https://doi.org/10.26219/heal.aueb.4847

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Athens University of Economics and Business
Department Of Informatics

Demand Estimation and Stock Management with Sales Information

Skartsilas Panagiotis

Submitted in part fulfilment of the requirements for the degree of
Master in Business Mathematics of the
Athens University of Economics and Business, December 2022

https://doi.org/10.26219/heal.aueb.4847

https://doi.org/10.26219/heal.aueb.4847

Dedication

This thesis is dedicated to my favorite human on earth. Anthouli without you that project would have stopped on 10th page..

https://doi.org/10.26219/heal.aueb.4847

https://doi.org/10.26219/heal.aueb.4847

Contents

https://doi.org/10.26219/heal.aueb.4847

https://doi.org/10.26219/heal.aueb.4847

Chapter 1

Introduction

1.1 The Importance of Inventory Management

Nowdays retailers face numerous challenges when it comes to managing their inventory, particularly in industries where demand is highly uncertain. The decision of how much inventory to order is a critical one that can significantly impact the success or failure of a retail business. Ordering too little inventory can result in lost sales and missed opportunities, while ordering too much inventory can lead to high holding costs and reduced profitability.

One of the most well-known decision-making problems in the field of operations research is the newsvendor problem. The newsvendor problem originally arose in the newspaper industry, where retailers had to decide how many newspapers to order in advance without knowing the actual demand. Today, the newsvendor problem is relevant to many industries, including fashion, food, and electronics, where demand is highly uncertain.

There have been several examples of companies that have faced financial difficulties due to poor inventory management, including instances where poor newsvendor policy contributed to their financial problems. While it is difficult to isolate the impact of newsvendor policy on a company's bankruptcy, there have been studies and reports that highlight the importance of inventory management and its role in a company's financial health.

https://doi.org/10.26219/heal.aueb.4847

One example of a company that faced bankruptcy due to poor inventory management is Blockbuster. Blockbuster was a leading video rental chain that dominated the industry in the 1990s and early 2000s. However, as the demand for DVD rentals declined and the shift towards streaming services began, Blockbuster struggled to manage its inventory effectively. Blockbuster's newsvendor policy resulted in high inventory holding costs and lost sales due to insufficient inventory levels. Ultimately, Blockbuster filed for bankruptcy in 2010, with inventory management being cited as one of the key factors contributing to their financial difficulties.

The newsvendor model provides a valuable tool for retailers to make informed decisions about ordering inventory. The model allows retailers to determine the optimal order quantity of a perishable product with a fixed selling price, uncertain demand, and a known cost of ordering and holding inventory. The model considers the cost of holding excess inventory against the cost of not having enough inventory to meet demand.

The importance of the newsvendor problem for retailers cannot be overstated. The decision of how much inventory to order is a balancing act between the cost of holding excess inventory and the cost of not having enough inventory to meet demand. Retailers who order too much inventory face the cost of holding excess inventory, which can significantly impact their profitability. This includes the cost of storage, insurance, and the opportunity cost of tying up capital. On the other hand, retailers who order too little inventory face the cost of lost sales and the risk of losing customers to competitors.

In real life, the newsvendor problem has several practical implications. For instance, a retailer who orders too much inventory may have to sell the products at a discount to clear the excess inventory, which can result in lower profit margins. Similarly, a retailer who orders too little inventory may not be able to fulfill customer orders, resulting in lost sales and a damaged reputation.

Developing newsvendor models can help retailers make better decisions about ordering inventory. These models can help retailers determine the optimal order quantity by taking into account factors such as demand variability, lead time, and inventory holding costs. Newsvendor models can also help retailers identify the factors that affect demand and use this information

https://doi.org/10.26219/heal.aueb.4847

to make better decisions about inventory ordering.

In this thesis, we will examine different newsvendor policies and how well they operate compared to each other. We will explore the classical newsvendor model with one product with uncensored and uncertain demand. Three policies are going to used, Uncensored, Burnetas-Smith and the Kaplan-Meier Our goal is to provide a comprehensive understanding of newsvendor models and how they can be used to help retailers make better decisions about inventory ordering. By examining the strengths and weaknesses of different newsvendor policies, we aim to provide insights into the most effective strategies for managing inventory in uncertain demand environments.

https://doi.org/10.26219/heal.aueb.4847

1.2 Thesis Outline

This thesis has been organized into seven chapters. This section outlines the description of each chapter:

- In Chapter 2, we are giving the theoretical backround of the fundamental newsvendor model. By doing that, we building a strong base to understand the topics on newsvendor policies and models. From that chapter we also going to understand how the policy on uncensored demand is constructed
- In Chapter 3, we are creating the Burnetas-Smith policy which is the first about the censored demand
- In Chapter 4, showing the second policy for the censored demand, known as Kaplan Meier estimator.
- In Chapter 5, giving an outline of the models and algorithms we are going to use in order to perform our comparison tests
- In Chapter 7, we are implement the actual numerical tests
- In Chapter 8, giving the whole program as it was a great part of this thesis' work

https://doi.org/10.26219/heal.aueb.4847

Chapter 2

Newsvendor with Normal Continuous Demand

We will start the main part of this thesis, giving the fundamental theoretical background of the problem that we are going to look into. Consider that we have a retailer which sells a single product and has to meet a random demand D on a single period of time. It is important to mention that all the results on current chapter concern the uncensored demand. That means that we observe exact values of demand in every period.
The following notation is going to used:

$$\begin{aligned} D &: \text{Demand (continuous random variable, non-negative)} \\ f(x) &: \text{Density function} \\ F(x) &: \text{Cumulative distribution function} \\ Q &: \text{Number of purchased units/our decision variable} \\ c_o > 0 &: \text{Overage Cost-Remaining Inventory at the end of the period} \\ c_u > 0 &: \text{Underage Cost-Cost per unit of unsatisfied demand} \end{aligned}$$

Our purpose is to develop an expression for the order level at the beginning of the period

https://doi.org/10.26219/heal.aueb.4847

Q(decision variable)

In order to do that, we are starting by developing a cost function which depends on the demand and the order level. Next we determine the expected value of that function and the value of Q that minimizes that.

2.1 Optimal Order Quantity

Suppose that in the start of the period, the order level is defined and then a demand arrives. If the demand happens to be greater than our inventory level ($D \geq Q$) then we have unsatisfied demand and we pay an underage cost, which can form like below:

$$c_u \cdot \max\{D - Q, 0\}$$

On the other side if the demand is less that our inventory, then we have costs due to that remained unused inventory we bought but we did not sell:

$$c_o \cdot \max\{Q - D, 0\}$$

We can define $G(Q, D)$ as the total underage and overage cost that we observe at the end of every period.

$$G(Q, D) = c_o \cdot \max\{D - Q, 0\} + c_u \cdot \max\{D - Q, 0\} \tag{2.1}$$

Our goal now is to derive the expected cost function:

$$\begin{aligned} G(Q) = E[G(Q, D)] &= c_o \int_0^{\infty} \max\{0, Q - x\} f(x)dx + c_o \int_0^{\infty} \max\{0, x - Q\} f(x)dx \\ &= c_o \int_0^{Q} (Q - x) f(x)dx + c_u \int_Q^{\infty} (x - Q) f(x)dx \end{aligned}$$

https://doi.org/10.26219/heal.aueb.4847

We now have an expression of the the cost function G. Our purpose is to find the value of Q that maximizes that function, or equivalent take the first derivative:

$$\begin{aligned} \frac{dG(Q)}{dQ} &= c_o \int_0^Q 1 f(x)dx + c_o \int_Q^\infty (-1) f(x)dx \\ &= c_o F(Q) - c_u(1 - F(Q)) \end{aligned}$$

and from the second derivative:

$$\frac{d^2G(Q)}{dQ^2} = (c_o + c_u) f(Q)$$

which is greater than zero for every $Q \geq 0$ and as a result cost function is convex. The optimal solution of the first derivative is where the latter equals to zero:

$$G'(Q^*) = (c_o + c_u) F(Q^*) - c_u$$

or equivalently

$$F(Q^*) = \frac{c_u}{c_o + c_u} \tag{2.2}$$

Where $\dfrac{c_u}{c_o + c_u} \geq 0$ because c_o, c_u are positive numbers.
We call the $\dfrac{c_u}{c_o + c_u}$ critical ratio and it is the probability of satisfying all the demand of the period if order quantity is Q^*. For example if we have order quantity 25 and our critical ratio is 0.82, then the probability that the period demand is lower or at least equal 25, is 0.82.

https://doi.org/10.26219/heal.aueb.4847

2.2 Maximize Profit

In the previous section, we derived an expression for the optimal order quantity that minimizes the expected costs. However, as retailers, it is crucial to know our expected profit, especially under the optimal order quantity (Q^*). For this purpose we are going to follow the method and the proofs are presented on [3] (Silver et. al.)

For the needs of the profit function we have to define four more constants:

$$u : \text{Each unit purchase cost}$$
$$p : \text{Sold unit revenue}$$
$$g : \text{Salvage Value of disposed units}$$
$$B : \text{Lose of goodwill cost for unsatisfied demand}$$

Now, let as consider that we have a stock of order quantity Q and we observe demand D, the profit (P) under these parameters is:

$$P(Q, x_O) = \begin{cases} -Qv + px_O + g(Q - x_O), & \text{if } x_O \leq Q \\ -Qv + pQ - B(x_O - Q), & \text{if } x_O \geq Q \end{cases}$$

If our order quantity cannot satisfy the demand (second case) then our income comes from the sale of Q products but we have also to pay a cost of goodwill for the units we could have sold. If our order quantity can satisfy the demand completely (first case) then our income comes from the sale of D products and we have to pay a cost from possible unsold items.

Now if we set as x_O the current period demand, the expected value of profit can be expressed as

https://doi.org/10.26219/heal.aueb.4847

$$
\begin{aligned}
E\left[P(Q)\right] &= \int_0^\infty P(Q,x_0)f(x_0)dx_0 \\
&= \int_0^Q \Big[-Qv + px_0 + g(Q-x_0)\Big] f(x_0)dx_0 + \int_Q^\infty \Big[-Qv + pQ + B(x_0-Q)\Big] f(x_0)dx_0 \\
&= -Qv + gQ\int_0^Q f(x_0)dx_0 + (p-g)\int_0^Q x_O f(x_0)dx_0 + pQ\int_Q^\infty f(x_0)dx_0 \\
&- B\int_Q^\infty (x_0-Q)f(x_O)dx_0 \\
&= -Qv + gQ\int_0^Q f(x_0)dx_0 + (p-g)\Big[\int_0^\infty x_0 f(x_0)dx_0 - \int_Q^\infty (x_0-Q)f(x_0)dx_0 \\
&- Q\int_Q^\infty f(x_0)dx_0\Big] + pQ\int_Q^\infty f(x_0)dx_0 - B\int_Q^\infty (x_0-Q)f(x_0)dx_0 \\
&= -Qv + gQ\int_0^Q f(x_0)dx_0 + (p-g)\left[\hat{x} - \int_Q^\infty (x_0-Q)f(x_0)dx_0 - Q\int_Q^\infty f(x_0)dx_0\right] \\
&+ pQ\int_Q^\infty f(x_0)dx_0 - B\int_Q^\infty (x_0-Q)f(x_0)dx_0 \\
&= (p-g)\hat{x} - (v-g)Q - (p-g+B)\int_Q^\infty (x_0-Q)f(x_0)dx_0
\end{aligned}
$$

in conclusion:

$$
E\left[P(Q)\right] = (p-g)\hat{x} - (v-g)Q - (p-g+B)\int_Q^\infty (x_0-Q)f(x_0)dx_0 \tag{2.3}
$$

Before we move on further is useful there to try maximize the above equation in order to find an expression for the optimal order quantity using the four new parameters we defined earlier. We start from

$$
\frac{dE[P(Q)]}{dQ} = 0 \tag{2.4}
$$

and we need to calculate the derivative of the 2.3. In that calculation we have to find the derivative of the quantity $\int_Q^\infty (x_0-Q)f(x_0)dx_0$.

