POLITECNICO DI TORINO

Department of Management and Production Engineering

DIGEP

M.Sc in Management & Engineering

Double Degree with ESCP Business School

Master's Thesis

Impact of different sources of uncertainty on the Newsvendor Model

Assess model robustness under different economic parameters and demand curves

Relator:
Prof.ssa Alfieri Arianna

Co-Relators:
Prof.ssa Pastore Erica
Prof.ssa Succi Chiara

Student:
Debernardi Luca

Academic Year 2024

Master's Thesis written in LaTeX
to assess the robustness of
the Newsvendor Model,
Luca Debernardi
S288881
.

Impact of different sources of uncertainty on the Newsvendor Model

Luca Debernardi - s288881

2024

Contents

List of Figures

Part I

Thesis' Overview

Chapter 1

Introduction

1.1 Why this topic?

The Master Thesis is the apex of any academic journey: it should encompass the totality of expertise and knowledge acquired during such a long and unique period.
My path started in 2017 with a Bachelor in Management Engineering here, at the Polytechnic University of Turin, because I wanted to focus my studies on two main topics: Business and Computer Science.
At the end of my B.Sc, completed in September 2020 with a final mark of 108/110, I wanted to strengthen even more this dualism of skills. To do so, as well as continue the standard course of study, I decided to enroll in a double degree program with the french ESCP Business School, the oldest and among the most prestigious management universities in the world.
After a few months, I also had the opportunity to start working full time in the IT department of Amplifon S.p.A., an Italian multinational company, where for the following 3 years I covered multiple roles both in the Global Strategy & Governance and in the Retail Excellence Big Data & BI teams. Now, while writing this document, I work as a Global Strategic Pricing Analyst at Thermo Fisher Scientific.
During my Master, I tried to exploit my academic knowledge in the professional world, by choosing roles that matched my skill-set. For my final thesis, I desired to solve a real business problem using the acumen and skills gained over these years and assess the Newsvendor robustness through simulation was the perfect challenge.

1.2 Brief introduction to the Newsvendor Model

The Newsvendor Model, rooted in operations research and inventory management theory, offers a systematic framework for determining optimal inventory levels in the face of uncertain demand. Originally derived from the problem faced by a newspaper vendor deciding how many copies of a newspaper to order for sale the next day, this model has found widespread application across various industries, ranging from retail and manufacturing to service sectors [DeMarle(2019)]. If the vendor orders too few, they risk running out of stock and missing out on potential sales. On the other hand, if they order too many, they incur unnecessary costs due to excess inventory that may not sell before becoming obsolete.

The logical basis is pretty simple: by balancing the costs of over-stocking and under-stocking, the model provides decision-makers with insights into how to optimize their inventory policies to maximize expected profits according to a predetermined demand distribution (in the literature mainly normal, exponential or uniform). Moreover, its adaptability to different scenarios, such as perishable goods, seasonal demand, and varying production lead times, further enhances its relevance and applicability in real-world settings.
The model is based on three economic parameters:

- **Unit Price:** (up) It is the price at which is possible to sell each product by satisfying the demand in standard conditions. It is earned only when the transaction is performed.
- **Unit Cost:** (uc) It is the cost associated to the purchase of every and each item, regardless if it will be sold or not.
- **Resale Value:** (sv) This is the residual value of the item once all the demand has been satisfied. It is usually close to 0 or extremely discounted if compared to the original unit price.

The logic behind the Newsvendor Model, according to the theory, is valid for any kind of distribution and dependent on the economic parameters underneath. In the simplest version, the model says that the optimal Q is the quantity that corresponds to a cumulative probability of the demand distribution equal to:

$$\frac{up - uc}{up}$$

This ratio is completely independent from the distribution, whose impact occurs only on the resulting quantity.

1.3 Research question and main objectives

The research question and main goal of this thesis can be easily guessed from the title:

How does the optimal quantity and resulting expected profit change in function of the economic parameters and the demand distribution type and moments?

To properly answer this question, the problem should be broken in three smaller milestones, to be tackled using a waterfall approach:

- The first objective is to confirm empirically that the theoretical model is valid under any circumstances, verifying it for different types of distribution and economic parameters.
- The second objective is to demonstrate what the model is already implicating: the Newsvendor is only dependent on the demand distribution up to the cumulative value of probability obtained from the economic parameters.
  In the simplest case, if we have a unit price of \$10 and a unit cost of \$5, the resulting cumulative probability is:

  $$P = \frac{up - uc}{up} = \frac{10 - 5}{10} = 50\%$$

  If the distribution is symmetrical, the resulting quantity would be the median (and therefore mean) value, regardless of the amount of standard deviation and kurtosis of the actual curve.
- The third objective is to assess the impact that a change in the demand distribution can have on the expected profit, even in the scenario where it does not impact directly the optimal quantity. This last point is extremely critical in the decision making process, providing insights on the level of risk associated to the chosen economic parameters and the expected demand curve peculiar to the model.

Chapter 2

Literature Review

2.1 (Lack of) existing literature on Newsvendor robustness

Even though the Newsvendor is a well known model in the optimization of supply chain management, supported by countless publications and studies (explicating particular scenarios or more comprehensive versions, as well as documents on particularly significant case studies), the literature addressing the sensitivity of the model is notably scarce.
The articles found can be categorized in three main groups:

- General articles summarizing the state of the research regarding the Newsvendor. Such articles very rarely cover the concept of sensibility due to the lack of general source studies on the topic;
- Highly specialized articles addressing narrowly focused research questions. These articles are fragmented, rely on complex mathematics, and provide specific demonstrations that are valid only under strict assumptions;
- Empirical studies that offer high level guidance and insights on Newsvendor variability. Their actionability is compensated by a lack of numerical examples and precise impact assessment.

This chapter discusses a comprehensive review of the existing literature, examining the limited studies available on the topic, their methodologies, findings and alternative approaches.

[Khanra et al.(2014)Khanra, Soman, and Bandyopadhyay] is probably the closest article published regarding the sensitivity in the Newsvendor model existing in the literature. It is an attempt of addressing the implications of parameter estimation errors and sub-optimal decisions. It focuses on the identification of conditions for symmetry and skewness of cost deviation, in order to determine whether ordering more or less than the optimum is advisable. According to the article, the mean of the demand distribution is the most important parameter in the definition of the optimal order quantity. However, skewness impacts need to be studied further to understand the impact over Newsvendor performance.

Other articles approach the robustness and the sensitivity problem from a different viewpoint: [Krishnendu Adhikary and Kar(2018)] for example extends the Newsvendor theory to a distribution-free scenario, setting-up the demand as a fuzzy-random variable. Interestingly, using this methodology, both mean and standard deviation of the demand are considered known, whereas the related probability distribution function is not. All the distribution specific impacts are neglected in this approach, that therefore needs to be completed by a complementary analysis that focuses on the significance of the distribution type, keeping equal mean and deviation.

[Borgonovo(2010)], even though not directly related to the Newsvendor model, portrays an important example of methodology. In fact, its purpose is to accurately assess the impact that discrete changes in the inputs parameters could have on the output of a model. This result can be obtained by leveraging integral function decomposition and sensitivity measures, enabling for the identification of the contribution of each individual parameter, as well as their group effects and interactions. Finally, the methodology facilitates communication of the sensitivity analysis results to decision-makers, through structured settings that highlight key drivers and provide insights into the model structure.

Instead, [Qin et al.(2011)Qin, Wang, Vakharia, Chen, and Seref] offers an overview of the latest researches on the Newsvendor Model. Even though providing many useful insights, it is another proof of the almost total lack of research on the topic of sensitivity and robustness for such a widespread model. In fact, it mainly focuses on three areas:

- the first one is customer demand and its interrelation with price and marketing effort;
- different supplier pricing policies, especially discount schemes over certain thresholds;

- buyer risk profile, that usually can be classified as risk averse and/or "risk" takers, instead of purely risk neutral.

[Hedayatinia et al.(2020)Hedayatinia, Lemoine, Massonnet, and Viviani], focus their analysis on a Newsvendor problem where the retailer determines both the selling price and the order quantity, considering stochastic and price-dependent demand. Additionally, the retailer can sell unsold units at the end of the sales season.
It also introduces an analytical model for optimizing retailer's decisions and identifies conditions for simultaneously finding optimal quantity and price through numerical methods.

[Andrew Butters(2019)], provides three comparative statistics regarding the level of demand uncertainty, specifically for the Newsvendor Model. Even in this article it is shown that two distributions, having the same mean but different standard deviations, result in a situation where both the expected profit and quantity sold drop in the most disperse one.
This result is key in the understanding the general behavior of the Newsvendor and its robustness. In fact, in a symmetrical distribution, a change in the standard deviation would both increase the probability of having a lower actual demand to occur, but also increase by the same amount the probability of a higher one. Even though the net effect on the distribution mean is absent, the impact of a higher variability on the Newsvendor must be kept into account, being its negative effect in the left tail of the distribution not compensated by the same change in the right one.
Unfortunately this article does not provide real examples or an estimate of the impacts of such change in variability.

[Jammernegg et al.(2022)Jammernegg, Kischka, and Silbermayr] is an interesting analysis focused on the behavioral Newsvendor. Even though it is still tackling the problem of the Newsvendor robustness, the level of specification is so significant that its contribution to a generalization of Newsvendor sensitivity is close to be negligible.
The same can be said for [Liu et al.(2022)Liu, Letchford, and Svetunkov], that approaches Newsvendor problems of different mathematical complexities.
Both of them offer valid insights, but they are also example of the extreme specification of the vast majority of the articles available on the topic.

What is mainly missing in the existing literature is an overview on the sensitivity and robustness of the Newsvendor model. Starting from the dif-

ferent insights gathered from these articles, we will analyze the impact that the variation of multiple parameters has on the model, ranking the by their importance and understanding their mutual effects.
This thesis aims to fill the gaps in the literature by providing actionable insights to the decision makers regarding the risk sensitivity of the Newsvendor, by calculating the numerical impacts of changes in mean and standard deviation, studying the effects that a skewed distribution could have on the optimal values and the importance of the residual value, parameter often neglected in the existing research on variability.

Chapter 3

Design of Experiment

This chapter explains the methodology and implementation of this thesis. It outlines the underlying logic and provides the foundation necessary to understand the insights and conclusions presented throughout the dissertation.

3.1 Why simulation?

Assessing the robustness and effectiveness of the Newsvendor Model requires a methodological approach able to captures its dynamic and stochastic nature. Simulation offers the flexibility to model all the intricate interactions typical of real world problems, while accounting for diverse sources of uncertainty, providing a more realistic representation.
Unlike static analytical models, simulation replicates dynamic scenarios where demand and supply conditions vary. It allows for validating model assumptions, testing the robustness of results, and conducting sensitivity analysis to assess the impact of parameter uncertainties on decision outcomes.
The possibility to immediately run what-if analysis is an additional tool to test the impact that a decision taken ex-ante has once the real conditions get defined, with the objective to measure model accuracy and reliability.
Simulation biggest strength is its capacity to capture the interplay of different uncertainties, may be the distribution type and moments or economic parameters, making it the best way to test real-world cases too complex to be analyzed using a pure analytical approach.

3.2 Simulation structure

The simulated approach has to run multiple Newsvendors in order to compare their optimal values and profit curves. Each simulation has potentially

different economic parameters and underlying demand distributions, both in terms of type and moments. Such distribution is completely independent from the economic parameters chosen for the Newsvendor.
The overall structure of the simulation can be seen in figure 3.1.

The Model class encompasses multiple Newsvendors, which are evaluated collectively and compared to derive conclusions. Each Newsvendor is characterized by its economic parameters and a specific distribution.
Further details on the Python implementation of the simulation can be found in appendix A.

3.2.1 Implemented distributions

In order to properly compare different distributions, they must be constructed in order to have precise values of mean, standard deviation, skewness and kurtosis. The distributions implemented are the following:

- **Uniform:** Constant distribution defined in all its points by the same probability. It is the easiest distribution to study and provides the most neutral results.
- **Normal:** Most common distribution, used as a benchmark to test the behavior of all the others.
- **Exponential:** Extremely skewed distribution, used to study the behavior on the Newsvendor for non-symmetrical distributions.
- **Gamma:** Dynamic distribution defined by the shape parameter α and the scale parameter β. Gamma distribution will be used to study in chapter 9 the impact of changes in skewness (γ): ranging from $\gamma = 2$, where it equals the exponential, to $\gamma = 0$, where it closely approximates the normal distribution.
- **Beta:** Generalization of the gamma function, it needs 2 parameters α and β as well. Even though it is considered a whole family of continuous probability distributions due to its extreme versatility, it will be used in this thesis only in its symmetrical instances, in order to study in chapter 10 the impact of the kurtosis on Newsvendor robustness. In fact, it can cover values of kurtosis (κ) between $\kappa = 0$, where it equals the normal distribution, and $\gamma = -1.2$, where it converge to the uniform one.

3.2.2 Definition of optimal Newsvendor results

The most important values to be found in any Newsvendor are the optimal quantity and the related optimal expected profit. Such optimal values are computed in two different ways, in order to compare the result obtainable from the theory to the experimental one derived from the simulation.

The **analytical approach** is rooted in the Newsvendor theory, that will be analyzed more in detail in the section 5.4. It is based on the well known formula derived from Littlewood's rule:

$$q_1' = D^{-1}\left(\frac{up - uc}{up - sv}\right)$$

where D is the cumulative distribution of the demand, up the unit price, uc the unit cost and sv the residual value of the item once all the demand has been satisfied. Consequently, to compute the optimal quantity, it is enough to know the 3 economic parameters and the inverse of the probability function. When performed analytically, this formula provides the optimal theoretical result of the Newsvendor, unfortunately without any further indication on its sensitivity.

