n(\beta-\mathbb{E}[X_1|X_1<\beta])=\\ &=(1-p)^r[c(r-1)-\beta]+(1-u)^n(\beta-\mathbb{E}[X_1|X_1<\beta])=\\ &=(1-p)^r c(r-1)-\beta(1-p)^r+(1-u)^n(\beta-\mathbb{E}[X_1|X_1<\beta]).\end{aligned}$$

To prove equation (2.20) we have to show that

$$-\beta(1-p)^r+(1-u)^n(\beta-\mathbb{E}[X_1|X_1<\beta])\leq 0.$$

Assuming $n\geq r$, since $(1-p)=\mathbb{P}(X_1=0)$ and $X_1=0$ when $X_1<\beta$, the above inequality can be proven as shown below

$$\begin{aligned}&\mathbb{E}[\beta-X_1|X_1<\beta]\mathbb{P}(X_1<\beta)^n-\beta(1-p)^n\leq\\ &\leq\mathbb{E}[\beta-X_1|X_1<\beta]\mathbb{P}(X_1<\beta)^n-\beta\mathbb{P}(X_1<\beta)^n=\\ &=\mathbb{P}(X_1<\beta)^n[\mathbb{E}[\beta-X_1|X_1<\beta]-\beta]=\\ &=-\mathbb{E}[X_1|X_1<\beta]\mathbb{P}(X_1<\beta)^n\leq 0.\end{aligned}$$

Therefore we can establish

$$D_n\leq D_r=(1-p)^r(r-1)c,$$

and demonstrate (2.8) since the maximum value of the equation is given for small values of p. By definition, p is minimised if $\beta = 0$. So we can reach to the final statement

$$D_n \leq D_r = (1-p)^r (r-1)c \leq (1-c)^r (r-1)c,$$

and by substituting $r = 1 + [\frac{1}{c}]$ we prove the theorem

$$D_n \leq D_r \leq \left[\frac{1}{c}\right] \, c \, (1-c)^{[1/c]+1}.$$

□

Conclusion

In this thesis, our aim was to present two different online selection games, the Secretary Problem in Chapter 1 and the Prophet Inequality in Chapter 2. Our objective was to understand their impact on decision-making processes and its application on real-life situations. We analysed how information can be managed and assessed in the most efficient way and looked into different versions of the problems. Furthermore, we were also interested in comparing the performance of two individuals at opposite ends of the knowledge spectrum: one with no prior information and the other with complete knowledge of future events.

Our analysis demonstrated that the implementation of the optimal stopping rule can significantly improve decision accuracy in online selection scenarios. We found that the strategies based on probability theories outperformed traditional methods. Specifically, we have discussed this issue in the context of finance and considered the situation where a stock analyst is asked to make an important decision after observing a sequence of price changes taking place during a given month. His task is to identify the peak day of the month during which stock prices will be highest. By applying Alexander's 5 percent filter technique for predicting successive price changes we confirmed the efficiency of the random walk model and supported the theory of independence of price movements. These findings are crucial as they provide a new framework for approaching real-time decision-making problems. However, it is important to keep in mind that the models we used for

the resolution of our problem is always based on the assumption of independent random variables.
To continue, knowing the peak of a random walk, which in our case is the stock price, is crucial in our examination. In actual practice, this information allows the short-selling investment strategy to achieve the highest possible profit. The short position strategy involves borrowing shares of a stock and selling them on the market to then buy them back later at a lower price. Being able to recognise when the stock price is going to be the highest allows the investor to earn from the decline in the stock price during the holding period. In simple words, the revenue from selling at a high price and buying back at a low price will be the highest.

In conclusion, this thesis offers mathematical proofs and demonstrations that enhance decision-making processes and improve strategies for maximising revenue. It also identifies the optimal stopping time and calculates the probability of success for the proposed strategy.

Appendix

Lemma A.1. *If X is a non-negative and real-valued random variable. Then*

$$\mathbb{E}[X] = \int_0^{+\infty} \mathbb{P}(X > x)dx$$

Theorem A.2 (Law of Total Probability)**.** *If $B_1, B_2, B_3, ...,$ is a partition of the sample space S, then for any event A we have*

$$\mathbb{P}(A) = \sum_i \mathbb{P}(A \cap B_i) \sum_i \mathbb{P}(A|B_i)\mathbb{P}(B_i)$$

In the special case where the partition is B and B^c

$$\mathbb{P}(A) = \mathbb{P}(A \cap B) + \mathbb{P}(A \cap B^c)$$

Lemma A.3. *For any $a \in \mathbb{R}$ with $a \neq 1$*

$$\sum_{i=0}^{n-1} a^i = \frac{1 - a^n}{1 - a}$$

Theorem A.4 (Law of Large Numbers)**.** *If $X_1, X_2, ...$ is a sequence of i.i.d. random variables, each with finite mean μ. Then*

$$\frac{X_1 + ... + X_n}{n} \xrightarrow{a.s.} \mu \qquad \text{for } n \to \infty$$

$$\mathbb{P}(\lim_{n\to\infty} \frac{X_1 + ... + X_n}{n} = \mu) = 1$$

Bibliography

[1] E. S. Cahn, *A difference prophet inequality for bounded i.i.d. variables, with cost for observations*, The Annals of Probability **20** (1992), no. 3, 1222–1228.

[2] S. N. Ethier, *The doctrine of chances: Probabilistic aspects of gambling*, Probability and Its Applications, Springer, Berlin and Heidelberg, 2010.

[3] E. F. Fama, *Random walks in stock market prices*, Financial Analysts Journal **51** (1995), no. 1, 75–80.

[4] T. S. Ferguson, *Who solved the secretary problem?*, Statistical science **4** (1989), no. 3, 282–289.

[5] M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm, *Automated online mechanism design and prophet inequalities*, AAAI, vol. 7, 2007, pp. 58–65.

[6] M. Hlynka and J.N. Sheahan, *The secretary problem for a random walk*, Stochastic Processes and their Applications **28** (1988), no. 2, 317–325.

[7] R. Kleinberg and S. M. Weinberg, *Matroid prophet inequalities and applications to multi-dimensional mechanism design*, Games and Economic Behavior **113** (2019), 97–115.

[8] B. Lucier, *An economic view of prophet inequalities*, ACM SIGecom Exchanges **16** (2017), no. 1, 24–47.

[9] D. G. Luenberger, *Investment science*, Oxford University Press, 2014.

[10] S.M. Ross and F. Morandin, *Probabilità e statistica per l'ingegneria e le scienze*, Idee e strumenti, Apogeo Education, 2015.