https://doi.org/10.26219/heal.aueb.4847

If we apply the Leibnitz's rule:

$$\frac{d}{dy}\int_{a(y)}^{a_2(y)} h(x,y)dx = \int_{a_1(y)}^{a_2(y)} [\frac{d}{dy}h(x,y)]dx_0 + h(a_2(y),y)\frac{da_2(y)}{dy} - h(a_1(y),y)\frac{da_1(y)}{dy}$$

we find that

$$\int_Q^\infty (x_0 - Q)f(x_0)dx_0 = F(Q) - 1 \tag{2.5}$$

and

$$\frac{dE[P(Q)]}{dQ} = 0 \Leftrightarrow -(\upsilon - g) - (p - g + B)[F(Q) - 1] = 0$$

or finally

$$F(Q) = \frac{p - \upsilon + B}{p - g + B} \tag{2.6}$$

Remember now that we need to find the expected profit for the optimal quantity Q^*. For Q^* the (2.3) gives:

$$\begin{aligned} E\left[P(Q)\right] &= (p-g)\hat{x} - (\upsilon - g)Q - (p-g+B)\int_Q^\infty (x_0 - Q)f(x_0)dx_0 \\ &= (p-g)\hat{x} - (\upsilon - g)Q^* - (p-g+B)\int_Q^\infty (x_0 - Q^*)f(x_0)dx_0 \\ &= (p-g)\hat{x} - (\upsilon - g)Q^* - (p-g+B)\int_{Q^*}^\infty (x_0 f(x_0)dx_0 + Q^*\int_{Q^*}^\infty f(x_0)dx_0 \\ &= (p-g)\hat{x} - (\upsilon - g)Q^* - (p-g+B)\int_{Q^*}^\infty (x_0 f(x_0)dx_0 + Q^*[1 - F(Q^*)] \end{aligned}$$

substituting $F(Q^*)$ with 2.6, we come up with the expression:

$$E[P(Q^*)] = (p-g)\hat{x} - (\upsilon - g)Q^* - (p-g+B)\int_{Q^*}^\infty x_0 f(x_0)dx_0 + (\upsilon - g)Q^* \tag{2.7}$$

$$= (p-g)\hat{x} - (p-g+B)\int_{Q^*}^\infty x_0 f(x_0)dx_0 \tag{2.8}$$

https://doi.org/10.26219/heal.aueb.4847

in this point, note that:

$$\begin{aligned}\int_0^{Q^*} x_0 f(x_0) dx_0 &= \int_{-\infty}^{Q^*} x_0 f(x_0) dx_0 \\ &= \int_{-\infty}^{+\infty} x_0 f(x_0) dx_0 - \int_{Q^*}^{+\infty} x_0 f(x_0) dx_0 \\ &= \hat{x} - \int_{Q^*}^{\infty} x_0 f(x_0) dx_0\end{aligned}$$

or

$$\int_{Q^*}^{\infty} x_0 f(x_0) dx_0 = \hat{x} - \int_0^{Q^*} x f(x_0) dx_0 \tag{2.9}$$

and using that expression we take from the (2.7):

$$E[P(Q^*)] = -B\hat{x} + (p - g + B) \int_0^{Q^*} x_0 f(x_0) dx_0 \tag{2.10}$$

https://doi.org/10.26219/heal.aueb.4847

2.3 Newsvendor with normal demand

We mentioned earlier that demand is a random variable that can be described from various distributions regarding the instences we are facing. The result from (2.2) can be applicable in all the distributions that demand can take. A specific distribution that is worth to examine is the normal distribution which we can observe and it is applicable in many real-world scenarios.

2.3.1 Optimal order quantity

As we said before, using the calculated value of the critical ratio $\dfrac{c_u}{c_o + c_u}$, we can find in which percentage of the normal distribution belongs. In other words we can find the value z of the distribution as we can see in the plot below. We now only need to find the specific point (Q^*) that we observe that z-value. It is known from the normal distribution, that stands $z = \dfrac{Q - \mu}{\sigma}$ or equivalently:

$$Q^* = \sigma z + \mu \tag{2.11}$$

where μ is the mean and σ the standard deviation of the normal distribution.
Now, from 2.6, substituting the z value we can take

$$\Phi(z) = \frac{p - v + B}{p - g + B}$$

and also we can express z as $\Phi^{-1}(z)$ and

$$Q^* = \Phi^{-1}\left(\frac{p - v + B}{p - g + B}\right)\sigma_x + \hat{x} \tag{2.12}$$

https://doi.org/10.26219/heal.aueb.4847

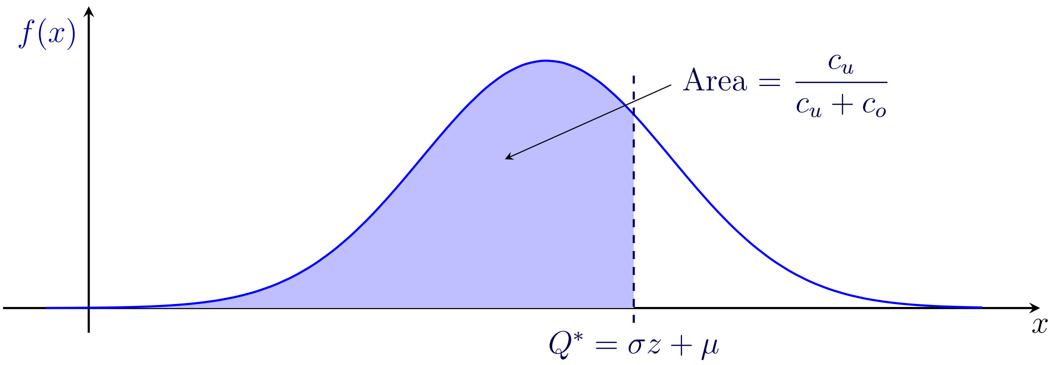

Figure 2.1: Graphical representation of optimal order quantity

2.3.2 Optimal Profit Function

Now, returning to the profit equation, we need to have an expression for the update normalized Q^*. If we substitute the (2.12) into (2.3):

$$E[P(Q)] = (p-g)\hat{x} - (u-g)Q - (p-g+B)\int_Q^\infty (x_0 - Q) f_x(x_0) dx_0 \Rightarrow$$

$$E[P(Q)] = (p-g)\hat{x} - (u-g)(\hat{x} + z\sigma_x) - (p-g+B)\int_{\hat{x}+z\sigma_x}^\infty (x_0 - \hat{x} - z\sigma_x)\frac{1}{\sigma_x\sqrt{2\pi}}\exp\left[\frac{-(x_0-\hat{x})^2}{2\sigma_x^2}\right]dx_0$$

we are mostly interested in the last quantity, which can be expressed as:

$$\int_{\hat{x}+z\sigma_x}^\infty (x_0 - \hat{x} - z\sigma_x)\frac{1}{\sigma_x\sqrt{2\pi}}\exp\left[\frac{-(x_0-\hat{x})^2}{2\sigma_x^2}\right]dx_0$$

it is useful there to substitute

$$u_0 = (x_0 - \hat{x})\sigma_x$$

now it is clear that $\frac{du_0}{dx_0} = \frac{1}{\sigma_x}$ or $dx_0 = \sigma_x du_0$.
Also, when $x_0 = \hat{x} + z\sigma_x$ then $u_0 = z$ and when $x_0 = \infty \Rightarrow u_0 = \infty$. So, the final value of the quantity is:

$$\sigma_x \int_z^\infty (u_0 - z)\frac{1}{\sqrt{2\pi}}\exp(-u_0^2/2)du_0 \tag{2.13}$$

https://doi.org/10.26219/heal.aueb.4847

To make things look simpler, we can set that quantity as:

$$G_u(z) = \sigma_x \int_z^{\infty} (u_0 - z) \frac{1}{\sqrt{2\pi}} \exp(-u_0^2/2) du_0 \tag{2.14}$$

and the expected profit becomes

$$E[P(Q)] = (p - g)\hat{x} - (\upsilon - g)(\hat{x} + z\sigma_x) - (p - g + B)\sigma_x G_u(k) \tag{2.15}$$

Remember that our objective there is to maximize the profit function. As we see from the above equation, the only term that depends on z, is the second term of the right-hand side, which we can set as $E[C(Q)]$:

$$E[C(Q)] = (\upsilon - g)z\sigma_x + (p - g + B)\sigma_x G_u(k) \tag{2.16}$$

on that point, from a standard normal distribution property (look on [3] appendix B), we can write:

$$\begin{aligned} G_u(z) &= \int_z^{\infty} (u_0 - z) f_u(u_0) du_0 \\ &= \int_z^{\infty} u_0 f_u(u_0) du_0 - z \int_z^{\infty} f_u(u_0) du_0 \\ &= f_u(z) - z(1 - \Phi(z)) \end{aligned}$$

and

$$\frac{E[C(Q)]}{\sigma_x} = (\upsilon - g)z + (p - g + B)\Big[f_u(z) - z(1 - \Phi(z)) \Big] \tag{2.17}$$

Which for the optimal values of Q^*, z^* becomes:

$$\frac{E[(Q^*)]}{\sigma_x} = (\upsilon - g)z^* + (p - g + B)\left[f_u(z^*) - z^* \left(\frac{\upsilon - g}{p - g + B} \right) \right] \tag{2.18}$$

$$= (p - g + B) f_u(z^*) \tag{2.19}$$

https://doi.org/10.26219/heal.aueb.4847

the last one stands because

$$\Phi(z) = \frac{p - \upsilon + B}{p - g + B} \Leftrightarrow 1 - \Phi(z) = \frac{\upsilon - g}{p - g + B}$$

and if we substitute 2.18 on the expected profit equation, we get:

$$\frac{E[P(Q^*)]}{\sigma_x} = \frac{(p - \upsilon)\hat{x}}{\sigma_x} - (p - g + B) f_u(z^*) \tag{2.20}$$

where $\hat{x}$ is the mean value of the mean distribution, σ_x the standard deviation and $f_u(z)$ the probability density function of the standard normal distribution.

https://doi.org/10.26219/heal.aueb.4847

Chapter 3

A first model for the cencored demand

3.1 Introduction

In the previous chapter, we introduced a news-vendor model for uncensored demand, where we assumed that we had full knowledge of the demand distribution in every period. However, in many instances, we do not know the demand distribution and may have censored demand samples, where demand values are less than our order quantity and their true value is uncertain. For example, suppose we set our inventory to 50 units for a period and sell all of them by the end of the period. At this point, we are unsure about the true level of demand, which could be 70, 100, or exactly 50. To address this problem, we use a model that does not require any specific distribution to calculate the period order quantity and can converge to optimal values even with censored demand. This model was developed by Burnetas and Smith in 2002. [5]

3.2 The model

To calculate the order quantity for each period, we add a small percentage of the previous order quantity. If we observe the exact demand, we subtract a small percentage of the previous order quantity instead. However, if we do not observe the exact demand, we add an uncertain value to the previous order quantity.

Burnetas and Smith, developed an adaptive ordering policy π that is consistent for all continuous demand distributions. The below policy is applicable even when there is no initial information on the stochastic properties of the demand.

Given a history H_{n+1}, policy π prescribes the order quantity for period $n+1$ according the following formula:

$$(\pi): \ Q_{n+1} = Q_n \left(1 - \frac{Y_n - R}{n}\right), \text{ n=1,2,..} \tag{3.1}$$

The order quantity in each period does not depend on the specific sales levels in previous periods but only on the random varuables Y_n, which indicate whether the demand was satisfied. In particular:

$$Y_n = \begin{cases} 1, & D_n < Q_n \\ 0, & D_n > Q_n \end{cases}$$

and R is the critical ratio as we defined later on:

$$R = \frac{c_u}{c_o + c_u}$$

It is also useful to mention that if $X_n = \dfrac{Q_{n+1} - Q_n}{Q_n}$ denotes the relative change in the order quantity from period n to $n+1$. Then, if $D_n \geq Q_n$:

$$Q_{n+1} = Q_n(1 + \frac{r}{n}) = Q_n + Q_n\frac{r}{n}$$

and

$$X_n = \frac{Q_n + Q_n\frac{r}{n} - Q_n}{Q_n} = \frac{r}{n}$$

https://doi.org/10.26219/heal.aueb.4847

On the other hand, if $D_n \leq Q_n$:

$$Q_{n+1} = Q_n(1 - \frac{1-r}{n}) = Q_n - Q_n\frac{1-r}{n}$$

and

$$X_n = \frac{Q_n - Q_n\frac{1-r}{n} - Q_n}{Q_n} = -\frac{1-r}{n}$$

and is now more clear that if all ordered product is sold out in period n, then the order for the next period is increased by a factor $\frac{r}{n}$, otherwise is decreased by $\frac{1-r}{n}$. In Burnetas-Smith(put the reference there) that Q_n converges in the R quantile of the demand distribution. Finally we can use that order quantity on the 2.3 expression to calculate the profit of each period.

https://doi.org/10.26219/heal.aueb.4847

Chapter 4

Kaplan-Meier

We will continue our discussion by introducing a second model that can handle censored demand. While the Burnetas-Smith method we discussed earlier only takes into account the last ordered quantity to calculate the current period's order up to a certain level, the model we will now explore considers all historical sales data. This method is the most extensively studied approach to deal with truncated or censored observations and is known as the Kaplan-Meier estimator. By using this widely recognized model, we can create an adaptive algorithmic approach that can handle stochastic demand without requiring knowledge of the distribution.

https://doi.org/10.26219/heal.aueb.4847

4.1 The KM estimator

The Kaplan-Meier estimator is a non-parametric method for estimating the survival function of a population based on right-censored data. It is also known as the product limit estimator and it is based on the definition of the survival function. Given a population of size N, the survival function at time t, denoted by $S(t)$, is defined as the probability that a randomly selected individual from the population will survive beyond time t. Formally,

$$S(t) = Pr(T > t)$$

where T is the time-to-event variable for an individual in the population.

In practice, the exact time of event for some individuals in the population may be unknown due to censoring. In such cases, the Kaplan-Meier estimator provides an estimate of the survival function based on the observed data.

The methodology of the Kaplan-Meier estimator involves dividing the follow-up time into discrete intervals and counting the number of individuals at risk of death at each interval. The number at risk is defined as the number of individuals who have not experienced the event and who have not been censored. The estimate of the survival function at each interval is calculated as the product of the previous estimate and the proportion of individuals still at risk of dying at that interval.