The **iterative approach**, based on simulation, is the empirical value obtained by calculating, for each possible purchasable quantity in the domain, the expected value of the profit $E[Pr(q_p)]$ and select the maximum.
The expected value of the profit, given the quantity purchased q_p, is the average of all the profits associated to each possible value of the demand d, weighed by the probability that each of them has to be the realized actual demand.

$$E[Pr(q_p)] = \sum_{d=0}^{\infty}\Big(Pr(q_p, d) \cdot P[D = d]\Big)$$

The exact profit associated to a combination of purchased quantity and realized demand is

$$Pr(q_p, d) = min(q_p, d) \cdot up - q_p \cdot uc + max(q_p - d, 0) \cdot sv$$

where the first addendum is the profit associated to the units sold, the second is the cost associated to the purchased quantity and the third is the residual value of the overstocked items.

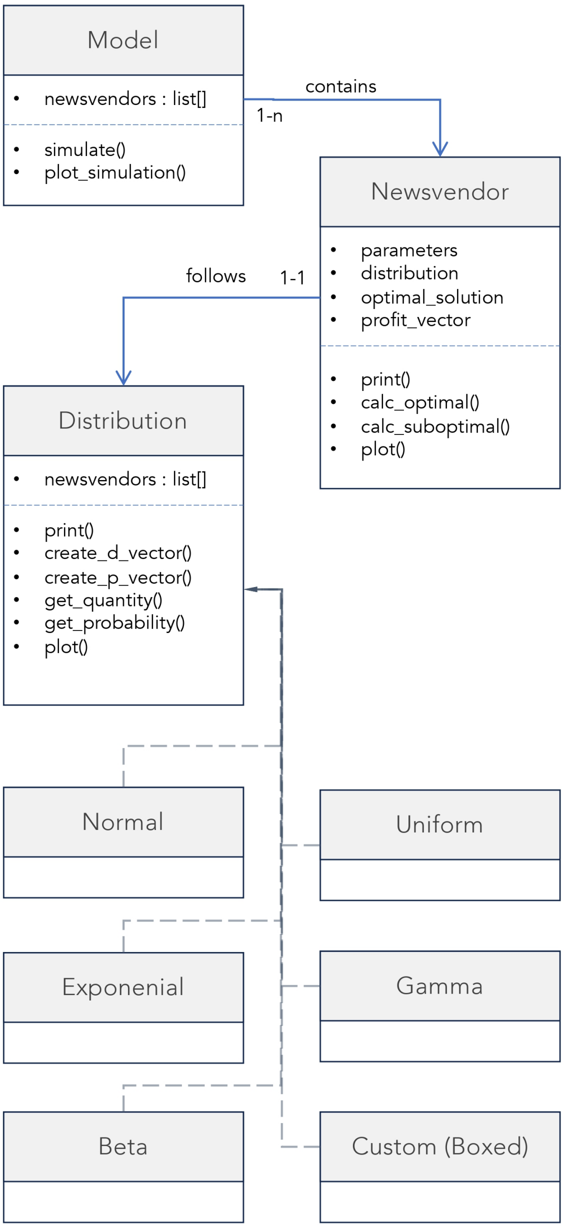

Figure 3.1: Simplified UML diagram portraying the relationship between all the different classes created for the project and their main attributes and methods

Part II

Alignment of Empirical & Analytical Models

Chapter 4

Validation of the empirical model

4.1 Some simple scenarios

As a first step, before drawing any conclusions from the simulation, it is mandatory to ensure that its outputs really match the analytical results obtainable from the theoretical model by starting from some known cases.
In this test, the economic parameters chosen are not relevant. For the reader's convenience, the following values are set throughout the chapter:

- Unit Price (up) = €8.00;
- Unit Cost (uc) = €5.00;
- Resale Value (sv) = €1.00.

In the case of a residual value different from 0, the analytical formula for the optimal quantity (that will be explained in details inside section 5.4) can be formulated as:

$$q_1' = D_1^{-1}\left(\frac{up - uc}{up - sv}\right)$$

Because D_1^{-1} is the inverse of the cumulative distribution, it is highly dependent on the curve chosen.
What can instead be defined, regardless of the distribution, is the value of the critical fractile, the ratio inside the brackets, defined as:

$$P' = \frac{up - uc}{up - sv}$$

Throughout the thesis, the percentage derived from this ratio will also be referred to as the *optimal cumulative probability*. This value indicates the

point of the demand distribution where the optimal quantity occurs. While the optimal quantity is dependent on the distribution, the critical fractile remains constant regardless of the curve chosen. This constancy facilitates comparability between Newsvendors with different demand distributions and enables to generalize the findings obtained to different possible curves.
Fixing the economic parameters, the value of the optimal cumulative probability is

$$P' = \frac{up - uc}{up - sv} = \frac{8-5}{8-1} = \frac{3}{7} = 42.86\%$$

This means that, whatever the distribution chosen, the optimal value will always be the quantity at the 42.86^{th} percentile of the curve.

The most common distributions chosen to perform these checks are uniform, exponential and normal. They will be analyzed in the following sections of the chapter, in order to ensure the validity of the simulated results in multiple scenarios.

4.2 Uniform distribution

The uniform distribution is defined by a constant probability, that remains the same throughout the full domain of the curve. It is therefore **symmetrical**, having its mean and median coinciding in the middle of the curve.
The a and b parameters chosen as a numeric example are:

- **Minimum quantity** (a) equal to 0 units;
- **Maximum quantity** (b) equal to 2,000 units.

Having defined a and b, the mean of the distribution can be easily computed as

$$\frac{b-a}{2} = \frac{2,000-0}{2} = 1,000 \text{ units}$$

4.2.1 Analytical solution for the optimal quantity following a uniform distribution

At first sight, having any point of the curve the same probability, it could be reasonable to believe that the middle point of the distribution, 1000 units, is

the less risky estimate of the purchasable quantity.
However, according to the model, the optimal quantity should be:

$$q_1' = D_1^{-1}\left(\frac{up - uc}{up - sv}\right) = D_1^{-1}(42.86)$$

In the case of a uniform distribution, the probability is equally spread throughout the curve. The cumulative probability in its mean, 1,000 units in this specific case, is therefore 50%, computed as

$$P_{50\%} = \frac{q'_{50\%} - a}{b - a} = \frac{1,000 - 0}{2000 - 0} = \frac{1,000}{2,000} = 50\%$$

In the case of a cumulative distribution in 42.86%, the relationship becomes

$$P' = \frac{q_1' - a}{b - a} = \frac{q_1' - 0}{2,000 - 0} = \frac{q_1'}{2,000} = 42.86\%$$

that can be very easily rewritten as

$$q_1' = 2,000 \cdot 42.86\% = 857 \text{ units}$$

Compared to what initially believed as the "safest" estimate for the quantity, we are in the situation of a 14.2% drop in the suggested quantity.

Having obtained the optimal quantity, two immediate follow-up questions can be formulated:

1. *What is the expected profit if I decide to purchase 857 units?*
2. *How much would I have lost on average if I decided to purchase 1,000 units?*

4.2.2 Analytical solution for the associated expected profit

To compute the profit associated to the purchased quantity, the calculation is more complex. Considering that the purchased quantity must be an integer, the distribution in the following calculation is considered as discrete and the equation can be written as:

$$E[Pr(q_1')] = \sum_{d=a}^{b} \Big(Pr(q_1', d) \cdot P[D = d] \Big)$$

because the probability is a constant value, not depending on q_1', it is possible to rewrite the expression above as

$$E[Pr(q_1' = 857)] = \sum_{d=0}^{2{,}000} \Big(Pr(q_1' = 857, d) \cdot prob \Big) = prob \cdot \sum_{d=0}^{2{,}000} \Big(Pr(q_1' = 857, d) \Big)$$

where the constant probability can be calculated as 100% divided by the number of the possible occurrences (that, ranging from 0 to 2,000, corresponds to 2,001 possible quantities)

$$prob = \frac{100\%}{b - a + 1} = \frac{1}{2,001}$$

and the profit associated to a each possible realized demand is

$$Pr(q_1', d) = min(q_1', d) \cdot up - q_1' \cdot uc + max(q_1' - d, 0) \cdot sv$$

Therefore:

$$Pr(d) = min(857, d) \cdot 8€ - 857 \cdot 5€ + max(857 - d, 0) \cdot 1€$$

The complete equation is:

$$E[Pr(q_1' = 857)] = \frac{1}{2,001} \sum_{d=0}^{2{,}000} \Big(min(857, d) \cdot 8€ - 857 \cdot 5€ + max(857 - d, 0)€ \Big)$$

Because the summation of a sum is the sum of the summations, it is possible to compute the three different components one by one:

Calculating the **revenues** part first, we have a summation between 0 and 857

$$\sum_{d=0}^{857} (d) = \frac{857 + 0}{2} \cdot (857 - 0 + 1) = 367,653$$

and a constant 857 for each subsequent demand value up to 2,000

$$\sum_{d=858}^{2{,}000} (857) = (2,000 - 857) \cdot (857) = 979,551$$

The resulting total for the turnover is:

$$367,653 + 979,551 = 1,347,204$$

The **costs** are instead quite simple to calculate, being a constant over the summation

$$\sum_{d=0}^{2{,}000}(857) = (2,000 - 0 + 1) \cdot 857 = 1,714,857$$

Finally, the **residual value** can be calculated similarly to the revenues, being a reversed summation between 0 and 857

$$\sum_{d=0}^{857}(d) = \frac{857 + 0}{2} \cdot (857 - 0 + 1) = 367,653$$

but becoming 0 for any value bigger than the optimal

$$\sum_{d=858}^{2{,}000}(0) = (2,000 - 857) \cdot (0) = 0$$

The resulting total residual value is:

$$367,653 + 0 = 367,653$$

Putting all together, the expected profit at 857 units purchased is:

$$E[Pr(q'_1 = 857)] = \frac{1}{2,001}\left(1,347,204 \cdot 8€ - 1,714,857 \cdot 5€ + 367,653€\right)$$

$$= \frac{1}{2,001}\left(2,571,000€\right) = 1,285€$$

Using the analytical approach for such a simple distribution, it is possible to answer the first question: the expected profit on average for this specific instance of the Newsvendor is 1,285€.

4.2.3 Alignment of simulated findings

Now that both quantity and profit have been calculated following the analytical solution, it is time to verify if the simulated one obtainable in Python is close enough.
The results are displayed in the picture 4.1, composed as follows:

- **Blue curve:** Plot of the probability distribution followed by the demand;
- **Yellow curve:** Plot of the profit distribution resulting;
- **Green vertical line:** Optimal quantity calculated by following the analytical approach;
- **Red vertical line:** Optimal quantity calculated through simulation.

As we can see, both optimal quantity and expected profit exactly match the analytical solution.
In addition, the simulated approach offers a great advantage over the analytical solution: the profit curve displays the expected profit for each possible chosen quantity, enabling to gather an immediate insight over the profit sensitivity for different quantities.

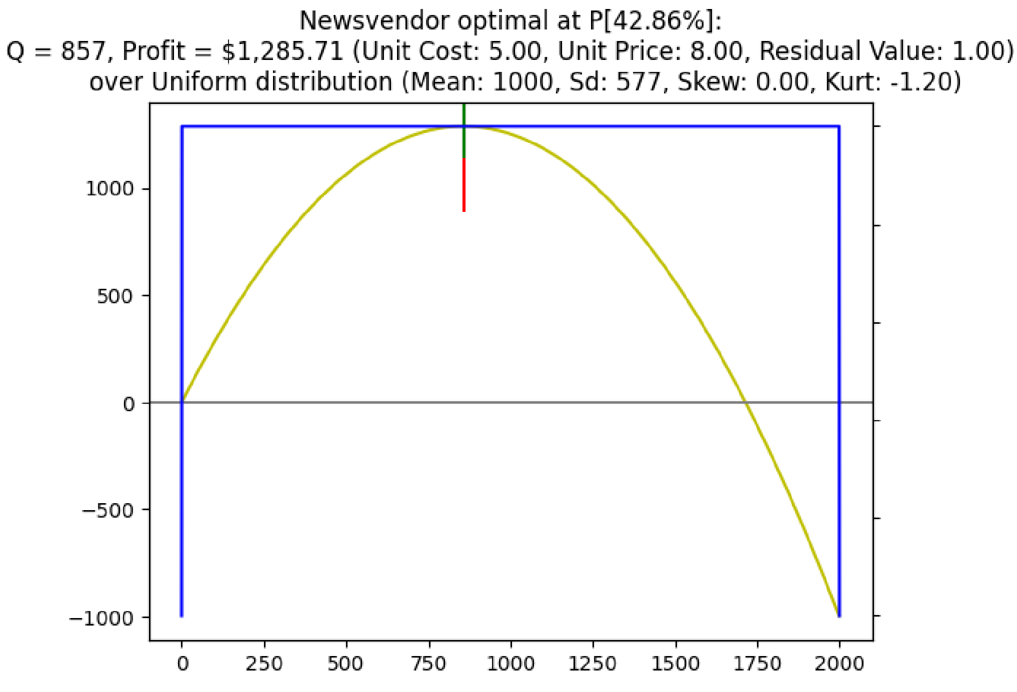

Figure 4.1: Plot of the probability (in blue) and profit (in yellow) in case of uniform demand distribution

By using the simulated approach, it is therefore possible to answer also the second question:

How much would I have lost on average if I decided to purchase 1,000 units?