The Kaplan-Meier estimator provides a step-wise estimate of the survival function over time. It is a widely used method in medical research, particularly in the analysis of time-to-event data in clinical trials, and is a robust method that does not require knowledge of the underlying distribution of the time-to-event variable. However, it has some limitations, such as a lack of ability to estimate the hazard function and make inferences about the parameters of the distribution of the time-to-event variable.

In our case instead of survivals and individuals at risk we have truncated demand in each period. We are going to use the Kaplan-Meier estimator in a myopic policy which is based in the construction of the empirical CDF mentioned before.

4.2 The Model

The model that we are going to build is similar to the first we discussed for the uncensored demand. At the beginning of each period $t = 1, 2, ...$ a stocking decision Q_{km} is made on how many units to order. After the order is placed, Q_{km} units of supply arrive instantaneously and then a random demand D_t occurs. As in previous instances we assume that the per unit ordering cost is zero.

Now let F be the cumulative density function (CDF) of D, that is, $F(d) = P[D \leq d]$. Let $\bar{F}(d) = 1 - F(d)$ be the complementary CDF (CCDF) of D. For each $r \in (0, 1)$, let

$$\psi(r) = inf\{d : \bar{F}(d) \leq 1 - r\}$$

be the r-quantile of D.

As we have seen, in order to minimize the expected cost one should order the $r^* = \dfrac{b}{b+h}$-quantile of D, or equivalently the quantity $\psi^* = \psi(r^*)$ which called the newsvendor quantile.

In our problem structure, from every period's demand, we can only observe the sales or the quantity $Z_t = min\{D_t, Y_t\}$. We will also use an indicator function (δ_t) as in chapter 3 to indicate whether the demand is met$(\delta_t = 1)$ or not $(\delta_t = 0)$.

Now, jumping on the actual model implementation, let $\{D_{t=1}^{\infty}\}$ be the sequence of indepedent and identically distributed random variables according to distribution D. Let be $\{Y_t\}_{t=1}^{\infty}$ be a sequence of censoring variables. For each t, let $Z_t = min\{D_t, Y_t\}$ and $\delta_t = I[D_t \leq Y_t]$.

https://doi.org/10.26219/heal.aueb.4847

4.3 Algorithm

Consider T observations

$$\{(Z_1, \delta_1), ..., (Z_{t-1}, \delta_T\}$$

To construct the KM estimator, we order the T observations from smallest to largest according to their Z_t value. If two observations are equal, this with $\delta_t = 1$ are placed before the censored with $\delta_t = 0$. Let

$$\{(Z_{(1), \delta_{(t-1)}), ..., Z_{(T)}\delta_{(T)}})\}$$

be the order observations of the current period. Then for each $1 \leq t \leq T$, $Z_{(T) \leq Z_{(t+1)}}$, and if $Z_{(t)} = Z_{(t+1)}$ then $\delta_{(t)} \geq \delta_{(t+1)}$.
As we mentioned before, our main purpose there is to construct an empirical CDF taking in account the censored observations. To do that we can use the following product formula and compute the KM estimator for the CCDF $\bar{F}$. For each $d \leq Z_{(T)}$:

$$\bar{F}_T(d) = \prod_{t:Z_{(t)} \leq d} \left(\frac{T - t}{T - t + 1} \right) \tag{4.1}$$

https://doi.org/10.26219/heal.aueb.4847

Chapter 5

Methodology

In this chapter, our main goal is to analyze the methods and models that we will apply in the rest of the pages. We want to construct a system for a single retailer who sells a specific product over certain periods of time. As we have seen in previous chapters, the retailer's objective is to specify every period's order quantity based only on historic demand data. Therefore, we consider the single-period inventory model with demand sizes in consecutive inventory cycles being independent and identically distributed normal random variables with the same mean and variance. We will use Monte-Carlo simulation to generate period demands, varying the number for comparison purposes. After generating each demand value, we will calculate the optimal order quantity for that period and then the expected value of the profit.

We will use three policies based on three specific models, which we will compare with each other. The first model assumes that the demand is uncensored, and we have full information about its exact value, which is normally distributed.

The second and third policies are built for the case where the demand is censored, and we have zero information about the demand's distribution.

For each policy, we will provide a brief description of how it works and an algorithmic approach that we used to implement it.

https://doi.org/10.26219/heal.aueb.4847

Before we proceed to defining the policies, it is important to describe how our initial system will work. It will be based on a demand generator that generates samples from a normal distribution. The number of samples generated will correspond to the number of periods we are investigating. For example, if we need to analyze 500 periods, we will generate 500 samples with the same mean and standard deviation.

Using these random samples and applying the policies discussed in the previous sections, we will calculate the optimal order quantity and the expected profit for each period. Our main goal is to compare the classical newsvendor policy, the Burnetas-Smith policy, and the Kaplan-Meier policy with the optimal newsvendor quantile. To achieve this, we will use two graphical figures as our comparison tool.

The first figure will display the order-up-to level calculated by each policy against the period number. The second figure will display the profit values for each period. To understand how the mean and standard deviation of the distribution affects the results, we will vary these values and compare the results. We will also vary the constants that affect the ratio and newsvendor quantile.

To ensure precise and reliable results, we will repeat the simulation process multiple times and calculate the mean value of each repetition. The algorithmic approach for this simulation is as follows and will be repeated for each policy definition.

In the next algorithm we can see the actual implementation of the system that we are going to use as our environment.

Algorithm 1 System

```
Initialize constants and calculate newsvendor quantile
Generate T samples from normal distribution
for every sample do
    Calculate optimal quantity and profit for every one of discussed policies
end for
Calculate optimal Newsvendor quantity and profit for the normal demand
Store calculated values in a structure and repeat N times
Calculate the mean value of that structure
Based on previous step plot every policy's quantity and profit along with the optimal ones
```

https://doi.org/10.26219/heal.aueb.4847

5.0.1 First Policy π_0

Our first policy is the one we firstly define in the second section. We consider that the demand is uncensored and we do not re-stock our unsold items.

Let $D_1, D_2, ..., D_T$ be a sequence of random variables representing demand for T inventory cycles. We now will use two unbiased estimators to calculate the mean and standard deviation of all historical data we have until the current period:

$$\hat{\mu}_T = \frac{\sum_{t=1}^{T} D_t}{T}$$

and

$$\hat{\sigma}_T = \sqrt{\sum_{t=1}^{T} (D_t - \hat{\mu}_T)^2}$$

Now we use the values above to calculate every period's optimal order quantity by

$$\hat{Q}^*_{T+1} = \hat{\mu}_T + z_r \hat{\sigma}_T$$

It is known that the estimators of sample mean and standard deviation converge to the theoritical values and so does the optimal order quantity. So, as we are getting more and more historical data and observed sales we are expecting to see that 5.0.1 converges to the normal demand news-vendor quantity 2.12.

The algorithm that implements that policy is shown below. We have to mention that algorithm it has to run in every period repeatedly.

Algorithm 2 Newsvendor policy

1: Calculate estimator for mean
2: Calculate sigma estimator
3: Calculate the value from expression 5.0.1

https://doi.org/10.26219/heal.aueb.4847

5.0.2 Second policy π_1

Suppose now, that our information about the demand distribution is zero and our observation are not the whole demand but the original sales quantity. So, we are going to use the policy (π_1) for the instance that the demand is censored. In the first of those "censored-data friendly" policies we can only use the last calculated order-up-to level in order tho estimate the next period's order quantity. The expression is the following, where the Q_n is the previous period quantity, Q_{n+1} the current, Y_n is the indicator (censored or uncensored) and R the news-vendor quantile ratio.

$$Q_{n+1} = Q_n \left(1 - \frac{Y_n - R}{n}\right), \text{ n=1,2,...}$$

In Burnetas-Smith paper [5] we can see how that adaptive method it converges to the optimal news-vendor order-up-to level quantity. More specifically if Q^* and R^* be the optimal order quantity and the corresponding average profit, respectively, under perfect information (we will see later how are we going to approach the period's profit). The performance measure of Burnetas-Smith method for an adaptive price and order quantity policy is convergence with probability one to Q^* and R^* when the demand follows any probability distribution including the normal distribution. Let us here define that an adaptive policy π will be called consistent within a class of demand distributions if

$$\mathbf{P}^\pi \left[\lim_{n\to\infty} Q_n = Q^* \right] = 1$$

for any distribution in the specific class. In the Burnetas-Smith (give reference there) the convergence is shown by the proof of the upcoming lemma, which provides a sufficient condition for consistency in terms of convergence of the order quantities. When the optimal order quantity is unique, convergence of the order quantity to the optimal value implies convergence in profits.

https://doi.org/10.26219/heal.aueb.4847

Lemma 1 *Assume that there exists $Q_1^* \leq Q_2^*$ such that $R(Q) = R^*$ for $Q_1^* \leq Q \leq Q_2^*$. Then, any adaptive policy π that satisfies*

$$\boldsymbol{P}^{\pi}\left[Q_1^* \leq \lim_{n\to\infty} inf\, Q_n \leq \lim_{n\to\infty} sup\; Q_n \leq Q_2^*\right] = 1$$

is consistent

As in previous policies we also give here an implementation of the algorithm:

Algorithm 3 Burnetas-Smith Policy

if period demand <= order-up-to level **then**
    $Y = 0$
else
    $Y = 1$
end if
order quantity = previous quantity $* \left(1 - \frac{Y - \text{quantile}}{\text{period} + 1}\right)$

https://doi.org/10.26219/heal.aueb.4847

5.1 Third policy π_2

5.1.1 Kaplan-Meier policy

In this section we are going to describe an adaptive data-driven policy that use the KM estimator as we have seen it section 4. First of all, it is obvious that we need a starting point for our observation procedure. For that purpose we are going to use the optimal newsvendor quantity for the random distributed demand as we have defined it earlier in these pages. Now consider that we have a fixed time period horizon which is consisting of (k) periods of time. In the start of each period we make a decision of our period inventory ($Q_k m$) and then a demand comes up. The actual observation which is denoted as (Z_t) of any period is the smallest quantity between period demand and order-up-to level: After setting all the observed sales in order, we use the product 4.1 to estimate the corresponding estimated probability for every sale point ($Z_{(1)}, .., Z_{(T)}$). In other words in every period (k) we are building a new CCDF. Once that CCDF is created we use the newsvendor quantile to calculate the optimal point as regarded to the estimated CCDF. Now this is the next period's decided order up to level quantity.

The paper of Levi et al. [6] it is suggested a slight modification in the original algorithm and method of KM estimator. More specifically we maintain the largest uncensored sales data point in a temporary variable and we keep track of the number of times that Y_t^{KM}-the estimated $\frac{C_o}{C_o + C_u}$ quantile under the Kaplan-Meier empirical distribution-equals to the largest data point, and among those time periods, how often we get censored observations from using the largest data point exceeds $\frac{C_o}{2(C_u + C_o)}$. The last condition suggests that our order-up-to levels are too low and when this happends, we double the magnitute of the largest data point and use that as the next period order-up-to level.

Also because of the lack of data on the early periods, it is handy to maintain a single fictitious data point during the first 20 periods. This data point is the same as the initial quantity we used on the start of the method. That last modification help us to make a more accurate model, without affect the convergence of the method.

In the Levi et al. [6] it is showed that the method of KM estimator converges to the true values

https://doi.org/10.26219/heal.aueb.4847

of quantity and profit. That stands with the help of the following theorem:

Lemma 2 *Let* $D_{t t=1}^{\infty}$ *be a sequence of non-negative, integer-valued, indepedent and identically distributed random variables according to a distribution D. Let* $Y_{t t=1}^{\infty}$ *be a sequence of censoring variables, such that t, the censoring variable* Y_t *is* $F_t = \sigma(D_1, ..., D_{t-1})$ *measurable. For each t, let* $Z_t = min \;\; D_t, Y_t$ *and* $\delta_t = \boldsymbol{I}[D_t \leq Y_t]$*. Let* $R = max\ d : P[Y_t \geq d] = 1$*. Then for each* $d \in [0, R]$ *the KM estimator converges surely to the true respective values as* $t \to \infty$

And as before, that implies that the profit function converges in the real value as well for a period close to infinity.

That algorithm for the Kaplan-Meier policy is a little bit more complicated and expensive.

Algorithm 4 Kaplan-Meier

```
put min(perioddemand, orderlevel) in right position
Calculate CCDF for current period, using 4.1
find observation closest to newsvendor quantile
if period less than 20 then
    maintain initial order quantity
else
    find observation closest to newsvendor quantile
end if
if (times that largest quantity = order up to level) / (times that get censored observations) => C_o / (2(C_o + C_u)) then
    double order-up-to-level
end if
```

https://doi.org/10.26219/heal.aueb.4847

5.1.2 Profit Expression

In that point we can finally define an expression for the profit. The bellow function calculates each period's profit, and as we saw in section 2 depends only in order-up-to level quantity and the real demand value:

$$P(Q, x_O) = \begin{cases} -Qv + px_0 + g(Q - x_0), & \text{if } x_0 \leq Q \\ -Qv + pQ + B(x_0 - Q), & \text{if } x_0 \geq Q \end{cases}$$

where Q is order-up-to level quantity calculated in the last period and x_0 the actual demand of the current period "T".
As we mentioned before, that function converges to the theoretical value of the profit in the instance of normal distributed demand, which is calculated from 2.20.

https://doi.org/10.26219/heal.aueb.4847

Chapter 6

Numerical results

On this chapter we are going to present the actual results from the policies implementation. As we mentioned before, in order to have a better picture of the performance, we will generate three different plots for each instance:

- Quantity plot
- Percentage difference from quantity newsvendor benchmark
- Profit plot

and each plot contains all three policies along with the newsvendor benchmark.