In figure 4.2 it is explicated what would be the impact:

- **Quantity:** from the optimal value of 875 found previously, the quantity has been fixed to 1,000 units (the purple vertical line in the chart), an increase of 16.7%;
- **Expected Profit:** looking at the profit curve, it is apparent that 1,000 units is still close to the optimal quantity and the decrease in profit is almost negligible.

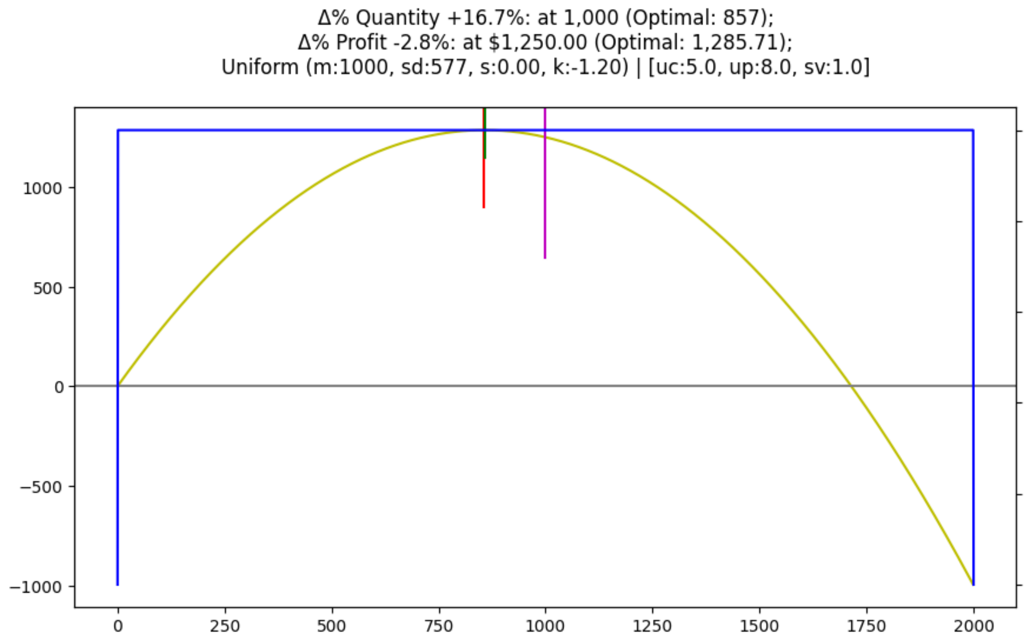

Figure 4.2: Plot that portrays the sensitivity of the expected profit when fixing the quantity to 1,000 units

However, it is important to notice that the profit curve decreases significantly the farther away from the optimal value that, by construction, is only dependent on the economic parameters. In fact, if they are chosen to result in an optimal cumulative probability of 10%, the related optimal quantity of 200 units would be extremely far away from the mean value of the distribution, that would remain equal to 1,000 units.

4.2.4 Changing the economic parameters

In the previous example, the Newsvendor model chosen was a particular scenario in which the optimal quantity was not too distant from the mean of the distribution. But what happens if we decide to set the economic parameters to result in a much lower optimal quantity?
Using

- Unit Price (up) = €6.50;
- Unit Cost (uc) = €5.00;
- Resale Value (sv) = €1.00;

the resulting cumulative probability P' is

$$P' = \frac{up - uc}{up - sv} = \frac{6.5 - 5}{6.5 - 1} = \frac{1.5}{5.5} = 27.27\%$$

and the optimal quantity can be computed as

$$q_1' = 2,000 \cdot 27.27\% = 545 \text{ units}$$

This quantity is just a little more than half of the distribution mean: in figure 4.3 it is shown that the expected profit dramatically drops by almost 70% to 125€, from an optimal value of around 409€. As this case makes obvious, the mean of the distribution says nothing about the optimal quantity that should be purchased, which is instead solely dependent on the chosen economic parameters.

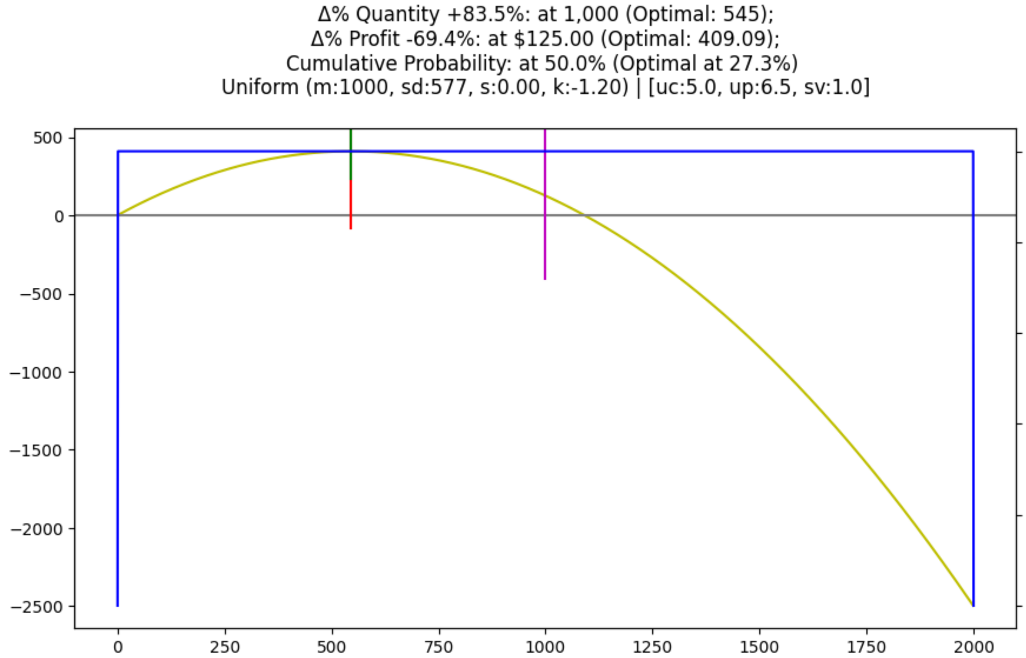

Figure 4.3: The picture portrays that the sensitivity of the profit can vary significantly by changing the economic parameters

4.3 Exponential distribution

The exponential distribution is defined by a decreasing probability. It is therefore **not symmetrical**, having its mean, median and mode at different points of the curve.
The desired measures for this curve, keeping the mean constant with the previous case, are:

- **Mean** equal to 1,000 units;
- **Standard Deviation** equal to 200 units.

4.3.1 Analytical solution following an exponential distribution

The cumulative probability of the baseline case remains equal to the value found for the uniform distribution: 42.86%.

The exponential distribution has a probability density function defined as

$$PDF = \lambda e^{-\lambda x}$$

that, if integrated, results in a cumulative distribution function equal to

$$CDF = 1 - e^{-\lambda x}$$

Mean (μ) and standard deviation (σ) are instead defined as

$$\mu = \frac{1}{\lambda}$$

$$\sigma = \frac{1}{\lambda}$$

By construction, the parameter λ can be calculated in the same way from both mean and standard deviation, that should therefore be identical. If the standard deviation is fixed to 200 units, the mean should be 200 as well. However, it is possible to obtain the desired curve by shifting each point by $1,000 - 200 = 800$ units, that would be the new mode and starting point of the distribution.

$$\lambda = \frac{1}{\sigma} = \frac{1}{200} = 0.005$$

Now that the λ is defined, it is possible to rewrite the CDF:

$$CDF = 42.86\% = 1 - e^{-0.005x}$$

Using simple algebra, it is possible to find x as:

$$x = \frac{\ln\left(1 - 42.86\%\right)}{-0.005} = \frac{-0.55}{-0.005} = 112$$

The optimal quantity computed analytically is therefore:

$$q' = 800 + 112 = 912 \text{ units}$$

The expected profit associated to a quantity of 912, in the case of exponential distribution, is €2,545, significantly higher than the profit obtained with the uniform distribution.

4.3.2 Alignment of simulated findings

The results of the simulation are displayed in the figure 4.4.

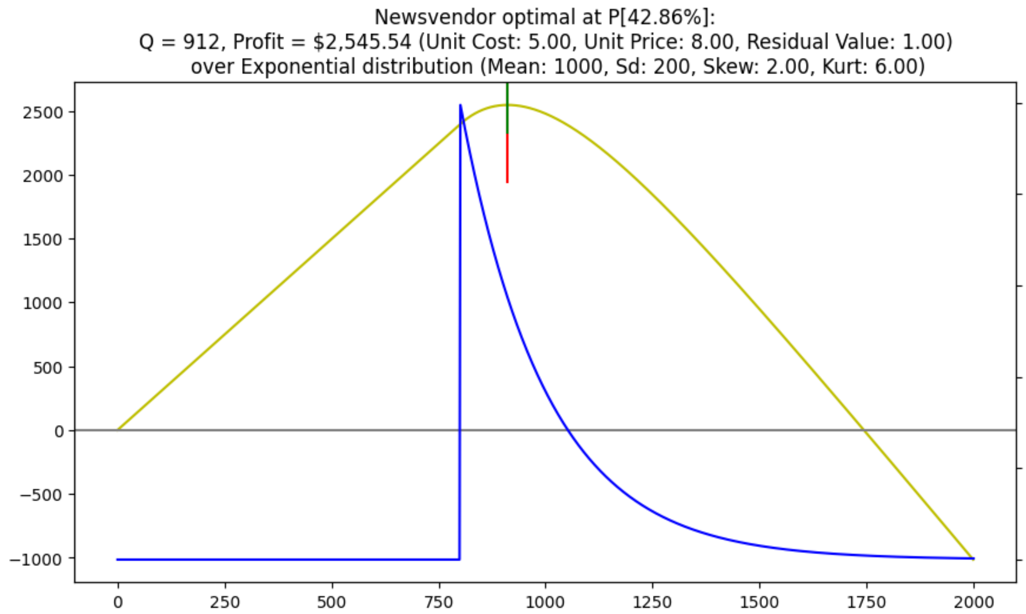

Figure 4.4: Plot of the probability (in blue) and profit (in yellow) in case of exponential demand distribution

Also in this context, the simulated solution confirms is validity and reliability, exactly matching the values obtained with the analytical approach. Compared to the same chart in case of uniform distribution, the yellow curve is way steeper, making the profit more sensitive to changes in quantity (even though starting from a way higher baseline).

In the figure 4.5 the sensitivity analysis performed for the uniform distribution is repeated also for the exponential one. Because the distribution is shifted by 800 units to the right, the optimal quantity in case of cumulative probability of 27.3% is 864, still very close to the distribution average (i.e., 1,000 units). In this case, compared to an optimal profit of €1,242, the gap with the optimal profit is only 12.1% lower (in the uniform distribution was around -70%)
This result is easily explainable by the significant shift to the right of the distribution, which make its starting point (cumulative percentage equal to

0%) happen at 800 units.
However, the most important takeaway from their comparison is that the distribution plays a significant role in the sensitivity of the profit curve, that cannot be summarized just by its mean.

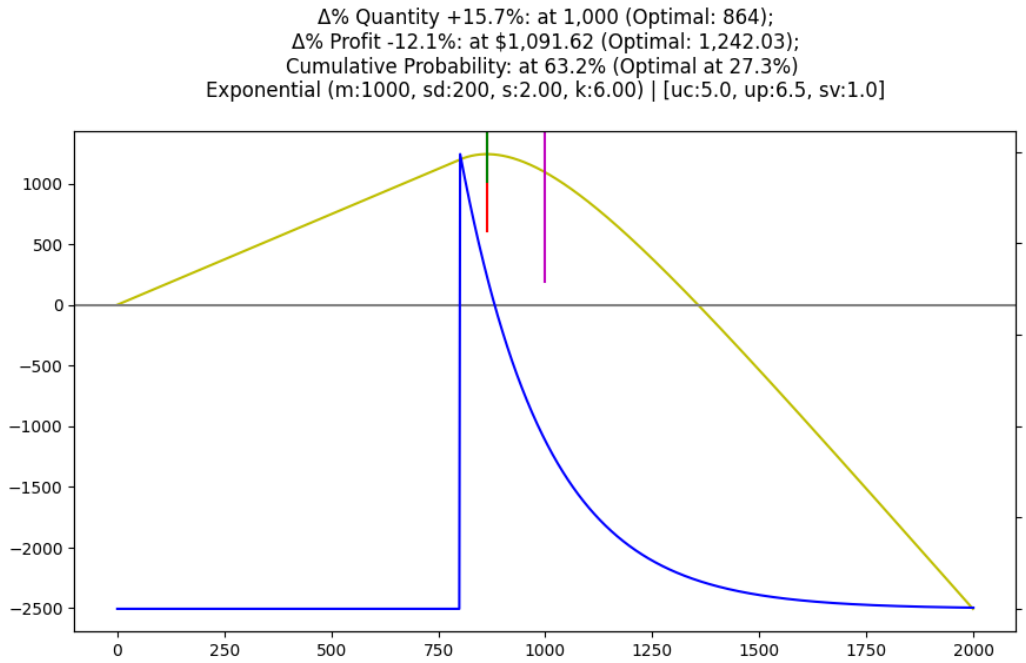

Figure 4.5: The impact of economic parameters in case of exponential distribution

4.4 Normal distribution

The normal distribution is a **symmetrical** distribution, having its mean, median (and mode) at the same quantity. The curve parameters, kept constant compared to the exponential, are:

- **Mean** equal to 1000 units
- **Standard Deviation** equal to 200 units

4.4.1 Analytical solution when following a normal distribution

Being the cumulative probability fixed to 42.86%, it corresponds in a standard normal distribution to a z of

$$z = N^{-1}(42.86\%) = -0.18$$

The quantity associated to z is obtainable through an easy transformation

$$z = \frac{q - \mu}{\sigma}$$

$$q = \mu + z \cdot \sigma = 1,000 - 0.18 \cdot 200 = 964 \text{ units}$$

The optimal profit associated is €2,450, slightly lower in comparison to the result obtained in the exponential case.