We are also interesting to see how the cost parameters affecting the performance of the models, and for that purpose we are going to make several tests with different parameter values. As for the system, in order to identify how the demand itself affects the models we are going to test with multiple mean and std values.

https://doi.org/10.26219/heal.aueb.4847

6.0.1 First set of figures

We are going to generate the first set of figures for a normal distribution with $\mu = 200$ and $\sigma = 100$.

$$D_t \sim N(200, 100)$$

For this example the overage cost is going to be $C_o = 9.5$ and the underage $C_u = 36$. The total number of periods is 500 and the problem instances (or repetitions) are 50.

a) figure of the difference between policies and optimal newsvendor quantity

b) figure of the actual quantity

c) figure of the profit functions

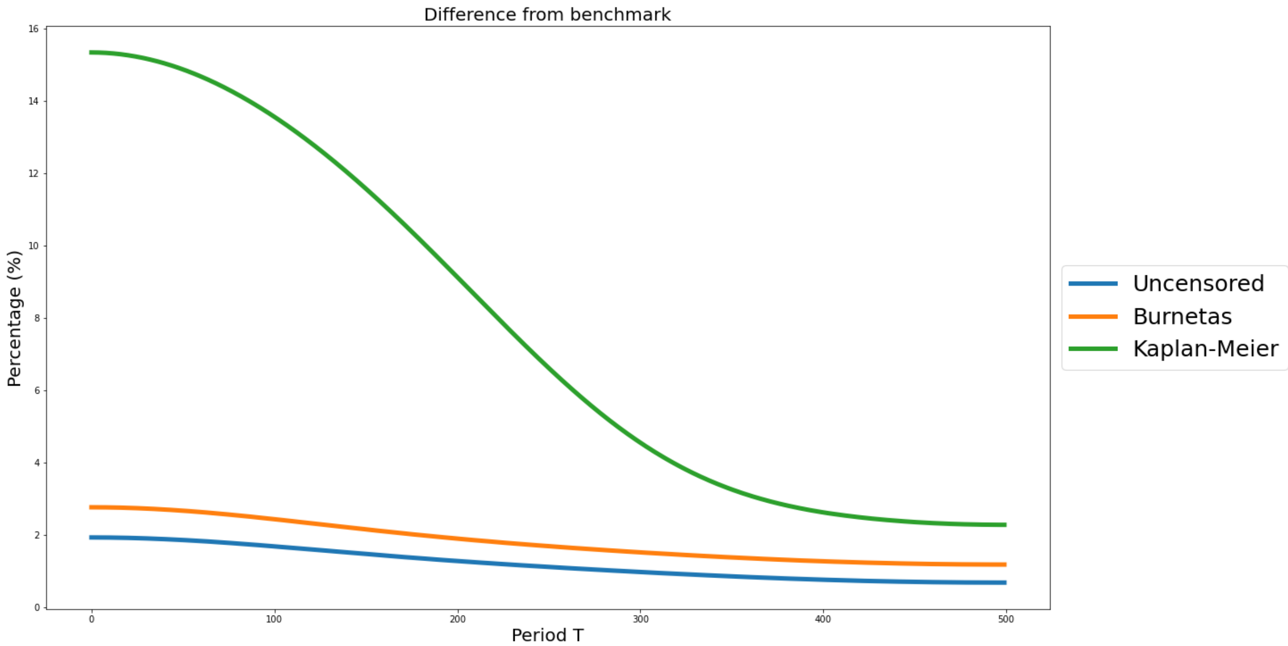

Figure 6.1: Difference from the Benchmark. Mean value is 200, standard deviation is 100 Underage cost is 9.5 and Overage cost is 31

https://doi.org/10.26219/heal.aueb.4847

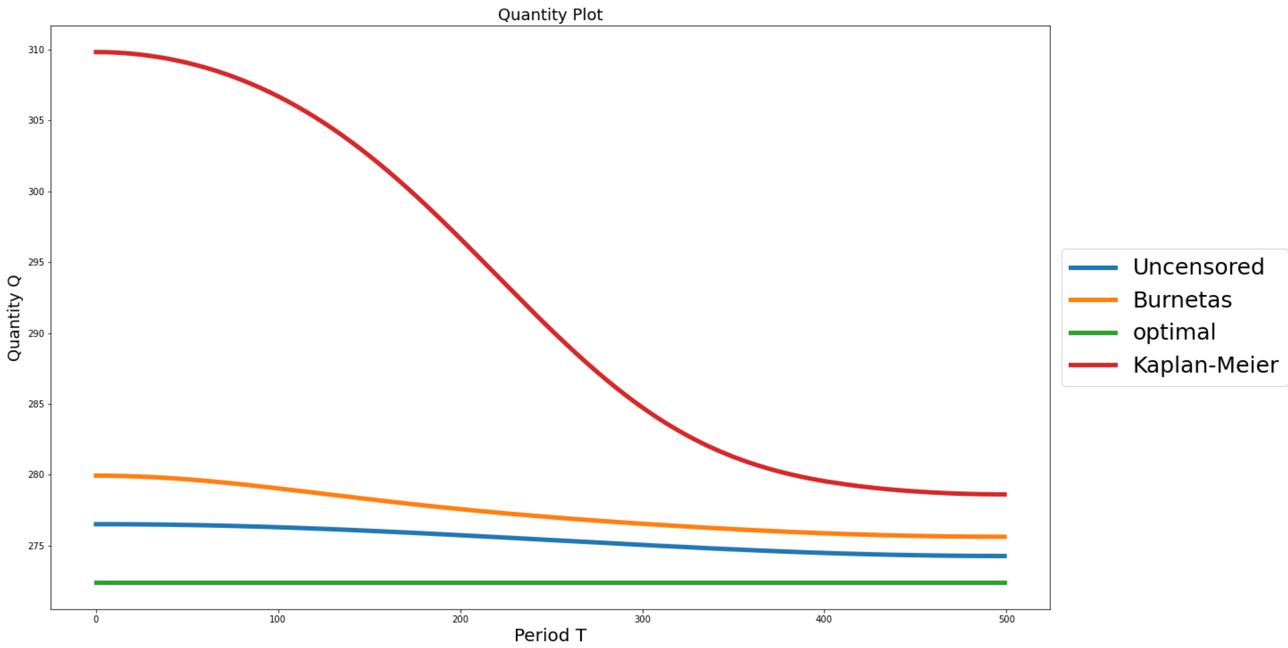

Figure 6.2: Quantity Plot. Mean value is 200, standard deviation is 100
Underage cost is 9.5 and Overage cost is 31. Burnetas and Uncensored policies are converge very early on benchmark

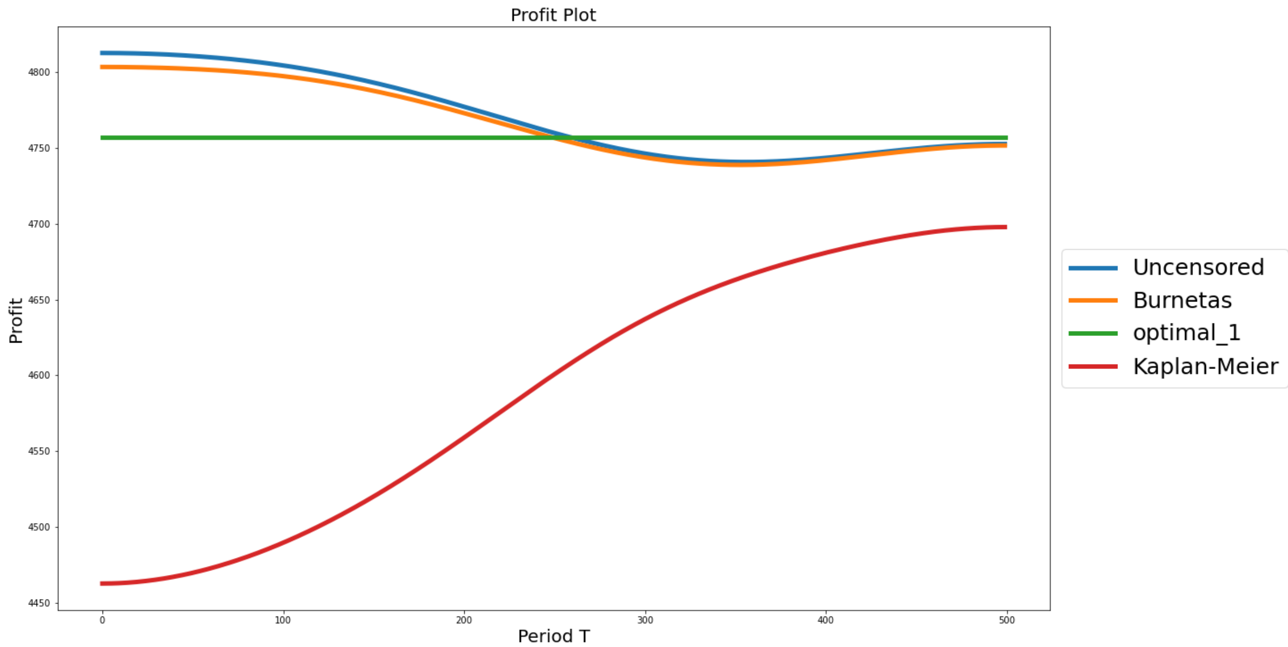

Figure 6.3: Profit Plot. Mean value is 200, standard deviation is 100
Underage cost is 9.5 and Overage cost is 31. Burnetas and Uncensored policies are converging after small fluctuation on profit benchmark

https://doi.org/10.26219/heal.aueb.4847

Upon examining the plots of our policies, it is evident that the Newsvendor and Burnetas-Smith models are performing exceptionally well, with early convergence to the Newsvendor benchmark. Specifically, both models closely match the benchmark from the first periods, with a difference of almost 0%. Regarding the Kaplan-Meier estimator, we observed poor performance during the initial periods, but as the number of periods approached 500, it also converged. This behavior was expected since the small amount of data at the beginning makes the estimator less efficient, despite maintaining the initial quantity.

In the next implementation, we used the exact same distribution parameters, problem instances, and total number of periods, but we slightly changed the cost values. We set the overage cost to $C_o = 21$ and the underage cost to $C_u = 19.5$. Our approach was to observe how well the policies perform when the costs are close to each other.

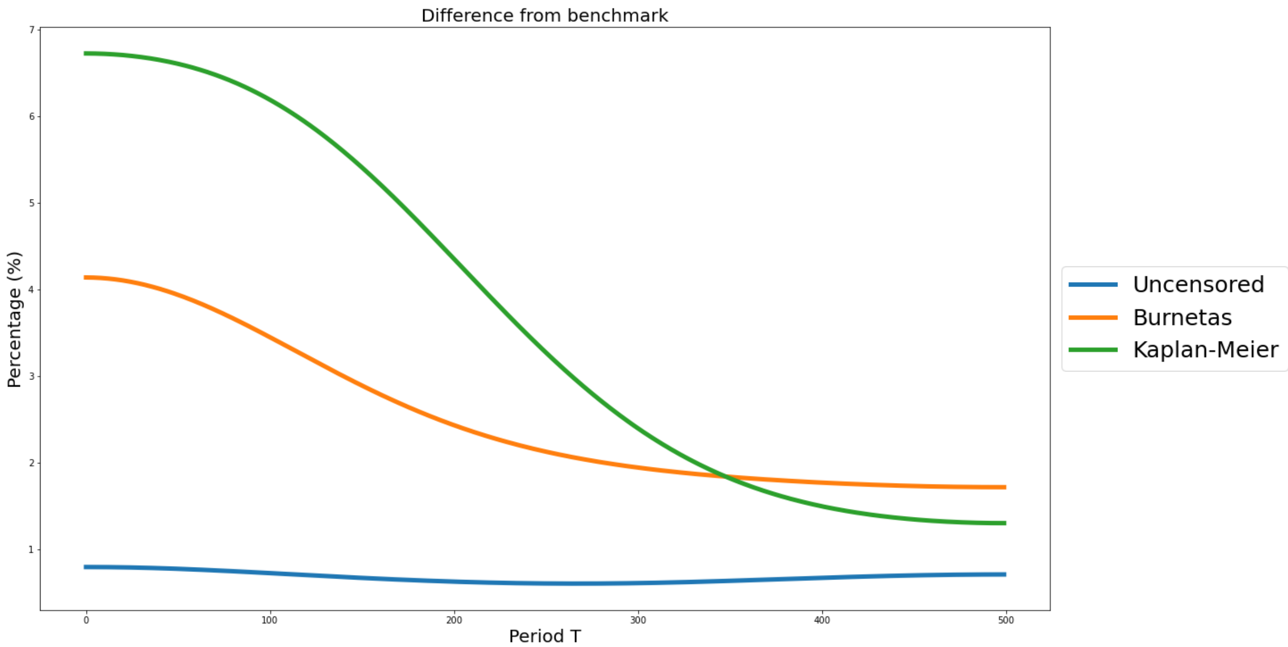

Figure 6.4: Difference from the Benchmark. Mean value is 200, standard deviation is 100 Underage cost is 19.5 and Overage cost is 21. Again uncensored policy is the best and Kaplan-Meier has the worst convergence. Burnetas-Smith has some deviation from last set. The percentage have risen on 4%

https://doi.org/10.26219/heal.aueb.4847

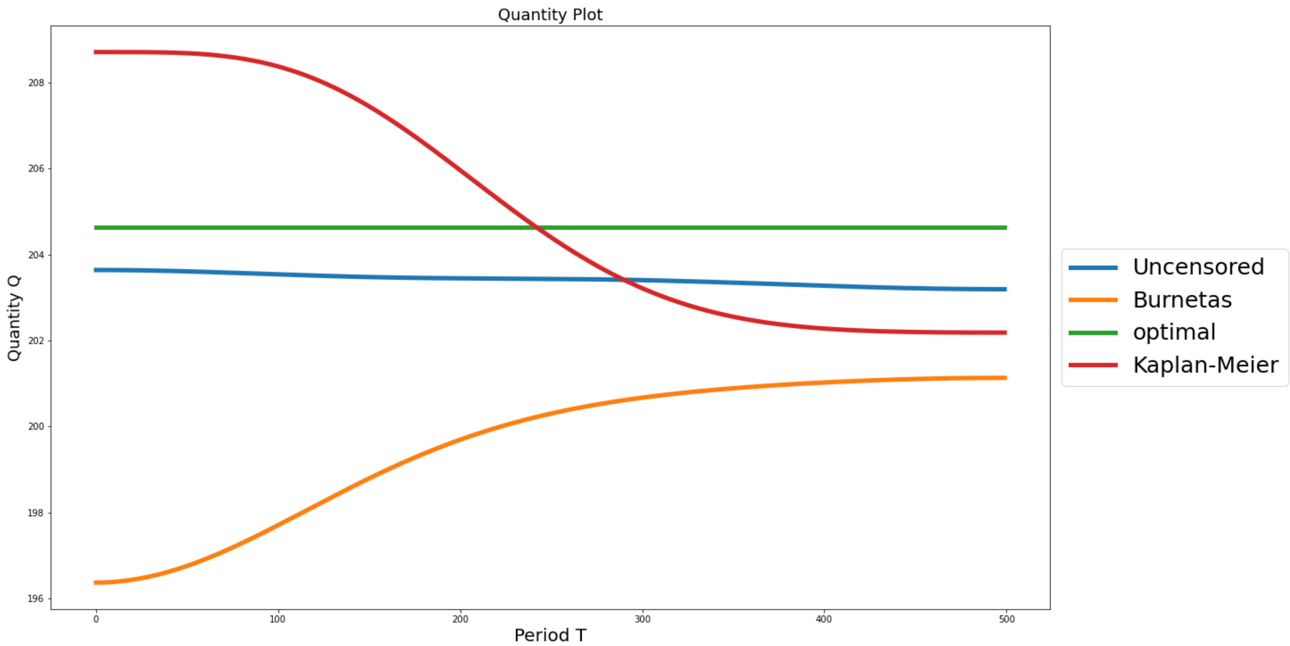

Figure 6.5: Quantity Plot. Mean value is 200, standard deviation is 100
Underage cost is 19.5 and Overage cost is 21

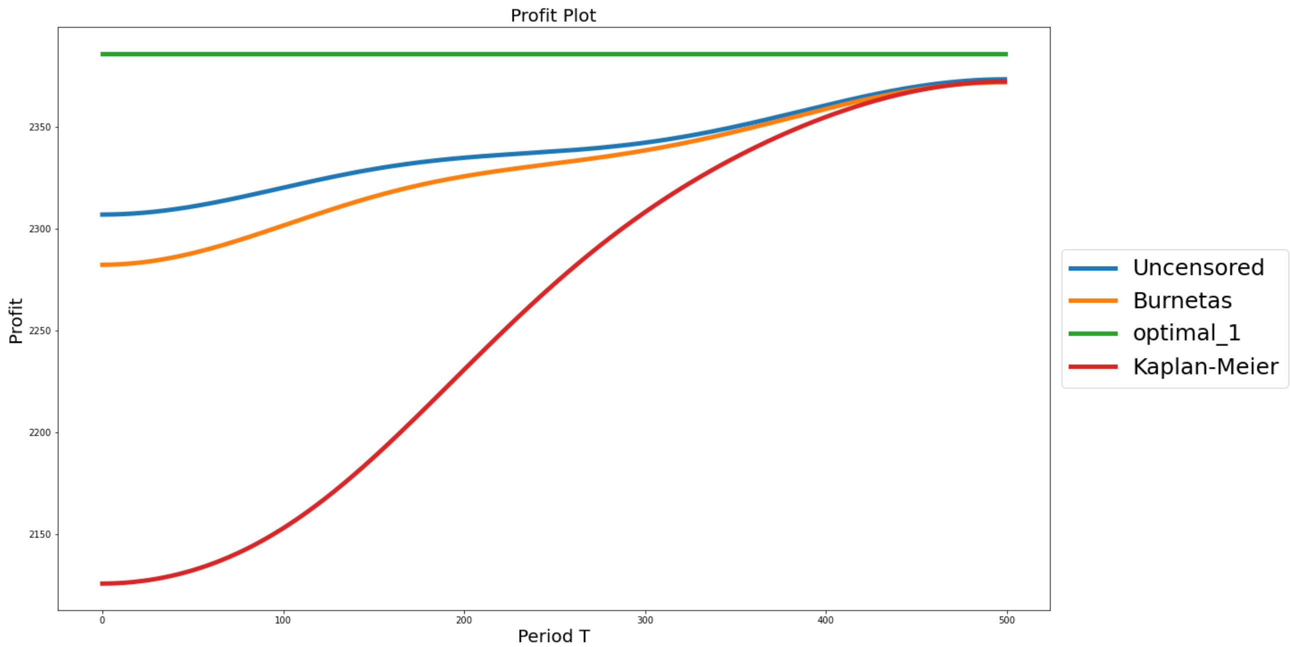

Figure 6.6: Profit Plot. Mean value is 200, standard deviation is 100
Underage cost is 19.5 and Overage cost is 21.

As we can see, the convergence rate of the uncensored policy is affected by the small difference between the costs. The Burnetas-Smith policy is slightly different compared to the last set. Regarding the Kaplan-Meier policy, we can observe a slight change in performance with a slower convergence rate. However, the difference from the benchmark is relatively small for all policies.

https://doi.org/10.26219/heal.aueb.4847

For the last test with the same distribution parameters, we will have a larger underage cost of $C_u = 29.5$ compared to the overage cost of $C_o = 11$. In this implementation, we observed some interesting results. When the underage cost is relatively larger than the overage, and the ratio is below 0.5, the Kaplan-Meier policy performs better than the Burnetas-Smith policy. The latter seems to converge very slowly to the benchmark, with a difference of over 10% during the first periods. On the other hand, Kaplan-Meier is closer to the benchmark from the beginning. One can conclude that relatively larger values of underage costs help Kaplan-Meier perform better but have the opposite influence on the Burnetas-Smith policy.

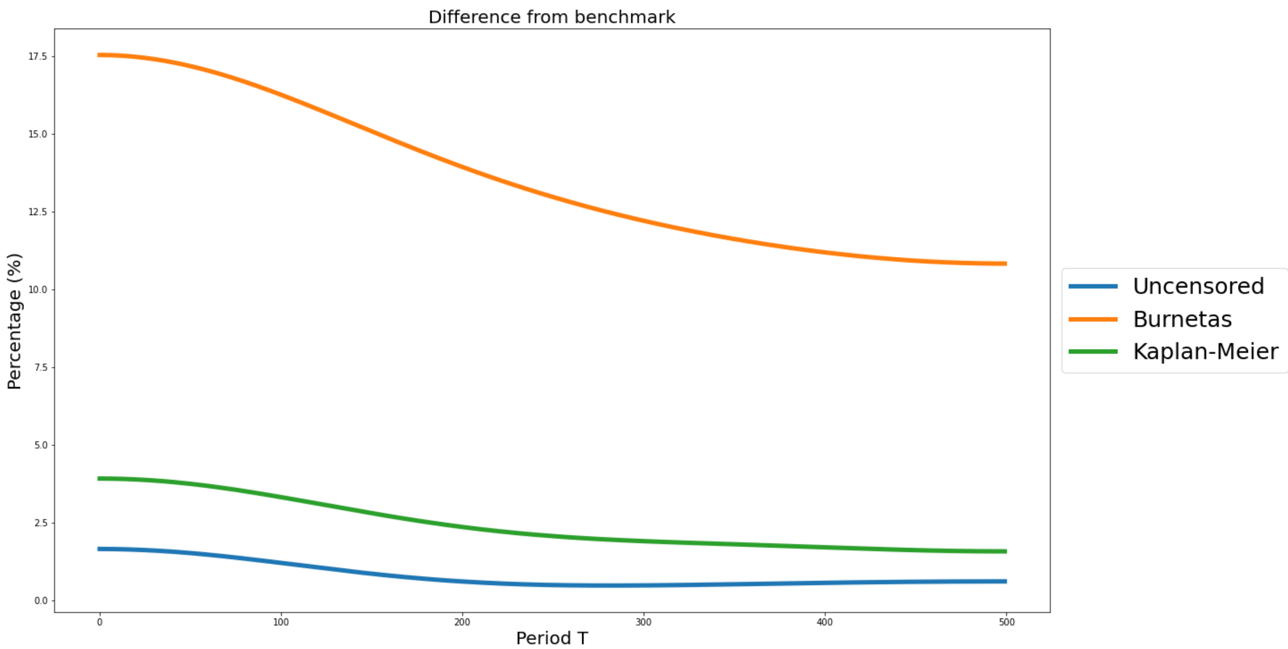

Figure 6.7: Difference from the Benchmark. Mean value is 200, standard deviation is 100 Underage cost is 29.5 and Overage cost is 11. Here Kaplan-Meier and Uncensored give better results. Burnetas has the biggest difference from the benchmark

https://doi.org/10.26219/heal.aueb.4847

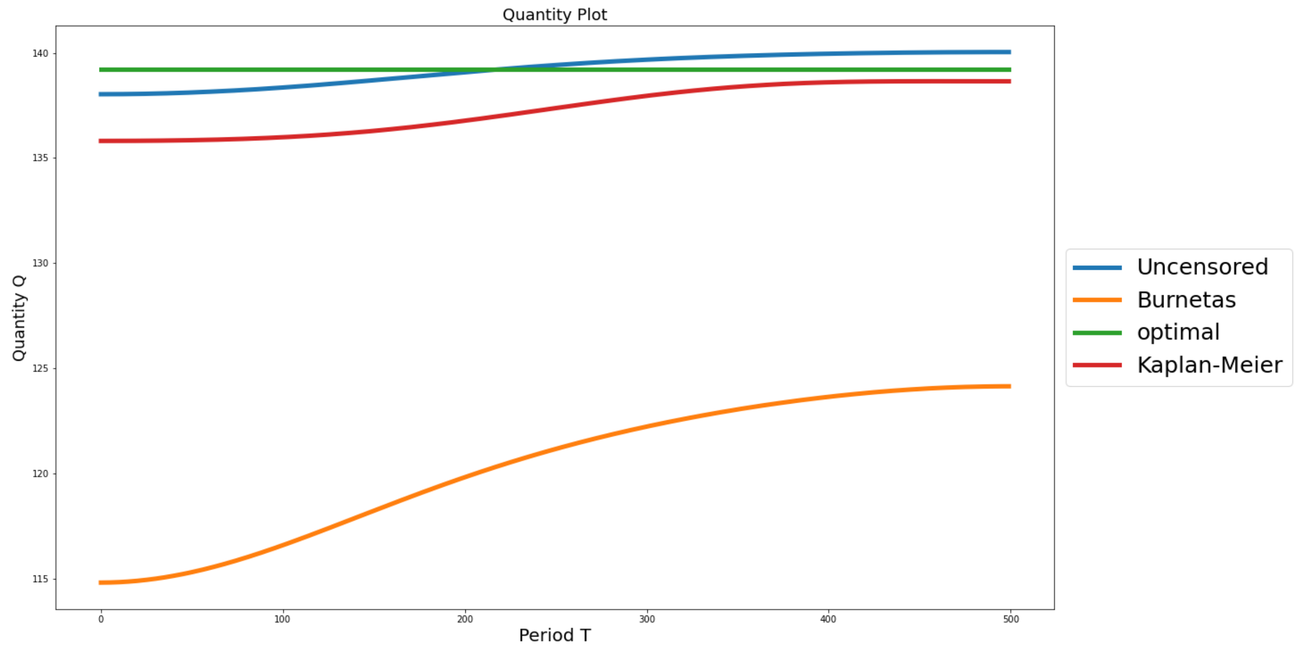

Figure 6.8: Quantity Plot. Mean value is 200, standard deviation is 100
Underage cost is 29.5 and Overage cost is 11. It is clear that Burnetas-Smith is the policy with the worst score in matter of quantity in comparison to the optimal newsvendor quantity