4.4.2 Alignment of simulated findings

In figure 4.6 it is once again confirmed the validity of the simulated solution.

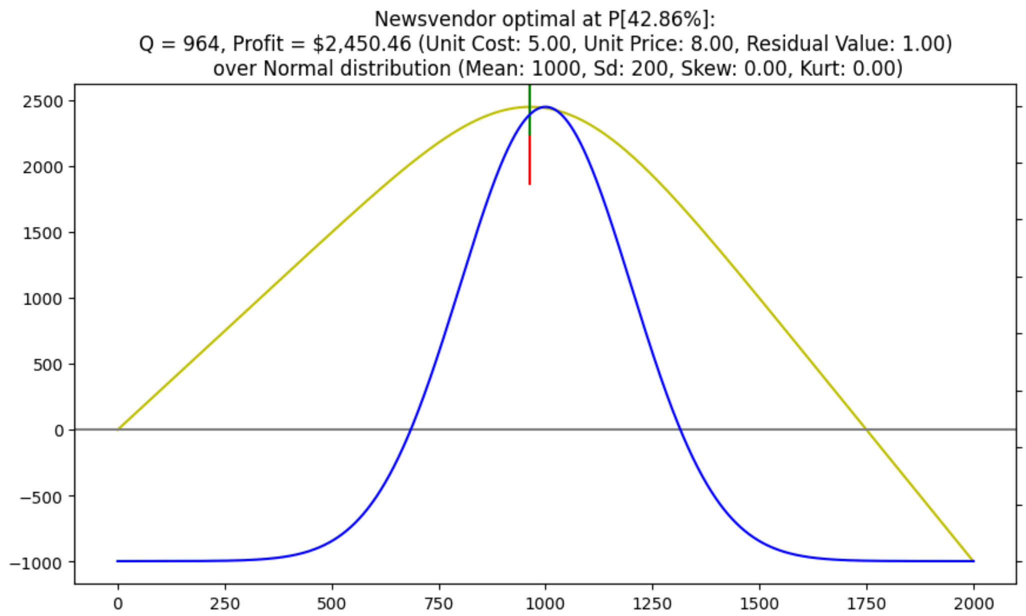

Figure 4.6: Plot of the probability (in blue) and profit (in yellow) in case of normal demand distribution

The results are quite aligned to those obtained in the previous subsection analyzing the exponential distribution, both in terms of profit and quantity. The same can be told for the second Newsvendor in figure 4.7 as well.
The bottom line of this section is that, even though the mean alone is not sufficient to estimate the sensitivity of the model, together with standard deviation they start to draw a decently accurate picture. However, even when combined, the information they provide is not enough to perfectly find the optimal solution of the Newsvendor.

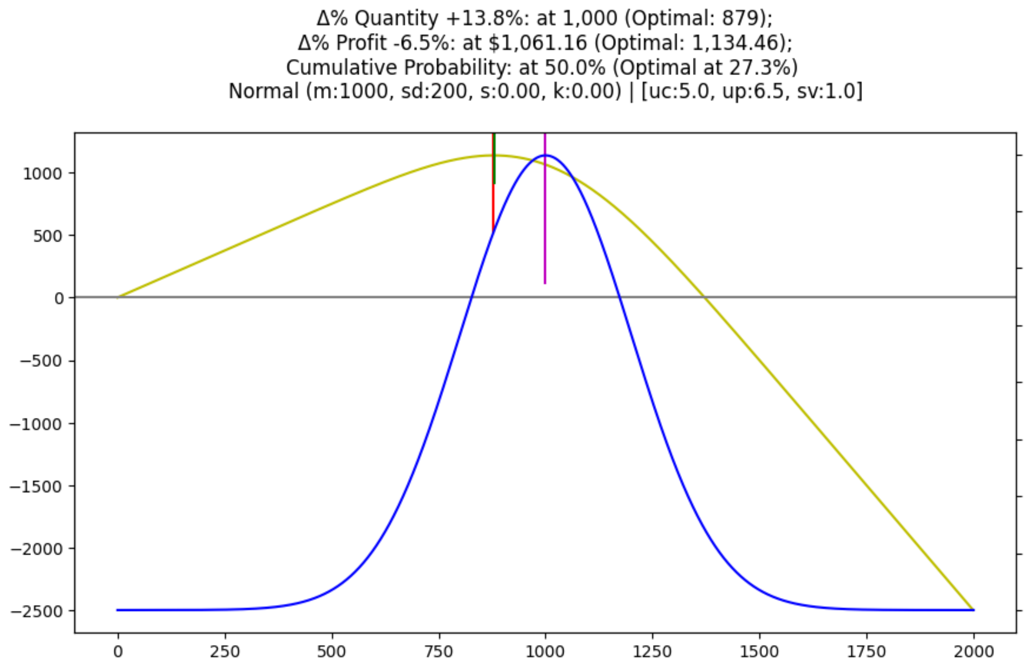

Figure 4.7: The impact of economic parameters in case of normal distribution

Chapter 5

Theoretical model validation both by simulation and mathematical proof

This chapter aims to find an answer to the first main question:

Is the Newsvendor theoretical model applicable to any distribution and economic parameters?

Thus far, it has been demonstrated that the simulation can accurately replicate analytical results, with the added advantage of being applicable to any type of distribution.
In the subsequent sections, we will evaluate ad-hoc distributions to ensure the robustness of the theoretical model also in specific and counter intuitive scenarios.

5.1 A counter-intuitive example based on special distributions

The key concept is that the critical fractile does not depend on the distribution structure, but only on the economic parameters associated to the model. This means that, according to the model, if the optimal cumulative probability is set to 40%, two distributions sharing the first half of the distribution will have the same optimal quantity, regardless of how different the second half is.

Starting from this premise, the objective of this section is to try to identify a scenario in which the model does not provide the correct results due to the particular type of distribution:

Distribution 1: figure 5.1

- The first 50% of the curve is a uniform distribution in the range [400, 600] units;
- The remaining 50% is also a uniform distribution, in the range [1,000, 1,200] units.

Distribution 2: figure 5.2

- The first 50% of the curve is a uniform distribution in the range [400, 600] units (identical to figure 5.1);
- The remaining 50% is again a uniform distribution, but its domain is in the range [1,800; 2,000] units.

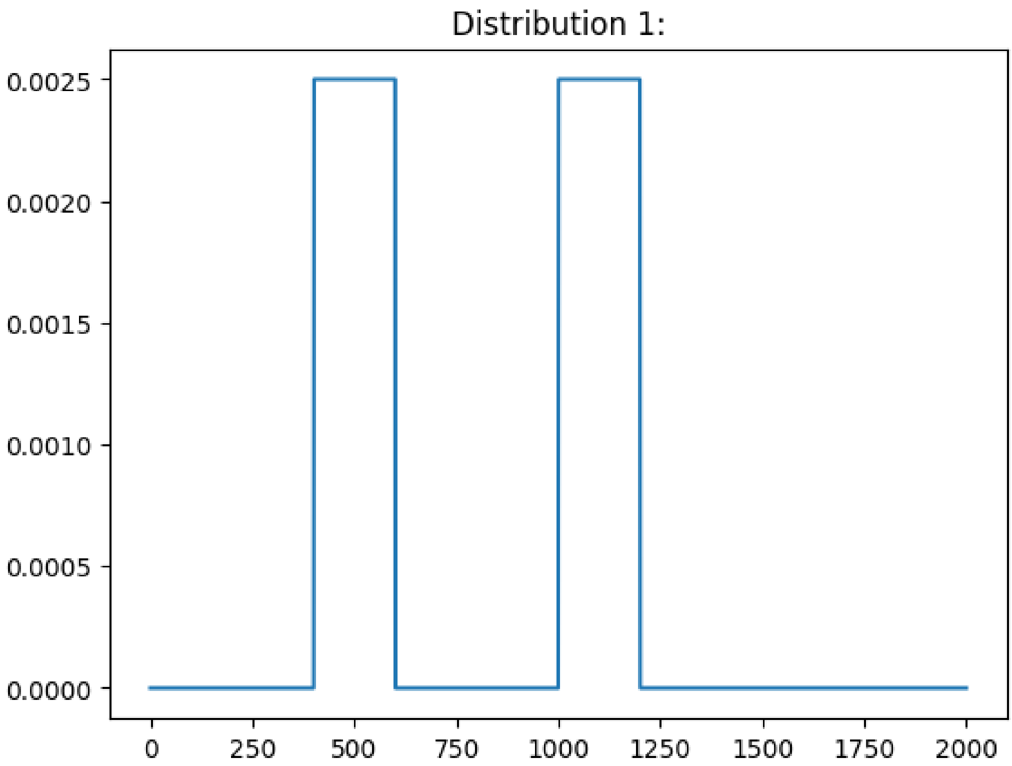

Figure 5.1: Distribution created as two separated uniforms with a range of 200 units to check if the Newsvendor theory works for any kind of distribution

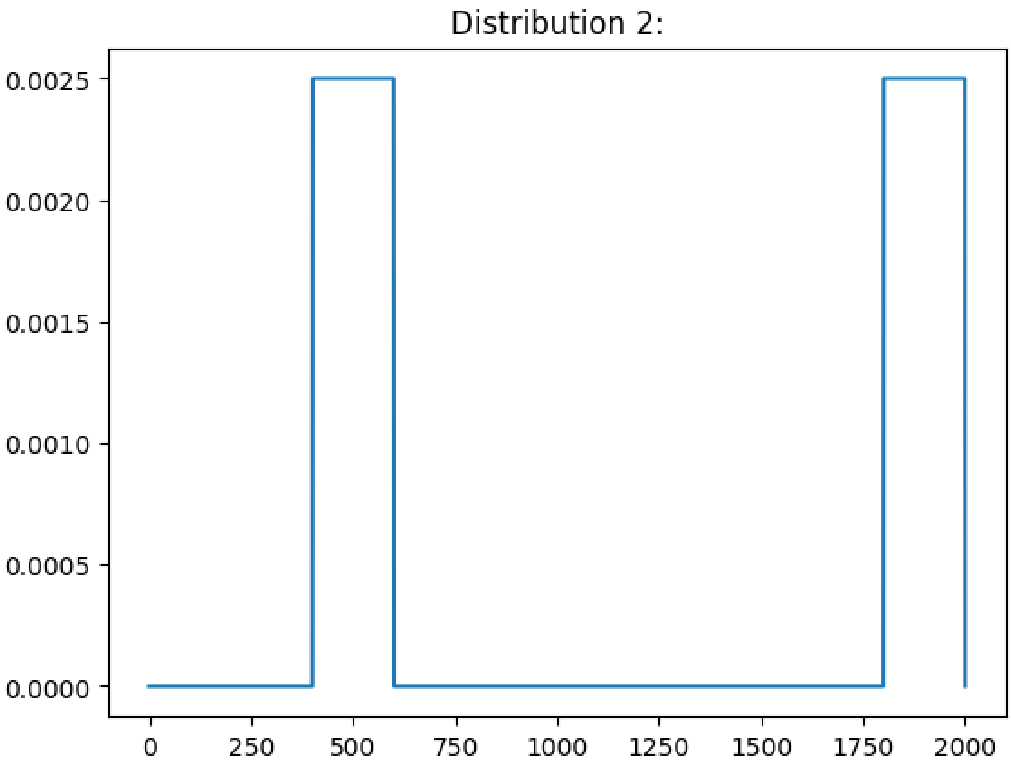

Figure 5.2: Benchmark distribution, identical to figure 5.1 in the first half, but having a higher mean for the second uniform component

At first sight, the second distribution seems to offer better conditions to the seller from every point of view.
In fact, in the first half the two distributions are identical, leading to the same realized quantity for any possible value of the critical fractile that falls within the range. However, in the remaining 50%, the second distribution provides a quantity 800 units higher for any value of the critical fractile above such threshold.

According to the model, if the economic parameters of both distributions are identical and set to result in a cumulative probability below 50%, the optimal quantity remains constant even though the expected value of the distribution is extremely different.

This behavior seems quite counter-intuitive: how can a distribution that is defined better or equal than another in any point, results in an identical

optimal quantity?
For the convenience of the reader, the economic parameters chosen for this test are set to the same values as before:

- Unit Price (up) = €8.00
- Unit Cost (uc) = €5.00
- Resale Value (sv) = €1.00
- Optimal Probability (P) = 42.86%

The simulation results, displayed in figure 5.3, clearly portrays the behavior of the Newsvendor in such circumstances.
First, it is confirmed that the optimal quantity actually remains the same between the two cases and the expected profit as well. It is worth noticing that the profit curve has the exact same behavior for the first half of both distributions, but they take extremely different paths immediately after they start to drift apart.
After a quantity of 1,000 units (the first point in which they differ), the second profit is extremely more stable and resilient to an overestimation of the optimal quantity than the first one. The profit of distribution 1) turns negative around 1,400 units, whereas the second distribution has a significantly positive profit throughout its whole domain [0, 2000].
However, the slope of the decrease is equal between units [1,100, 2,000] of the first Newsvendor and [1,900, 2,000] of the second (quantities that correspond to the two centers of the second uniform component in the two distributions). The same can be said for the range [1,000, 1,800] of the second distribution, where the slope remains the same of the range [600, 1,000], shared between the two distributions.