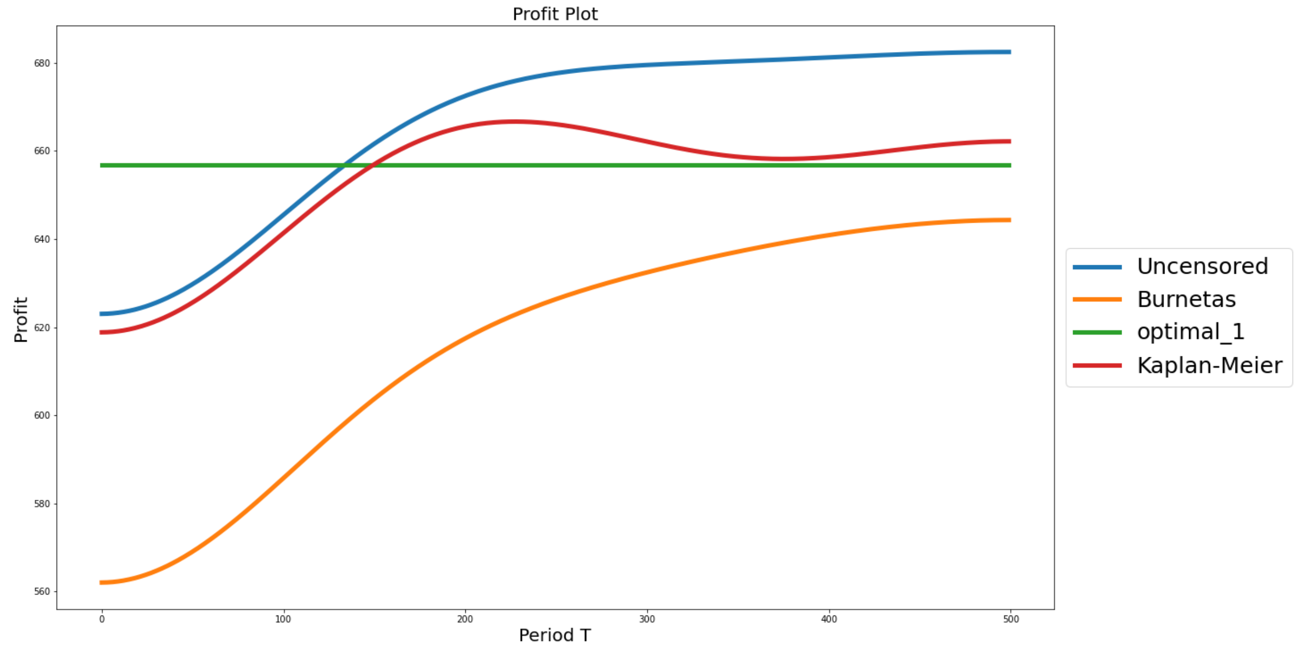

Figure 6.9: Profit Plot. Mean value is 200, standard deviation is 100
Underage cost is 29.5 and Overage cost is 11.

https://doi.org/10.26219/heal.aueb.4847

6.0.2 Second set of figures

We will now test how changes in the distribution parameters affect the performance of our models. For this purpose, we will fix the values of the cost parameters and only modify the mean and standard deviation values. We will fix the overage and underage cost at $C_o = 9.5$ and $C_u = 31$, respectively. We will use two different sets of values for the distribution parameters. Using various distribution parameters can be useful for a number if reasons.

Firstly, it allows us to assess the robustness of a model. By testing the model with different distribution parameters, we can see how the model performs under different scenarios and how sensitive the model is to changes in the underlying distribution. This can help us to identify any weaknesses in the model and to make improvements to ensure that the model is reliable and accurate.

Secondly, it can help us to gain a better understanding of the relationship between the distribution parameters and the model outputs. By varying the distribution parameters and observing the effect on the model outputs, we can gain insights into how the model works and how the different parameters affect the results. This can help us to refine the model and to develop a deeper understanding of the system being modeled.

We will perform the following test comparison for that purpose:

- $\mu = 200,\ \sigma = 100$
- $\mu = 200,\ \sigma = 50$

https://doi.org/10.26219/heal.aueb.4847

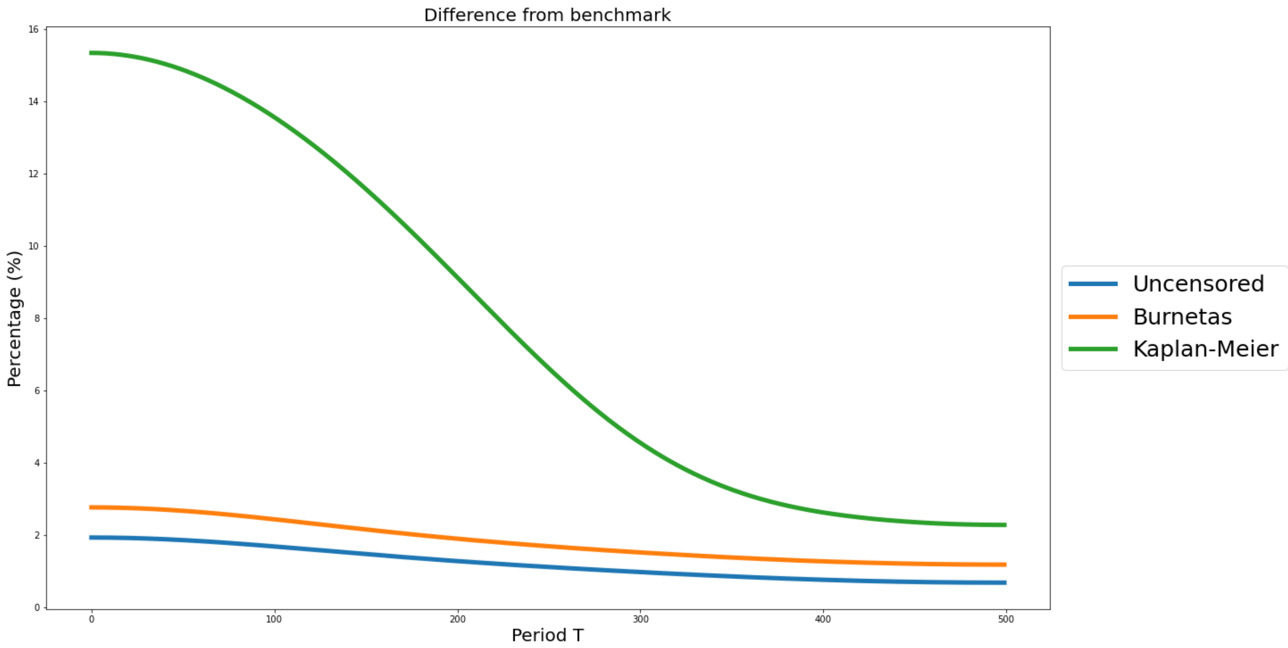

Figure 6.10: Difference from the Benchmark. Mean value is 200, standard deviation is 100 Underage cost is 9.5 and Overage cost is 31

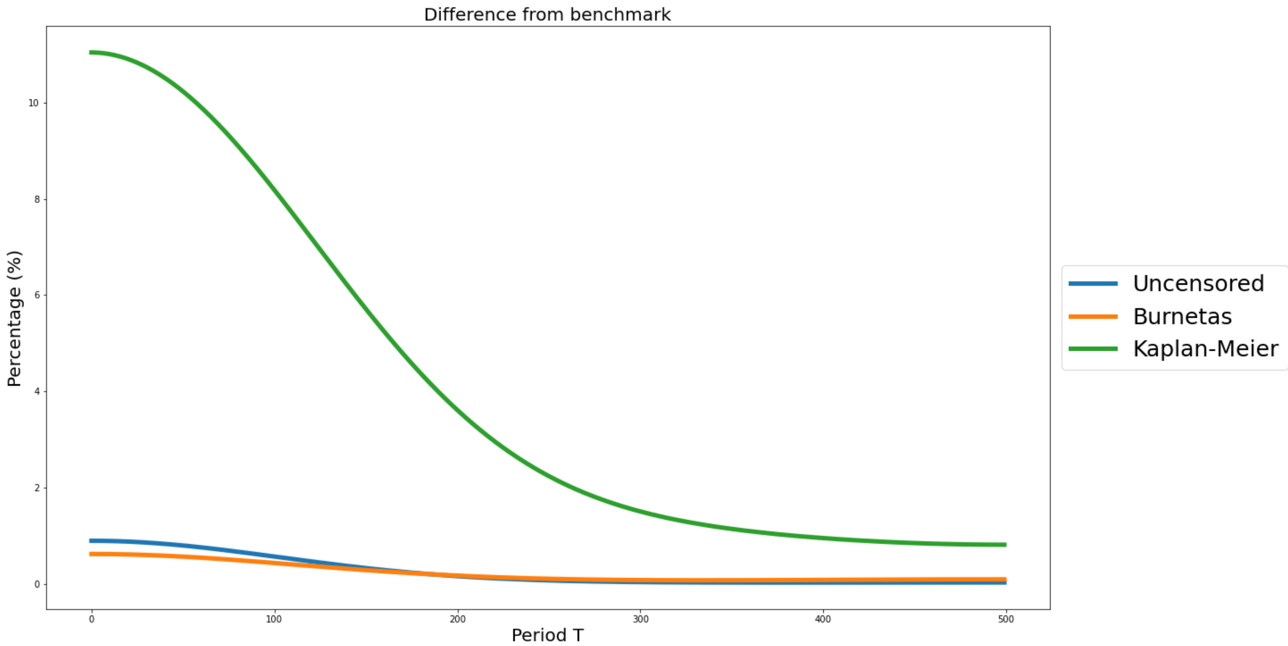

Figure 6.11: Difference from the Benchmark. Mean value is 200, standard deviation is 50 Underage cost is 9.5 and Overage cost is 31

https://doi.org/10.26219/heal.aueb.4847

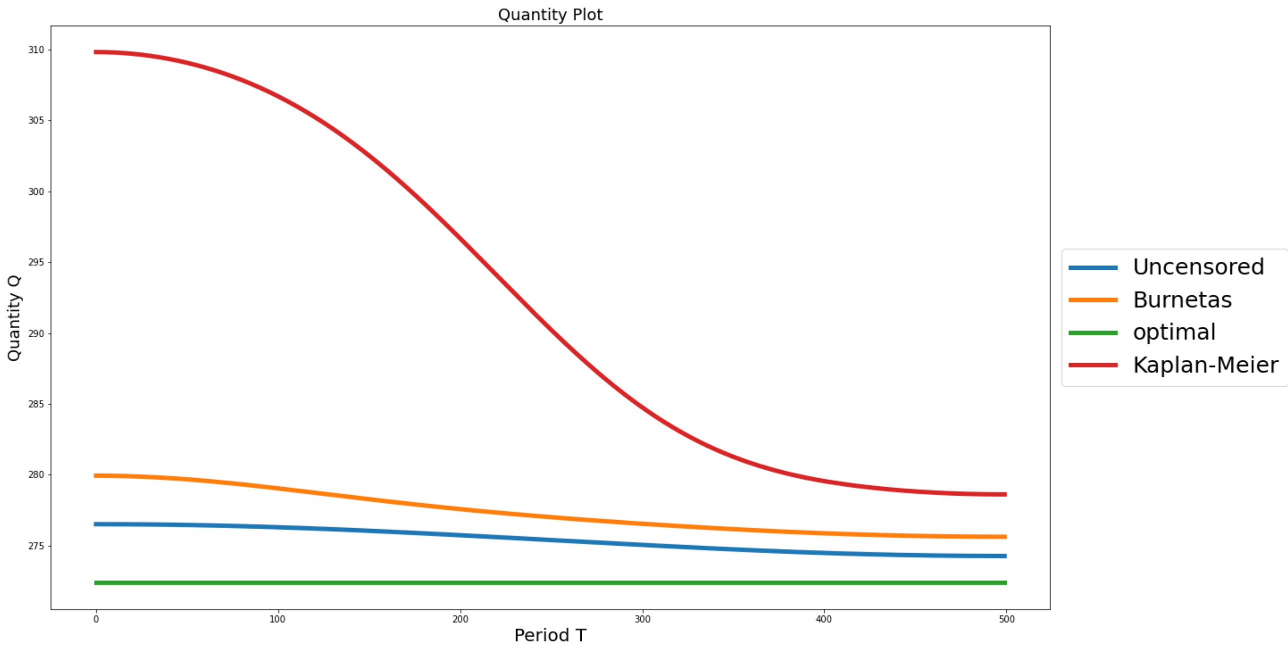

Figure 6.12: Quantity Plot. Mean value is 200, standard deviation is 100
Underage cost is 9.5 and Overage cost is 31. Burnetas and Uncensored policies are converge very early on benchmark

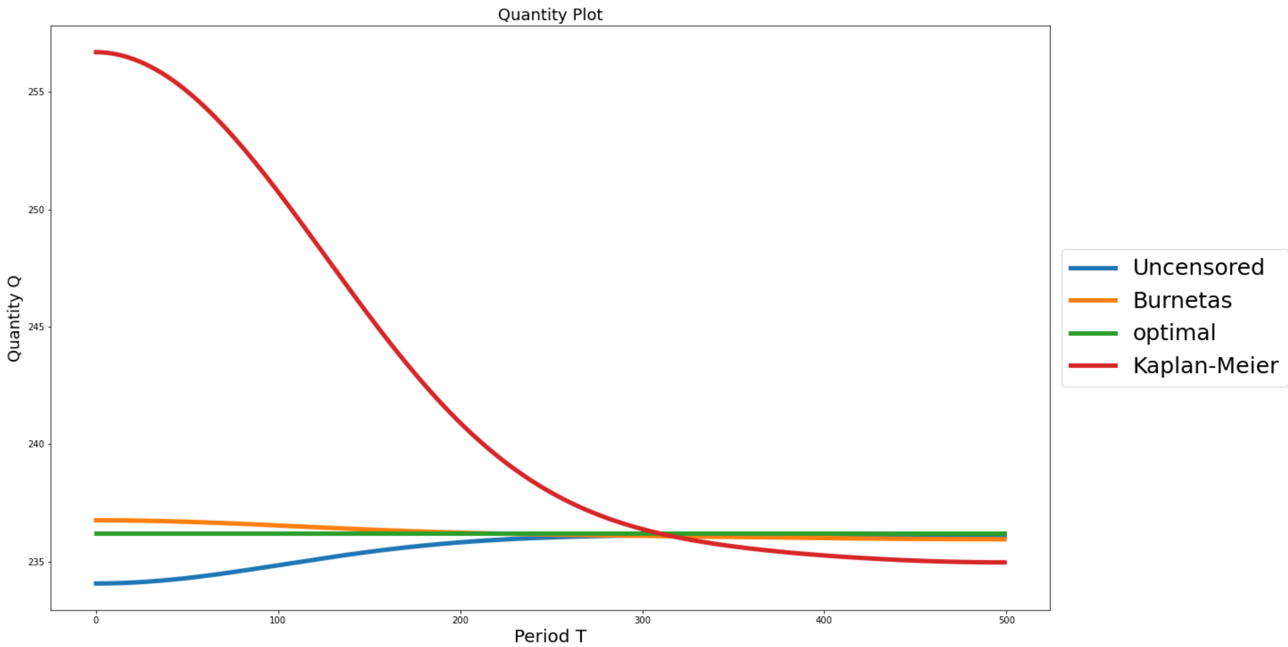

Figure 6.13: Quantity Plot. Mean value is 200, standard deviation is 50
Underage cost is 9.5 and Overage cost is 31. Burnetas and Uncensored policies are converge very early on benchmark

https://doi.org/10.26219/heal.aueb.4847

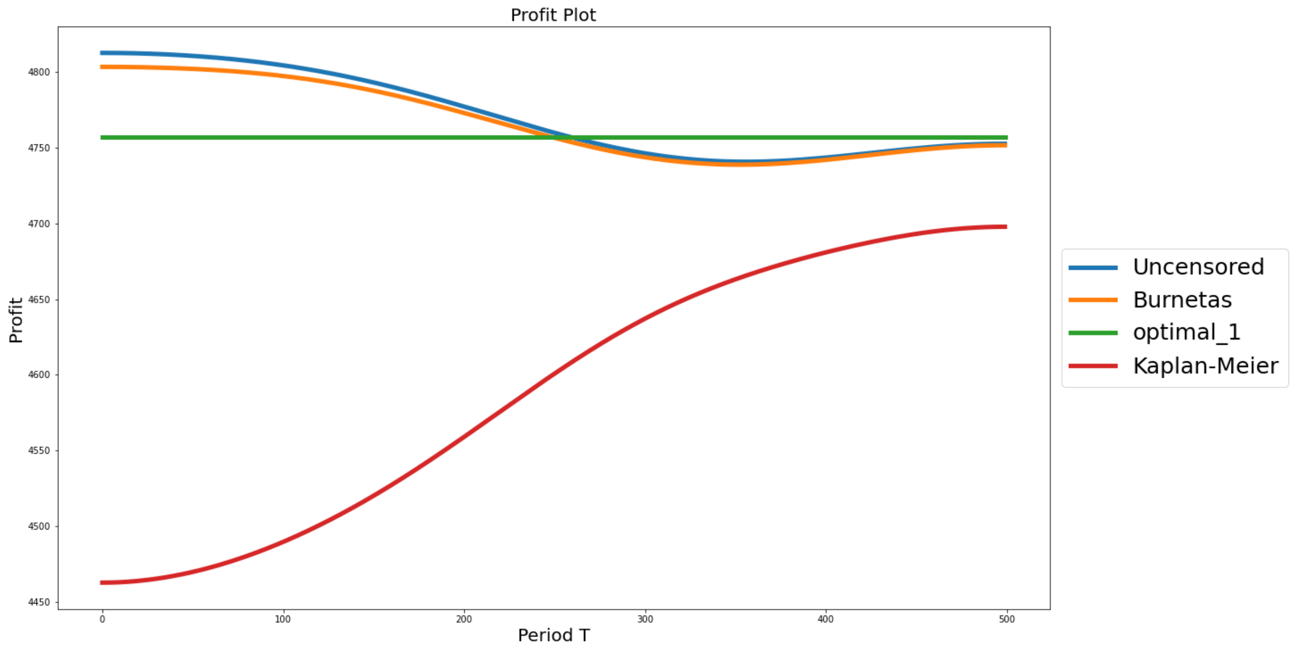

Figure 6.14: Profit Plot. Mean value is 200, standard deviation is 100
Underage cost is 9.5 and Overage cost is 31. Burnetas and Uncensored policies are converge very early on benchmark

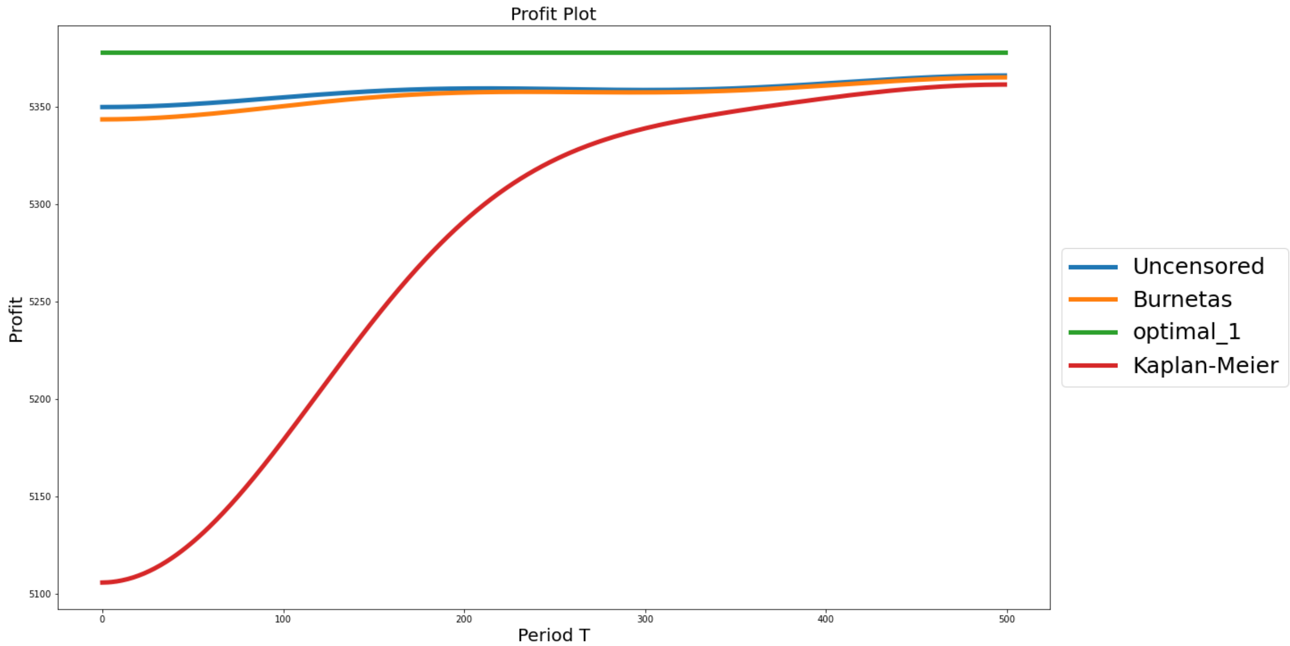

Figure 6.15: Profit Plot. Mean value is 200, standard deviation is 50
Underage cost is 9.5 and Overage cost is 31. Burnetas and Uncensored policies are converge very early on benchmark

https://doi.org/10.26219/heal.aueb.4847

To begin with, when comparing the benchmark figures, it is clear that there is not much difference in terms of how the policies operate and their convergence rates. On the other hand, we observed that when the standard deviation is larger and closer to the mean value, all policies have better convergence rates. Kaplan-Meier drops from 16% to 10%, and the other two policies drop close to zero.

Regarding the profit functions, the Burnetas-Smith and Uncensored policies tend to deviate from the benchmark as the standard deviation increases. The opposite behavior is observed for the Kaplan-Meier profit function, which seems to be closer to the benchmark with a higher standard deviation.

https://doi.org/10.26219/heal.aueb.4847

Chapter 7

Conclusion

Our study aimed to build a number of newsvendor policies and evaluate their effectiveness in a normally distributed demand environment. We constructed and tested three models, the uncensored in which we had complete information about the demand and used the policy that comes up from the normal demand newsvendor equations. The burnetas-Smith policy which is an adaptive method used for the censored instance and use only the last period order quantity to generate next's. And finally the Kaplan-Meier model, which estimates an empirical CDF based on censored data and then uses the newsvendor quantile to find the period's optimal quantity.

In our numerical examples we tested how these policies operated compared to the newsvendor benchmark. I general all three methods showed great results and in the most cases converged efficiently with divergence in the number of period that needed to accomplish that. Burnetas-Smith had the best score, as they were closer to the benchmark from the first periods of time. On the other hand Kaplan-Meier had slower convergence in most instances.

We can relate that with the fact that Kaplan-Meier estimator has a strong dependency on historical data to generate the survival function which are limited on initial periods. That leads the estimator to have a huge error on the early estimates and to operate less efficiently than the other policies.

On this point we have to mention that in one instance the KM estimate was better the Burnetas-

Smith, when the underage cost was relatively bigger than the overage a fact that leads to a ratio over 0.5. Because on the expression of Burnetas-Smith the quantile it is a negative value, corrects each period quantity and leading it into its 'true' value much slower. On the other hand quantile in KM's CDF is closer on the larger values and converges faster to the benchmark. That realization can be useful in order to pick policies depending on the ratio R which is always known in advance. Additionally, we also examined how changing the distribution parameters affected the policies' performance. We found that larger standard deviations and mean values tended to improve the policies' convergence rates, but the impact on the profit functions varied. Overall, our study provides insights into the effectiveness of different newsvendor policies in a normally distributed demand environment and highlights the importance of considering factors such as historical data and distribution parameters when selecting a policy.

https://doi.org/10.26219/heal.aueb.4847

Chapter 8

Code

As a great part of this work was the construction of an algorithm that could implemented the policies and generate results, we think it is useful to present that code in this chapter.