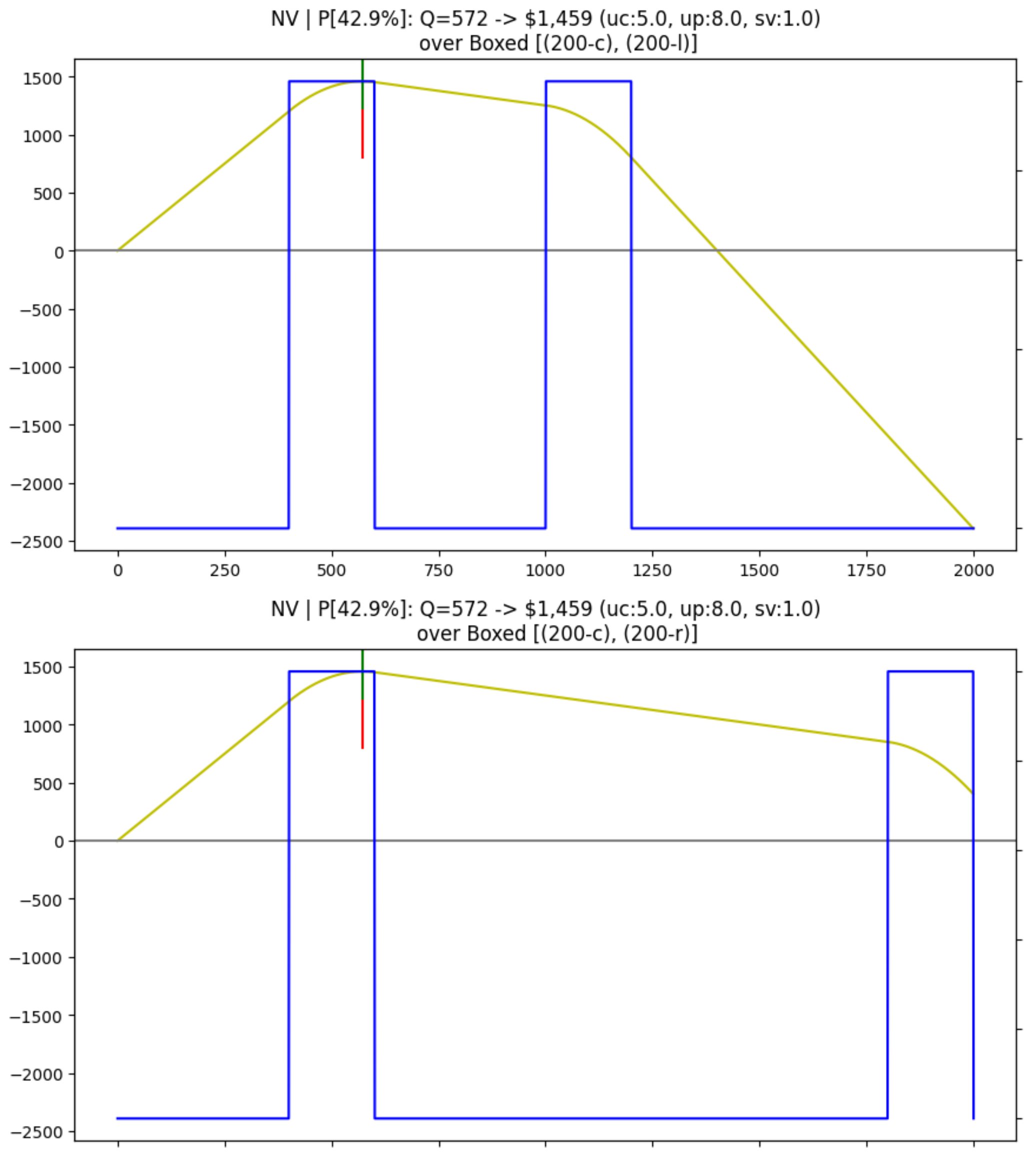

Figure 5.3: Empirical solution of the two Newsvendor that results in same optimal quantity and expected profit for both custom distributions

5.2 A graphical explanation of Newsvendor behavior

Trying to confute the theoretical validity of the Newsvendor for this type of distribution, the simulation ended up confirming once again its validity: the cumulative percentage of the distribution that correspond to the optimal quantity solely depends on the economic parameters and has nothing to do with the type of distribution chosen.
The purpose of the following subsections will be to generalize the findings obtained in the section 5.1, understanding all the possible scenarios arising from different Newsvendor models.

5.2.1 Impacts of distribution changes before the optimal quantity

If the distributions differ before the optimal quantity, but the cumulative probability indicated by the critical fractile occurs at the same quantity in all cases, the optimal quantity is expected to remain constant across scenarios. However, due to the varying spread of the density distribution before such quantity, the resulting expected profit is anticipated to differ.
In figure 5.4 are plotted three curves composed by three uniform distributions each. As expected, the optimal quantity is fixed to 710 units, but the expected profit increases as the left uniform shifts towards right, increasing the number of units sold on average.

5.2.2 Impacts of distribution changes at the optimal quantity

Changing the quantity associated with the optimal probability will obviously result in different optimal quantities and expected profits across the three scenarios.
In figure 5.5 it is possible to observe the specific behavior, where the optimal quantity almost doubles between the 710 units of the first chart and the 1,226 of the last one. The same cannot be said for the expected profit, that passes only from €1,781 to €2,125. This can be explained by the fact the first 33% is in common between the two distributions, and the extra profit arises only from the remaining 9% of cumulative probability.

5.2.3 Impacts of distribution changes after the optimal quantity

In figure 5.6 it is confirmed the behavior seen in section 5.1, where the changes in the distribution happen only after the optimal quantity, that does not varies. In addition, also the optimal expected profit remains identical in all the charts.

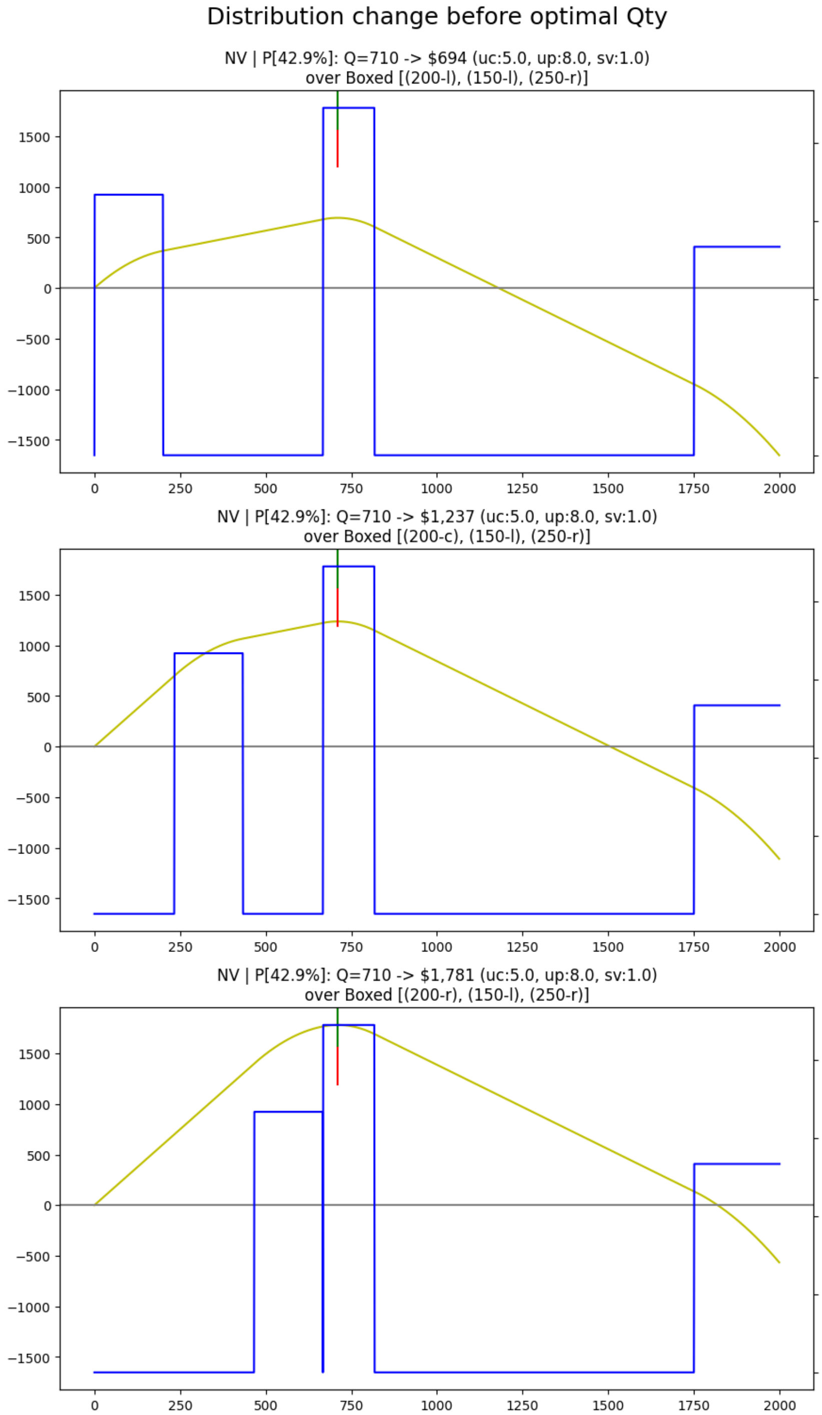

Figure 5.4: Impact on quantity and expected profit of a change in distribution happening before the optimal quantity

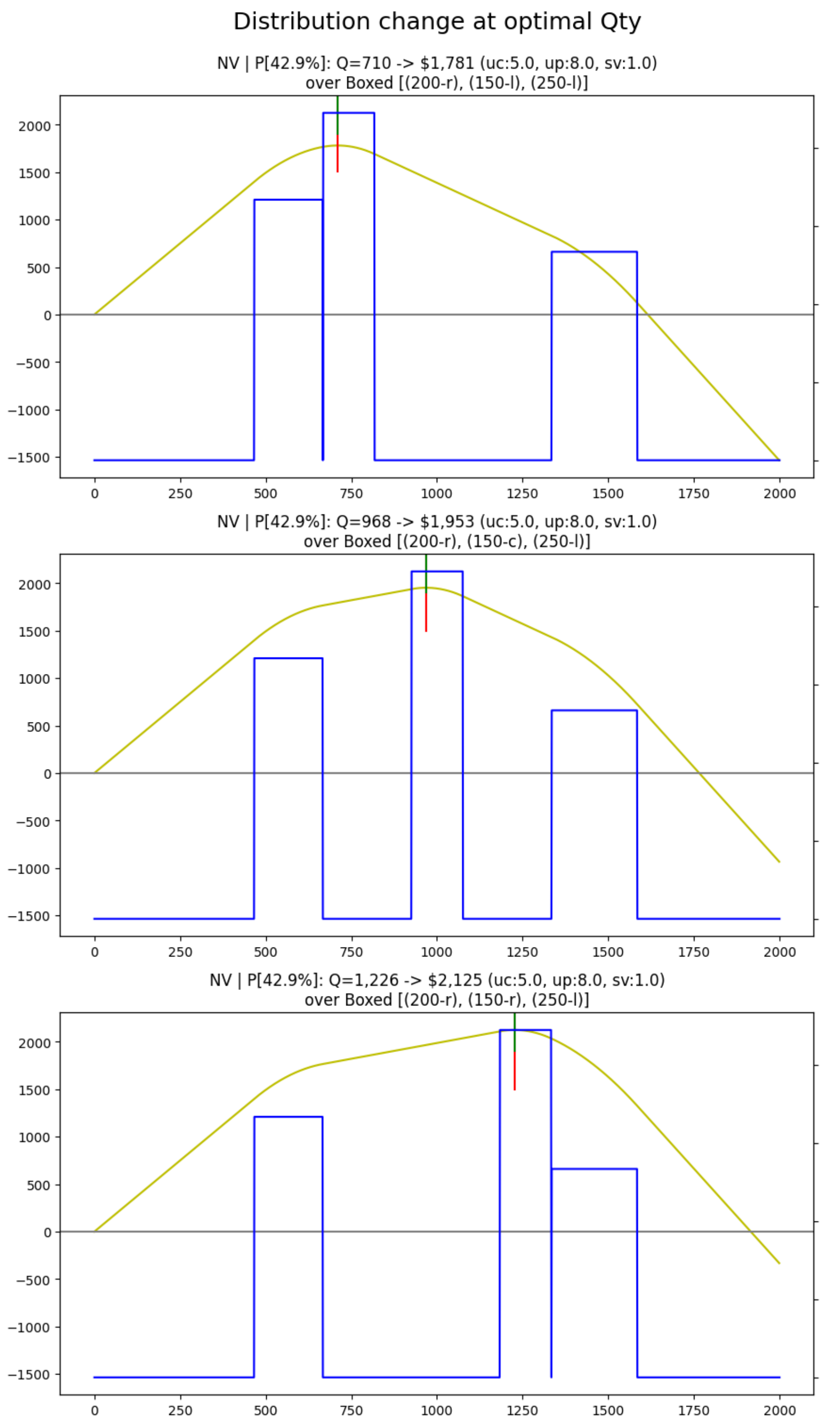

Figure 5.5: Impact on quantity and expected profit of a change in distribution happening at the optimal quantity

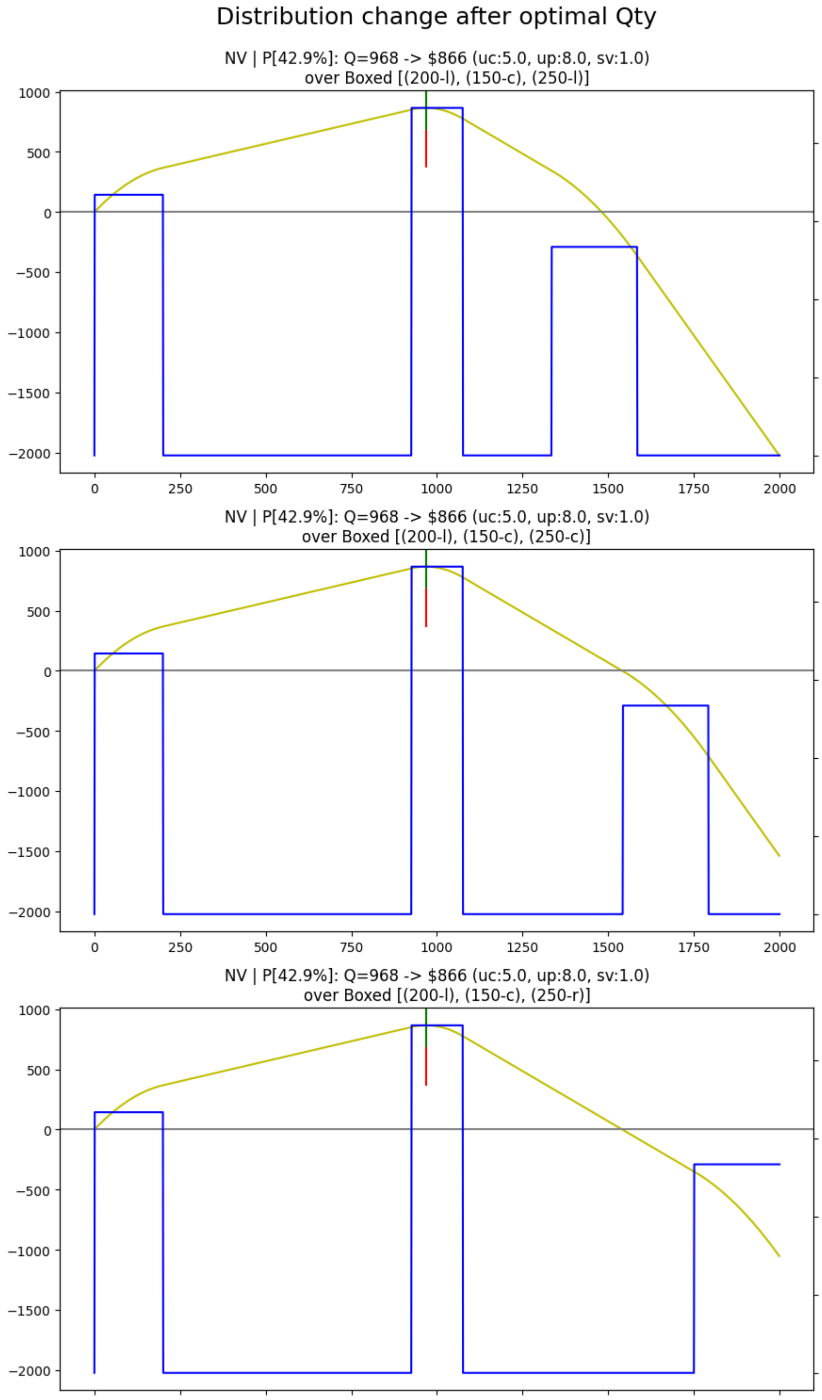

Figure 5.6: Impact on quantity and expected profit of a change in distribution happening after the optimal quantity

5.3 Mathematical explanation of profit behavior at distribution changes

The key takeaway from the previous section is that the optimal profit is influenced solely by the structure of the distribution up to the optimal quantity, remaining completely unaffected by its definition beyond that point.
Therefore, it is possible to state the following:

The expected profit associated to a quantity q' is impacted only by the portion of the demand distribution up to q'.

The explanation of this behavior is rooted on the mathematical construction of the model itself.
To better explain it, we will consider a simplified discrete distribution, defined as follows:

- 10 units in the left 10%;
- 30 units in the following 20%;
- 60 units for another 20%;
- 200 units after the median and for the remaining 50%.

To prove the statement, the expected profit (also known in financial contexts as EBIT, abbreviation for "Earnings Before Interests and Taxes") is reported in figure 5.7 at 4 different quantities (30, 60, 65, 100), decomposed in its three components: revenues, costs (COGS, abbreviation for "Cost Of Goods Sold") and residual value (Sale).
The final expected profit is then the average of the profits obtained at the different demand levels weighted by their realization probabilities.
In the table, a row is associated to a specific quantity (and represents the different ex-ante possibilities for the decision maker), whereas a column is linked to a specific demand level (representing the probabilities within the decision taken).

5.3.1 Effects of different quantities on revenues calculation

Given an actual purchased quantity q_p, the revenues are computed at each demand d level as:

$$Revenues(q_p, d) = min(q_p, d) \cdot up$$

The revenues are therefore function of both quantity and demand, however their relationship is a $min()$ function: fixed a quantity q_p, there is no additional positive impact for any demand exceeding the chosen quantity, even in the extreme cases. Looking at the first row of figure 5.7, between a demand of 10 and 30 there is a significant rise in revenues, but afterward it remains perfectly flat for any higher value of demand.
The same applies to the other rows as the demand exceeds the selected quantity.
This behavior is coherent with the conclusions drawn in the subsection 5.2.1, where different profits are associated to distribution changes before the optimal quantity. in fact, different demand levels before the selected quantity impact the $min()$ function and provide different results.
For the same reason, the $min()$ function also prevents any statistical revenues improvement once the selected quantity is exceeded, as shown in subsection 5.2.3

5.3.2 Effects of different quantities on costs calculation

Given an actual purchased quantity q_p, the costs are computed at each demand level d as:

$$Costs(q_p) = -q_p \cdot uc$$

The relationship shows a perfect linear dependence, where each additional quantity purchased decreases the profit by €uc. In fact, each row of the COGS table in figure 5.7 has constant costs, that correspond to the purchased quantity multiplied by the factor uc.
The effect is thus fixed and solely dependent on the purchased quantity, unaffected by the realized demand. Consequently, it acts as a probability-neutral component in profit calculation.

5.3.3 Effects of different quantities on residual value calculation

Given an actual purchased quantity q_p, the residual values are computed at each demand d level as:

$$\text{Residual Value}(q_p, d) = max(q_p - d, 0) \cdot sv$$

The result associated to this component of the profit calculation is quite similar to the revenues case:
In fact, even if the function is a $max()$, it does not return the maximum

between q_p and d, but between their difference and 0. If the demand significantly exceeds the quantity, the $max()$ prevents any negative impact on profit.
The sale table in figure 5.7 clearly highlights this behavior in each row.
The residual value, likewise the direct revenues, changes the profit for values of the demand lower than the quantity, but does not impact at all the remaining cases.

5.3.4 Combining all effects together

The previous sub-sections mathematically explain the behavior of the Newsvendor, broken down in its components.
The EBIT table in figure 5.7, clearly portrays in the first two rows that the profit is constant at any point of the demand distribution after the chosen quantity. Examining the first row reveals that regardless of the distribution's value after the first cell, the profit reaches its maximum at the chosen quantity of 30 units and cannot be enhanced.
Among the quantities, the highest expected profit belongs to the scenario with 60 units.
Comparing second and third rows:

- The revenues are identical for the first three columns (50% of the distribution), being slightly different only on the second half because of the extra 5 units;
- The costs are instead fixed and totally deterministic. The extra five units results in a certain extra cost of €25;
- The behavior of the resale value is instead a little more complex to be analyzed properly, because it is different across most rows and columns. The residual value has the lowest impact, because it simply weights out a part of the extra costs between the two rows by a factor of $\frac{sv}{uc} = \frac{1}{5}$. In fact, it reduces the extra costs from the €25 expected to just $25 \cdot \frac{4}{5} =$ €20.

This behavior is easily explainable by the fact that a Newsvendor with economical parameters

$$\text{unit price } up = 8, \text{ unit cost } uc = 5 \text{ and residual value } sv = 1$$

is equivalent in the computation of the optimal quantity to

$$up = 7,\ uc = 4 \text{ and } sv = 0$$

resulting in the same optimal probability of 42.86%.
In fact, the residual value can be considered as a discount on the effective cost, that also reduces the profit because we are selling at *up* something with an intrinsic value of *sv*. Therefore, the profit varies accordingly to two main variables:

- Each additional unit, in this specific case, provides an extra profit that is equal to its price. However, this positive effect is only effective when the demand is higher enough to allow for that extra unit to be sold. Therefore, the average impact is the price weighted by the probability it has to be effective, 50% in this case.

- On the other hand, every extra quantity has the certainty to increase the costs associated by the unit cost associated to the Newsvendor.

In this specific case, being

$$up \cdot prob = 7 \cdot 50\% = 3.5 < 4 = uc$$

the scenario at quantity 60 units has an higher profit than the one at 65 by

$$€(4 - 3.5) \cdot \Delta\text{units} = €0.5 \cdot 5 = €2.5$$

from €103 to €100.5.

	1	2	3	4	TOT
p	10.0%	20.0%	20.0%	50.0%	100.0%

price	€ 8.00
cost	€ 5.00
sale value	€ 1.00

Num	3.00
Den	7.00
P	42.9%

D1	10	30	60	200

Revenue

30	€ 80.00	€ 240.00	€ 240.00	€ 240.00
60	€ 80.00	€ 240.00	€ 480.00	€ 480.00
65	€ 80.00	€ 240.00	€ 480.00	€ 520.00
100	€ 80.00	€ 240.00	€ 480.00	€ 800.00

COGS

30	€ 150.00	€ 150.00	€ 150.00	€ 150.00
60	€ 300.00	€ 300.00	€ 300.00	€ 300.00
65	€ 325.00	€ 325.00	€ 325.00	€ 325.00
100	€ 500.00	€ 500.00	€ 500.00	€ 500.00

Sale

30	€ 20.00	€ -	€ -	€ -
60	€ 50.00	€ 30.00	€ -	€ -
65	€ 55.00	€ 35.00	€ 5.00	€ -
100	€ 90.00	€ 70.00	€ 40.00	€ -

EBIT

30	-€ 50.00	€ 90.00	€ 90.00	€ 90.00
60	-€ 170.00	-€ 30.00	€ 180.00	€ 180.00
65	-€ 190.00	-€ 50.00	€ 160.00	€ 195.00
100	-€ 330.00	-€ 190.00	€ 20.00	€ 300.00

W. EBIT

30.00	-€ 5.00	€ 18.00	€ 18.00	€ 45.00
60.00	-€ 17.00	-€ 6.00	€ 36.00	€ 90.00
65.00	-€ 19.00	-€ 10.00	€ 32.00	€ 97.50
100.00	-€ 33.00	-€ 38.00	€ 4.00	€ 150.00

W. EBIT

30.00	€ 76.00
60.00	€ 103.00
65.00	€ 100.50
100.00	€ 83.00

Figure 5.7: This table portrays the profit associated to 4 different quantities in a simplified scenario, by dividing its computation in the different components

5.4 Mathematical proof valid for any distribution

The purpose of this section is to provide a formal mathematical explanation, to the empirical behavior described previously. To do so, we will borrow the Littlewood's rule [Yeoman and McMahon-Beattie(2017)], developed by Ken Littlewood, that was an effective solution method for the seat inventory problem in case of two fare classes. Being Pr_1 and Pr_2 the two different fares and

$$Pr_1 > Pr_2$$

the relationship of the two prices. If the demand of the less profitable Class2 is served first if compared to the demand of Class1 (and therefore should be fixed beforehand), demand for Class2 should be accepted until the equation below is satisfied for x:

$$Pr_2 \geq Pr_1 \cdot P[D_1 \geq x]$$

where

- Pr_1 is the profit associated to Class1;
- Pr_2 is the profit associated to Class2;
- D_1 represents the cumulative probability function for the demand of Class1;
- x is the remaining capacity.

It means that additional demand for Class2 should be accepted until its profit is higher or equal than the profit of Class1, weighted by its probability to fill the capacity available. Rewriting the equation in a slightly different way, demand for Class1 should be reserved for the biggest quantity q for which

$$Pr_2 \leq Pr_1 \cdot P[D_1 \geq q]$$

$$\implies \frac{Pr_2}{Pr_1} \leq P[D_1 \geq q]$$

$$\implies \frac{Pr_2}{Pr_1} \leq 1 - P[D_1 \leq q]$$

$$\implies P[D_1 \leq q] \leq 1 - \frac{Pr_2}{Pr_1}$$

$$\implies D_1(q) \leq 1 - \frac{Pr_2}{Pr_1}$$

$$\implies q \leq D_1^{-1}\left(1 - \frac{Pr_2}{Pr_1}\right)$$

Therefore, the optimal quantity q_1 that should be reserved for Class1 can be obtained by applying the $floor()$ function to q

$$q_1 = \left\lfloor D_1^{-1}\left(1 - \frac{Pr_2}{Pr_1}\right)\right\rfloor$$

It is of extreme importance to notice that the calculations above are valid for any family of distributions, do not requiring symmetry or other particular properties.
Therefore, if it were possible to link the Newsvendor Model to the Littlewood's rule, this would be an additional confirmation of the model validity for any kind of existing distribution. In reality, creating a connection between the two is fairly easy, in fact The fare associated to Class1 is

$$Pr_1 = up$$

that is nothing else than the profit associated to the sale of an item (given the fact that it has been already purchased by the seller); whereas to Class2 is

$$Pr_2 = uc$$

that is an opportunity cost of not purchasing an additional product.
In fact, it is true that purchasing one less unit could be a loss of up, but it is also a sure saving of uc. The optimal quantity q_1 of the best scenario can be calculated as:

$$\begin{aligned} q_1 &= \left\lfloor D_1^{-1}\left(1 - \frac{uc}{up}\right)\right\rfloor \\ &= \left\lfloor D_1^{-1}\left(\frac{up - uc}{up}\right)\right\rfloor \end{aligned}$$

Instead, in the slightly more complex case where residual after-season discounted value is different from 0, the relationship between the two fares remains valid, but needs to be slightly adjusted as follows:

$$Pr_1 = up - sv$$

because the payment in cash of up by the final user, for an item with residual value of sv, results in a net profit of $up - sv$ for the seller company and

$$Pr_2 = uc - sv$$

given that an item is purchased at a price of uc by the seller company, and it has an intrinsic value of sv, the transaction results in a net cost of $uc - sv$. By purchasing one less unit, the company saves such amount.
The formulation of the optimal quantity q_1' in this more comprehensive scenario is a generalization of the previous result. It is defined as:

$$\begin{aligned} q_1' &= \left\lfloor D_1^{-1}\left(1 - \frac{uc - sv}{up - sv}\right)\right\rfloor \\ &= \left\lfloor D_1^{-1}\left(\frac{up - uc}{up - sv}\right)\right\rfloor \end{aligned}$$

Both proposed solutions (with and without residual value) align precisely with the theoretical formulation of the Newsvendor model found in the literature.