Code is splitted into three parts in order to be read and run more efficiently. The main part consists of the system that is capable of running the policies(2nd part) and the profict function (third part). We are starting from the policies, followed by the profit function and ending with the system implementation:

```
# ===========================================================
import numpy as np
from scipy.stats import norm
# ===========================================================
# Calculate the order quantity of a single period
# with newsvendor uncensored policy for known normal demand
# ===========================================================
def newsvendor(demand,ratio,period):
    estimated_mu = np.mean(demand[0:period+1])
    estimated_sigma = np.std(demand[0:period+1])
    result = estimated_mu + estimated_sigma*norm.ppf(ratio)
    return result
# ===========================================================
```

https://doi.org/10.26219/heal.aueb.4847


```
# Calculate the order quantity of a single period
# with Burnetas-Smith policy for censored unknown demand distribution
# ==========================================================
def Burnetas_Smith(demand,quantity,ratio,period):
    if demand[period] <= quantity:
        Y=1
    else:
        Y=0
    result = quantity*(1-(Y-ratio)/(period+1))
    return result
# ==========================================================
# Calculate the order quantity of a single period
# with Kaplan-Meier policy for censored unknown demand distribution
# ==========================================================
def Kaplan_Meier(demand,quantity,Z_t,period,CCDF,ratio,count_km,Co,Cu,
   count_censored,q_maintain,Q_nv):
    def insert_into_array(arr, new_element):
        num_rows, num_cols = arr.shape
        i = 0
        while i < num_cols and arr[0, i] < new_element[0, 0]:
            i += 1
        right_side = arr[:, i:num_cols]
        arr = np.delete(arr, np.s_[i:num_cols], axis=1)
        arr = np.hstack((arr, new_element, right_side))
        return arr, i

    def find_closest_index(arr, value):
        diff = np.abs(arr - value)
        return np.argmin(diff)

    if quantity<=demand[period]:
```

https://doi.org/10.26219/heal.aueb.4847


```
            delta = 0
        else:
            delta = 1

        col = np.array([[min(quantity,demand[period])],[delta]])
        Z_t , x = insert_into_array(Z_t, col)
        if delta==1:
            for i in range(x,Z_t.shape[0]):
                if Z_t[0,i]==Z_t[1,i] and Z_t[1,i]==0:
                    Z_t[:,x],Z_t[:,i]=Z_t[:,i],Z_t[:,x]
                    x = i
                    break

        for l in range(period):
            F=1
            for r in range (0,l):
                F = F*((period-r)/(period-r+1))**(Z_t[1,r])
            col_cdf = np.array([[F],[Z_t[0,l]]])
            CCDF[0,l]= F
            CCDF[1,l]= Z_t[0,l]

        F=1
        for r in range (0,period):
                F = F*((period-r)/(period-r+1))**(Z_t[1,r])
        col_cdf = np.array([[F],[Z_t[0,period]]])
        CCDF = np.append(CCDF,col_cdf,axis = 1)


        CCDF = CCDF[ :, CCDF[0].argsort()]

        position = find_closest_index(CCDF[0], 1-ratio)
```

https://doi.org/10.26219/heal.aueb.4847


```
    quantity = CCDF[1,position]


    if period <= 20:
        quantity = Q_nv
    else:

        if Z_t[0,period]==quantity:
            count_km += 1
            if delta==0:
                count_censored+=1
        if count_km!=0 and count_censored/count_km - Co/(2*(Cu+Co)) > 0:
            Z_t[0,period] = Z_t[0,period]*2
            quantity = Z_t[0,period]
            count_censored = 0
            count_km = 0
# ==================================================================
    return quantity,Z_t,count_censored,count_km,CCDF,q_maintain
# ==================================================================
```

8.0.1 Profits

```
import numpy as np
from scipy.stats import norm
# ====================================================
# Calculate the observed profit of a single period
# ====================================================

def profit(u,demand,quantity,p,g,B,period):
    if demand[period] <= quantity:
```

https://doi.org/10.26219/heal.aueb.4847


```
            result = -quantity*u+p*demand[period]+g*(quantity-demand[period])
        else:
            result = -quantity*u+p*quantity-B*(demand[period]-quantity)
        return result
# ==================================================
# Calculate the expected profit of the normal demand
# ==================================================
def expected_profit(p,u,g,B,mu,sigma,ratio,quantity):
    result = (p-u)*mu - (p-g+B)*(norm.pdf(norm.ppf((p-u+B)/(p-g+B))))*sigma
    return result
# ==================================================
def expected_profit_2(p,u,mu,sigma,quantity,g):
     expected_profit = (p-g)*mu*norm(mu,sigma).cdf(quantity)-\
                       (p-g)*sigma*norm(mu,sigma).pdf(quantity)-\
                       (u-g)*quantity*norm(mu,sigma).cdf(quantity)+\
                       (p-u)*quantity*(1-norm(mu,sigma).pdf(quantity))
     return expected_profit
```

8.0.2 Main System

```
import numpy as np
from scipy.stats import norm
from matplotlib import pyplot as plt
from scipy.ndimage.filters import gaussian_filter1d
import policies as pol
import profits as prof
import os
#initialize underage cost & overage cost
# u: unit purchase cost
# p: sold unit revenue
```

https://doi.org/10.26219/heal.aueb.4847


```
# g: salvage value of disposed units
# B: Lose of goodwill cost
# Cu: underage cost (Cu = p-u+B)
# Co: overage cost (Co = u-g)
# R: critical ratio
#initialize costs
p  = 40
u  = 30
g  = 0.5
B  = 1
Cu = p-u+B
Co = u-g
R  = Cu/(Cu+Co)
# ======================================================================
# mean and standard deviation, number of periods T
# initialize starting quantity whick going to be same in every iteration
mu    =  200
sigma = 50
T     = 500
# ======================================================================
# # The newsvendor policy is given from
# ======================================================================
Q_nv = mu + sigma * norm.ppf(R)
# ======================================================================
# Calculate profit function
# ======================================================================
# to run the simulation we repeat the process for N times
# in every iteration we generate a number of T samples
```

https://doi.org/10.26219/heal.aueb.4847


```
# from a random distribution
# A is a 3 dimentional matrix. Every level contains T values of order quantities
# for every level the first column is for the uncensored quantity
# second column contains the profit of uncencored policy
# the third is the quantity for the burnetas smith
# fourth is profit with the burnetas smith policy
# second and fourth is for the equivalent costs
N = 50
A = np.zeros((N,T,6))
for i in range(0,N):
    print(i)
# ==================================================================
#     # initialize the starting quantities
# ==================================================================
    Q_i   = Q_nv
    Q_bs  = Q_nv
    Q_km  = Q_nv
# ==================================================================
#     # generate all the period demands
# =================================================================
    demand = np.random.normal(mu, sigma, T)
    #Intialize estimated mean value and sandar deviation
    est_mu    = np.zeros(T)
    est_sigma = np.zeros(T)
# =================================================================
#     # Initials needed for the Kaplan Meier policy
#     #construct the list for the KM estimator
#     #First row contains the observed quantity
#     #Second row contains the indicator (censored or not)
#     #Third row contains the demand of current period
```

https://doi.org/10.26219/heal.aueb.4847


```
# ===============================================================
    Z_t = np.zeros((2,1))
    CCDF  = np.zeros((2,1))
    count_km = 0
    count_km_c = 0
    count_censored=0
    q_maintain = 0

    if Q_km < demand[0]:
        delta = 0
    else:
        delta = 1
        Z_t[0]=min(Q_km,demand[0])
        Z_t[1]=delta
        x=0
        CCDF[0] = 0.5**delta
        CCDF[1] = min(Q_km,demand[0])



# ==================================================================



    # start with the calculation of the policies
    # We starting with the Calculation of Burnetas-Smith policy and
    # mean & standard deviation for the Uncensored problem

    for k in range(0,T):

# ==================================================================
```

https://doi.org/10.26219/heal.aueb.4847


```
#           #1.Known demand uncensored newsvendor policy
# ==================================================================

        #Order quantity
        A[i,k,0] = pol.newsvendor(demand,R,k)
        #Period profit
        A[i,k,1] = prof.profit(u,demand,Q_i,p,g,B,k)
        #the next iteration order quantity
        Q_i = A[i,k,0]

# ==================================================================
#           #2.Burnetas-Smith policy
# ==================================================================
        A[i,k,2] = pol.Burnetas_Smith(demand,Q_bs,R,k)
        #Period profit
        A[i,k,3] = prof.profit(u,demand,Q_bs, p, g, B,k)
        #the next iteration order quantity
        Q_bs = A[i,k,2]
        #3. Kaplan-Meier Policy
        # we have already set up the first order quantity at the begining
        if k > 0:
            #Order quantity
            A[i,k,4],Z_t,cencor_count,count_km,CCDF,q_maintain = pol.Kaplan_Meier(demand, Q_km, Z_t, k, CCDF, R, count_km, Co, Cu, count_censored,q_maintain,Q_nv)

            #Period profit
            A[i,k,5] = prof.profit(u,demand, Q_km, p, g, B,k)

            #the next iteration order quantity
```

https://doi.org/10.26219/heal.aueb.4847


```
            Q_km   = A[i,k,4]
# ====================================================================
#        #expected profit
# ====================================================================
expected_profit   = prof.expected_profit(p, u, g, B, mu, sigma, R,
   Q_nv)
expected_profit_2 = prof.expected_profit_2(p, u, mu, sigma, Q_nv, g)
# ====================================================================
# #Smoothing
# ====================================================================
# Taking the mean value array of A
estimate_values = sum(A[j,:,:] for j in range(0,N))/N

# Create a smoothed version of the quantities
Burnetas_Smoothed = gaussian_filter1d(estimate_values[:,2], sigma=100)
Uncencored_Smoothed = gaussian_filter1d(estimate_values[:,0], sigma
   =100)
KM_Smoothed = gaussian_filter1d(estimate_values[:,4], sigma=100)

# Create a smoothed version of the profits
Burnetas_profit = gaussian_filter1d(estimate_values[:,3], sigma=100)
Uncencored_profit = gaussian_filter1d(estimate_values[:,1], sigma=100)
KM_profit = gaussian_filter1d(estimate_values[:,5], sigma=100)
```

https://doi.org/10.26219/heal.aueb.4847


```

# Create a smoothed version of the percentage differnce
Burnetas_Smoothed_perc = gaussian_filter1d(100*abs((estimate_values[:,2]-Q_nv)/Q_nv), sigma=100)
Uncencored_Smoothed_perc = gaussian_filter1d(100*abs((estimate_values[:,0]-Q_nv)/Q_nv), sigma=100)
KM_Smoothed_perc = gaussian_filter1d(100*abs((estimate_values[:,4]-Q_nv)/Q_nv), sigma=100)
# ======================================================================
base_path =  "C:\\Users\\user\\Documents\\      \\Diploma\\Plots for final\\"
file_name_quantity = 'quantity.png'
file_name_profit = 'profit.png'
file_name_benchmark = 'benchmark.png'
path = os.path.join(base_path,f"{mu}_{sigma}_{Co}_{Cu}")
if not os.path.exists(path):
    os.makedirs(path)
# ======================================================================

file_path_quantity = os.path.join(path, file_name_quantity)
file_path_profit = os.path.join(path, file_name_profit)
file_path_benchmark = os.path.join(path, file_name_benchmark)
# ======================================================================

# Create figure to compare the policies
x =np.ones(T)*Q_nv
x1 = range(0,T,1)
fig, ax = plt.subplots(figsize=(20, 10), layout='constrained')
ax.plot(Uncencored_Smoothed ,label='Uncensored',linewidth = 5)
```

https://doi.org/10.26219/heal.aueb.4847


```
ax.plot(Burnetas_Smoothed,label='Burnetas',linewidth = 5)
ax.plot(x1,x,label='optimal',linewidth = 5)
ax.plot(KM_Smoothed,label='Kaplan-Meier',linewidth = 5)
ax.set_xlabel('Period T',fontsize = 20)
ax.set_ylabel('Quantity Q',fontsize = 18)
ax.set_title("Quantity Plot",fontsize = 18)
ax.legend(loc='center left', bbox_to_anchor=(1, 0.5), prop={'size': 25})
plt.savefig(file_path_quantity)
# =============================================================================

# Create figure to compare the policies
x =np.ones(T)*expected_profit
x_2 = np.ones(T)*expected_profit_2
x1 = range(0,T,1)
fig2, ax2 = plt.subplots(figsize=(20, 10), layout='constrained')
ax2.plot(Uncencored_profit,label='Uncensored',linewidth = 5)
ax2.plot(Burnetas_profit,label='Burnetas',linewidth = 5)
ax2.plot(x1,x,label='optimal_1',linewidth = 5)
ax2.plot(KM_profit,label='Kaplan-Meier',linewidth = 5)
ax2.set_xlabel('Period T',fontsize = 20)
ax2.set_ylabel('Profit',fontsize = 20)
ax2.set_title("Profit Plot",fontsize = 20)
ax2.legend(loc='center left', bbox_to_anchor=(1, 0.5), prop={'size': 25})
plt.savefig(file_path_profit)
# =============================================================================

# Create figure to compare the policies
```

```
fig, ax = plt.subplots(figsize=(20, 10), layout='constrained')
ax.plot(Uncencored_Smoothed_perc ,label='Uncensored',linewidth = 5)
ax.plot(Burnetas_Smoothed_perc,label='Burnetas',linewidth = 5)
ax.plot(KM_Smoothed_perc,label='Kaplan-Meier',linewidth = 5)
ax.set_xlabel('Period T',fontsize = 20)
ax.set_ylabel('Percentage (%)',fontsize = 20)
ax.set_title("Difference from benchmark",fontsize = 20)
ax.legend(loc='center left', bbox_to_anchor=(1, 0.5), prop={'size': 25})
plt.savefig(file_path_benchmark)
```

https://doi.org/10.26219/heal.aueb.4847

Bibliography

[1] Klein, John P. and van Houwelingen, Hans C. and Ibrahim, Joseph G. and Scheike, Thomas H. (Eds.). *Handbook of Survival Analysis*. CRC Press, 2013.

[2] Li, Y., and Zheng, Y. “Estimating the optimal order quantity and the maximum expected profit for single-period inventory decisions.” *Applied Mathematical Modelling*, vol. 37, no. 24, pp. 10185-10195, 2013.

[3] Silver, E. A., Pyke, D. F., and Peterson, R. *Inventory Management and Production Planning and Scheduling*. Wiley, 1998.

[4] Huh, W. T., and Janakiraman, G. “The classical newsvendor model under normal demand with large coefficients of variation.” *European Journal of Operational Research*, vol. 216, no. 2, pp. 300-306, 2012.

[5] Burnetas, A. N., and Smith, C. E. “Adaptive ordering and pricing for perishable products.” *Management Science*, vol. 50, no. 5, pp. 710-722, 2004.

[6] Chen, X., and Simchi-Levi, D. “Adaptive data-driven inventory control policies based on Kaplan-Meier estimator.” *Operations Research*, vol. 62, no. 5, pp. 1071-1086, 2014.

[7] Nahmias, S. *Production and Operations Analysis*. McGraw-Hill Education, 2015.