Part III

Sensitivities in the Newsvendor

Chapter 6

Impact on optimal solution in case of pure mean changes

Throughout this chapter, the main focus will be to assess the sensitivity of the Newsvendor model to changes in the distribution mean.
The business meaning of the analysis, as well as its final objective, is to answer the following question:

If I was expecting a distribution centered at q units, what would be the impact to the optimal quantity and the maximum expected profit if the real probability function has instead its mean set to a different quantity q'?

This is a very important business question. In fact, in the Newsvendor not only it is impossible to predetermine the actual demand, but also its underlying distribution is only estimated.
Assessing the robustness of the Newsvendor model in the presence of mean estimation errors is crucial, as it allows decision-makers to evaluate the risk inherent in the model.

6.1 Changes in the mean keeping optimal probability near 50%

For this specific sensitivity study on mean uncertainty, the transformation is as simple as a shift of the distribution by x units. Figure 6.1 portrays the impact of a variation in the mean for many normal distribution having the same standard deviation ($\sigma = 150$ units), but different mean values.
The baseline is the third plot, having a mean of 1,000 units, and it is benchmarked versus many different distributions, ranging from a mean of 850 units

(-150) to 1,150 units (+150).

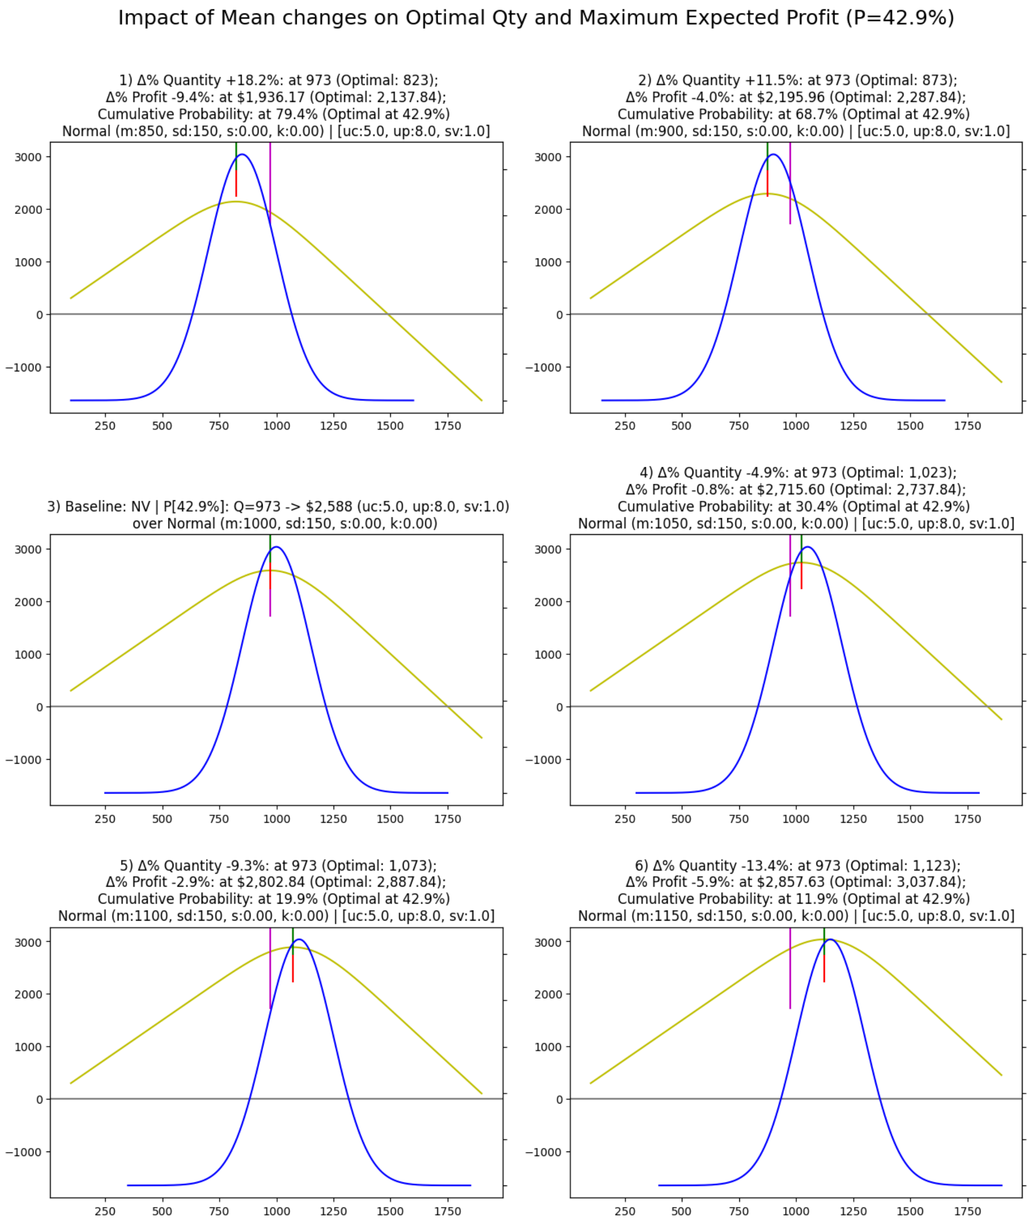

Figure 6.1: Impact of mean changes in case of a Normal distribution keeping the standard deviation constant

In the first chart it is showcased the scenario of expected distribution having mean of 1,000 units and optimal quantity equal to 973 units (at a

cumulative probability of 42.9%). However, the real distribution is shifted by -150 units, corresponding to a standard deviation. Thus, the real mean is set to 850 units and the actual optimal quantity 27 units away, at 823. Such a leftward shift in the distribution results in the selected quantity occurring at a cumulative probability of nearly 80%, significantly deviating from the optimal value of 42.9%. Consequently, the profit is reduced by 9.4% (approximately €200) compared to the optimal case of €2,138.
In the normal distribution, 68.2% of the distribution is concentrated in the symmetrical range $[\mu - \sigma, \mu + \sigma]$. As a result, if the distribution is shifted by one standard deviation to the left, and the reference benchmark quantity is close to a cumulative probability of 50%, the shift would bring the actual cumulative probability to nearly 80% at 973 units.
In conclusion, even with a significant discrepancy between the real optimal quantity and the purchased one, the expected profit remains reasonably close to the optimal value. However, the impact becomes more pronounced as the distribution deviates farther from the expected.
It is noteworthy that in this specific scenario, selecting a quantity of 1,500 units would results in an expected profit of €0. Such neutral profit quantity decreases more rapidly than anticipated by the pure leftward shift: in fact, because the distribution begins at a lower quantity, the model consequently accumulate a lower profit in its optimal. Therefore, although the threshold heavily depends on economic parameters, it is clear that the mean significantly influences both optimal quantity and expected profit determination.

The last chart in figure 6.1 represents the opposite scenario, where the real distribution is shifted again by a standard deviation, landing to 1,150 units for the actual mean and 1,123 units for the real optimal quantity.
In this case, the cumulative probability decreases by approximately 30 percentage points (pp) to 11.9%. Such value of cumulative probability is at the very beginning of the distribution, yet the associated profit is already €2,857, only 5.9% below the optimal profit of €3,038.
The expected profit is extremely high, and this occurs because the distribution is shifted towards right while keeping fixed the standard deviation.
The additional units sold at the optimum is

$$\Delta units = 1,150 - 850 = 300$$

and the profit associated to each unit sold is

$$up - uc = 8€ - 5€ = 3€$$

Therefore, the extra expected profit at the same cumulative percentage the distribution can be calculated as

$$300 \text{ units} \cdot 3€ = 900€$$

and the result is coherent with the maximum profits obtained in the two scenarios:

$$3,038€ - 2,138€ = 900€$$

In conclusion, even if the demand is underestimated and the expected profit will be lower than the optimal, it will still be better than the maximum expected profit in the baseline scenario due to this upwards shift in the profit curve.

Looking at the table in figure 6.2, it is possible to observe this behavior in many more scenarios. In fact, it is obvious that any increase in the value of the mean results in the same rise in the optimal quantity and a 3x increase in the € amount of the expected profit.
If we look at the behavior of the expected profit at a quantity of 973 (the optimal in the baseline case) for values in the range $[-10, 10]$, the behavior is similar to the optimal expected profit.
However, the farther away from the baseline case, the bigger the error accrued over the step of 3€ expected. At a Δ quantity of +20, the change in profit associated is +56.37€ (the optimal increases by +60€); whereas at a Δ quantity of -20, the Δ profit is -63.69€.

The increase and the decrease follow the behavior of the optimal profit, but having the additional effect of the profit decrease going farther away from the mean. In fact, the profit gain resulting from the increased quantity will be lower than the 3€ per unit expected, whereas the loss in case of lower quantity will be higher, due to the double effect of mean decrease and sub-optimal profit.
Interestingly enough, in first approximation, the profit gap between -20 and +20 is 120€. Considering the 40 units of difference, it results in an average change in profit of almost exactly 3€ for each unit. This is not a standard property of the Newsvendor (it is guaranteed only in case of symmetrical distribution and optimal quantity coinciding with the mean, thus at 50% cumulative probability), but it is an acceptable approximation, around the optimal, in many case of known distributions.
In the baseline case the distribution is symmetrical and the optimal probability is around 43%, resulting in a quite symmetrical profit distribution. Considered this symmetry, the profit at $+x$ is very similar to the one at $-x$

from the optimal and the profit change between a shift of $+x$ and $-x$ would be the pure upward lift of $2x \cdot 3$€.

x_shift	mean	Otimal_x	Max E[Profit]	y_shift	y_at_973	y_shift_973
+50	1,050	1,023	€ 2,737.84	€ 150.00	€ 2,715.60	€ 127.76
+45	1,045	1,018	€ 2,722.84	€ 135.00	€ 2,704.76	€ 116.92
+40	1,040	1,013	€ 2,707.84	€ 120.00	€ 2,693.51	€ 105.66
+35	1,035	1,008	€ 2,692.84	€ 105.00	€ 2,681.83	€ 93.99
+30	1,030	1,003	€ 2,677.84	€ 90.00	€ 2,669.73	€ 81.88
+25	1,025	998	€ 2,662.84	€ 75.00	€ 2,657.19	€ 69.35
+20	1,020	993	€ 2,647.84	€ 60.00	€ 2,644.21	€ 56.37
+18	1,018	991	€ 2,641.84	€ 54.00	€ 2,638.90	€ 51.06
+16	1,016	989	€ 2,635.84	€ 48.00	€ 2,633.52	€ 45.67
+14	1,014	987	€ 2,629.84	€ 42.00	€ 2,628.06	€ 40.22
+12	1,012	985	€ 2,623.84	€ 36.00	€ 2,622.53	€ 34.69
+10	1,010	983	€ 2,617.84	€ 30.00	€ 2,616.93	€ 29.09
+8	1,008	981	€ 2,611.84	€ 24.00	€ 2,611.26	€ 23.42
+6	1,006	979	€ 2,605.84	€ 18.00	€ 2,605.51	€ 17.67
+5	1,005	978	€ 2,602.84	€ 15.00	€ 2,602.62	€ 14.77
+4	1,004	977	€ 2,599.84	€ 12.00	€ 2,599.70	€ 11.85
+3	1,003	976	€ 2,596.84	€ 9.00	€ 2,596.76	€ 8.92
+2	1,002	975	€ 2,593.84	€ 6.00	€ 2,593.81	€ 5.96
+1	1,001	974	€ 2,590.84	€ 3.00	€ 2,590.83	€ 2.99
+0	1,000	973	€ 2,587.84	€ -	€ 2,587.84	€ -
-1	999	972	€ 2,584.84	-€ 3.00	€ 2,584.83	-€ 3.01
-2	998	971	€ 2,581.84	-€ 6.00	€ 2,581.81	-€ 6.04
-3	997	970	€ 2,578.84	-€ 9.00	€ 2,578.76	-€ 9.08
-4	996	969	€ 2,575.84	-€ 12.00	€ 2,575.70	-€ 12.15
-5	995	968	€ 2,572.84	-€ 15.00	€ 2,572.61	-€ 15.23
-6	994	967	€ 2,569.84	-€ 18.00	€ 2,569.51	-€ 18.33
-8	992	965	€ 2,563.84	-€ 24.00	€ 2,563.26	-€ 24.59
-10	990	963	€ 2,557.84	-€ 30.00	€ 2,556.92	-€ 30.92
-12	988	961	€ 2,551.84	-€ 36.00	€ 2,550.52	-€ 37.32
-14	986	959	€ 2,545.84	-€ 42.00	€ 2,544.04	-€ 43.80
-16	984	957	€ 2,539.84	-€ 48.00	€ 2,537.49	-€ 50.36
-18	982	955	€ 2,533.84	-€ 54.00	€ 2,530.86	-€ 56.99
-20	980	953	€ 2,527.84	-€ 60.00	€ 2,524.16	-€ 63.69
-25	975	948	€ 2,512.84	-€ 75.00	€ 2,507.07	-€ 80.77
-30	970	943	€ 2,497.84	-€ 90.00	€ 2,489.53	-€ 98.32
-35	965	938	€ 2,482.84	-€ 105.00	€ 2,471.52	-€ 116.33
-40	960	933	€ 2,467.84	-€ 120.00	€ 2,453.04	-€ 134.80
-45	955	928	€ 2,452.84	-€ 135.00	€ 2,434.10	-€ 153.74
-50	950	923	€ 2,437.84	-€ 150.00	€ 2,414.70	-€ 173.15

Figure 6.2: Table that portrays the impact of mean changes to better evaluate the sensitivity in case of a normal distribution and constant standard deviation

6.2 Changes in the mean in case of skewed Newsvendor

In the previous section, the analyzed Newsvendor model had economic parameters resulting in an optimal quantity near to the mean.
The follow-up question is therefore:

What is the sensitivity to mean changes when the optimal quantity is near to the extremes of the distribution?

Such a question is extremely interesting for two reasons:

- it enables the generalization to any Newsvendor Model of the findings observed in the previous section;
- it provides insights of the effects that such changes would have on the final expected profits, allowing decision makers to tackle with more confidence uncertain scenarios.

In figure 6.3 it is portrayed the behavior of two different Newsvendors having very different economic parameters:

- **NV1:** unit price: 6€, unit cost: 5€, residual value: 1€, resulting optimal probability = 20.0%;
- **NV2:** unit price: 10€, unit cost: 5€, residual value: 4€, resulting optimal probability = 83.3%.

The first Newsvendor provides a profit of 1€ for every unit sold, and has a net cost of 4€ for each one purchased. It is a very risky business and therefore it is very important to stay in the most certain part of the distribution: the beginning.
By opposite, the second one can rely on a profit of 5€ per unit sold and just 1€ dollar of overstock risk per unit. It is extremely advantageous and thanks to the very low risk the optimal quantity is far away in the right tail of the distribution.

The first point of interest, looking at the charts 5) and 6) in figure 6.3 (the baselines of the two Newsvendors), is the optimal quantity of the two models: 874 units for the first and 1,145 for the second. They are around 270 units apart, that corresponds to a 31% gap between them.
The profit however varies on a whole different level, ranging from the 790€

of the first model to 4,775€ of the second, a multiplying factor of 6x.

If compared to their respective benchmarks, in both charts 1) and 2) the mean has been shifted by 150 units on the left (to 850 units) and the optimal profit decreased by 150€ and 750€ respectively (aligned to the 1€ and 5€ of additional profit per extra unit sold). Even though the expected profit decrease (due to the vertical shift of the curve) is proportional to the profit associated to each unit sold, the same cannot be said for the loss due to the sub-optimal quantity chosen: the first Newsvendor has a loss from the optimal of 130€, whilst in the second scenario the loss is mitigated to 80€, which in percentage correspond to a -20% and -2% respectively.

Shifting the distributions to the left, the chosen quantity automatically occurs at a higher cumulative probability of the demand distribution. The profit curve in chart 2) is way less sensitive to an error in mean estimate in the right tail rather than in the left one:
Considering a point where the cumulative probability is almost 100% in both the distributions, the loss in profit for each additional unit, as seen in section 5.4, is:

$$NV1 : \Delta Profit = (1 - P) \cdot (up - sv) - (uc - sv) = 1€ - 5€ = -4€$$

$$NV2 : \Delta Profit = (1 - P) \cdot (up - sv) - (uc - sv) = 4€ - 5€ = -1€$$

These numbers represent the impact of an extra unit purchased on the expected profit once the distribution reaches its very end and the cumulative probability borders the 100%. It is obvious that the first distribution drops way faster than the second, regardless of the way lower optimal profit amount.

It is important to understand that those profit reductions happens only when the distribution is close to 100%:
In charts 4) and 2), the probability ranges between 95% and 97.5%, shifting the mean by 50 units. The difference between optimal and sub-optimal profit is $4,275 - 4,235 = 40€$ for a 100 units shift from the base case and $4,025 - 3,946 = 80€$ for 150 units. Therefore, the last decrease of 50 units makes the expected profit fall by $80 - 40 = 40€$, close to the 1€ of profit loss per extra unit purchased and not sold.
It is slightly lower than 1€ per unit because the cumulative probability is in the range [95%,97.5%], lower than 100%, that partially factors in the very high profit associated to each unit sold (6€). Using analytical methods, it is also possible to compute the average probability over the interval as:

$$1 - \frac{40}{50} = 1 - 0.8 = 0.2 = (1 - P) \cdot 6$$

that becomes

$$P = 1 - \frac{2}{60} = 96.7\%$$

perfectly in line with the expected range.

Generalizing the findings of the previous calculation, it is therefore possible to estimate the drop in expected profit due to a shift in the mean:

$$Profit_{Sub} = Profit_{BO} + \Delta Q \cdot \{(up - uc) - [(1 - P_{avg}) \cdot (up - sv) - (uc - sv)]\}$$

where $Profit_{Sub}$ is the sub-optimal expected profit that we are trying to estimate, $Profit_{BO}$ is the optimal profit of the baseline scenario and P_{avg} is the average between the cumulative probabilities at the beginning and the end of the range.
The component $(up - uc)$ is the upward (downward) shift of the profit curve due to each additional (decremental) quantity, whereas $[(1 - P_{avg}) \cdot (up - sv) - (uc - sv)]$ is the average decrease due to each unit of distance from the optimal point (the impact is always negative).
Trying to compute the profit of the third chart, the ΔQ is -100 units, the starting P is 20% and the ending P is 43%, their average is therefore 31.5%. The expected profit found is

$$Profit_{Sub} = 790 + (-100) \cdot \{(6 - 5) - [(1 - 31.5\%) \cdot (6 - 1) - (5 - 1)]\}$$

$$= 790 - 100 \cdot [1 - (68.5\% \cdot 5 - 4)]$$

$$= 790 - 100 \cdot [1 - (-0.57)] = 790 - 157 = 633$$

The approximated value found is just 2€ away from the one obtained through simulation (635€) and it is due to the non-uniformity of the underlying distribution over the range chosen. Its accuracy is therefore dependent on the size of the mean shift (the smaller, the better) and on the distribution (the more regular over the interval, the most reliable the approximation).

The final point of interest is the distribution of the profit curve and, in particular, its potential asymmetry when the cumulative probability is either close to 0 or 100%. The critical fractile, derived from the economic parameters, emerges as a crucial factor in assessing the sensitivity of the profit curve to variations in the mean.
Below a summary of the behaviors analyzed in this section:

- Newsvendors with optimal quantities at low probabilities exhibit the lowest optimal profits and have significantly steeper profit curves on

their right tails compared to their left ones. Therefore they are more sensitive to mean shifts towards left - as shown in chart 1), where the profit reduces by 130€ for a 150 units decrease in mean - rather than towards right, as in chart 9), where the gap is limited to -76€ for +150 units.
However, due to the low benchmark optimal profit, both the shifts are very impactful in percentage and could result in significant losses from the expected.

- Newsvendors with very high optimal probabilities display the opposite behavior: they provide the highest profits and show greater sensitivity to shifts of the mean towards right. For instance, in chart 2), a €80 profit reduction results from a decrease of 150 units, whereas in chart 10) it loses €145 for an increase of 150 units.
These drops in profit are similar, but switched between the tails, if compared to the previous case; even though the expected profit is on a whole different order of magnitude.

- Probabilities near 50% exhibit a more symmetrical profit curve, where the impacts of shifts of the mean towards the left or right side are comparable in magnitude.

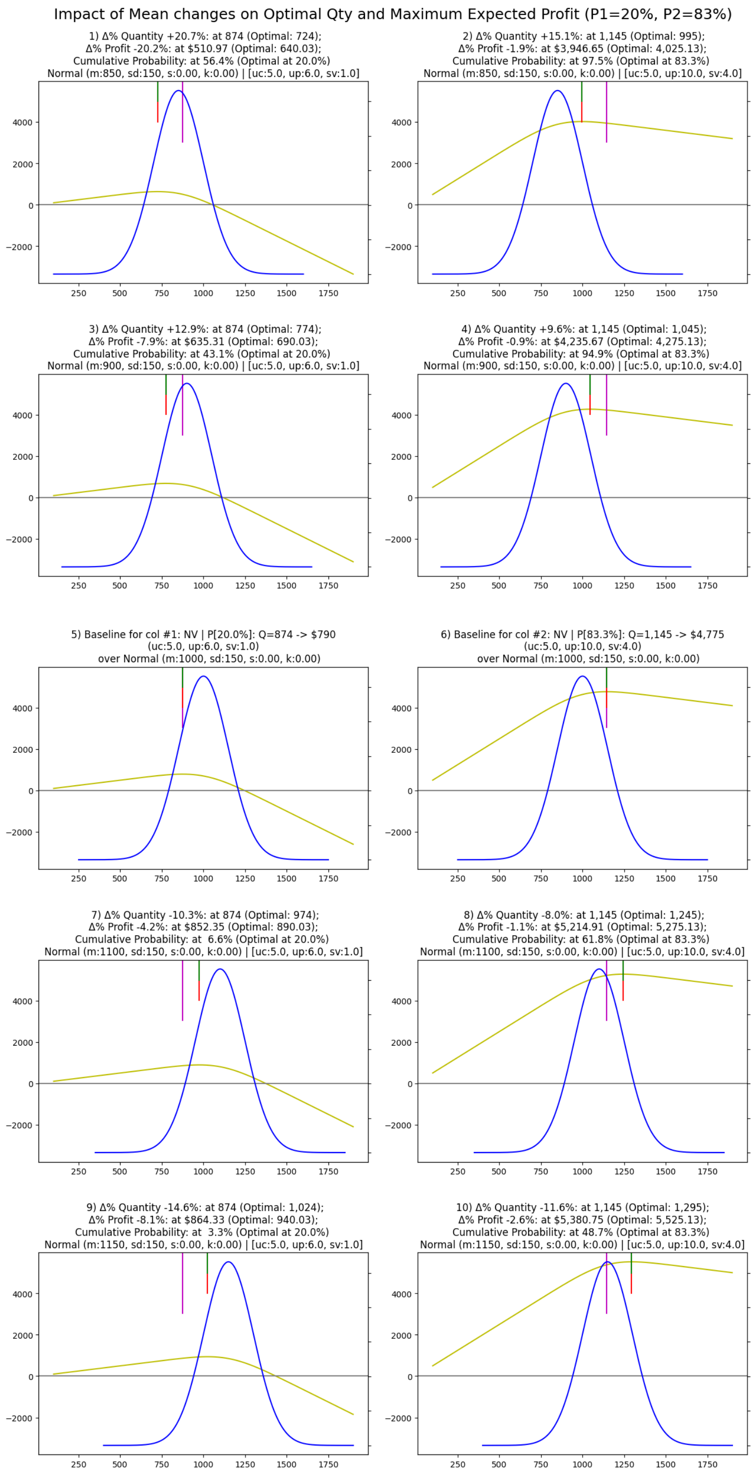

Figure 6.3: Impact of mean changes in case of a normal distribution on Newsvendors with extreme optimal probabilities

Chapter 7

Impact on optimal solution in case of pure standard deviation changes

Having analyzed the impact of various mean changes on different Newsvendors in the previous chapter, we now turn our attention to a different measure: the standard deviation. In fact, whereas the mean is a measure of central tendency, the standard deviation is a dispersion indicator, equally important in defining a distribution.
In this chapter, we will independently study, through a sensitivity analysis, the impact of the standard deviation on both optimal purchased quantity and profitability, keeping all other parameters constant.
This analysis holds significant business implications as well, summarized by the following question:

If I was expecting a distribution having a given variability (that can be measured as standard deviation), what would be the impact on the optimal quantity and the maximum expected profit if the real probability function has instead a different volatility?

Variability is a key risk indicator in many sectors (e.g., financial services) and it is very closely monitored and estimated.
It must be remembered that, during the thesis, the profit has never been computed as actual profit (the one associated to the effective purchased quantity and the realized actual demand) but only as the expected profit associated to the probability distribution.
Therefore, even before assessing the impact on the expected profit, it must be considered that a higher standard deviation increases the size of the interval in which the actual demand may fall into at a constant p-value.

A higher standard deviation is therefore detrimental by itself for the decision makers, increasing the level of risk associated to the model.

7.1 Changes in the standard deviation when optimal probability is near 50%

By changing only the standard deviation, the mean of the distribution remains the same for all the charts, at 1,000 units. In the case of a critical fractile returning a probability close to 50%, the optimal quantity occurs close to the mean. Therefore, we expect it to remain quite close to 1,000 units, slightly decreasing as the standard deviation increases. In figure 7.1, 6 different charts are portrayed having values of standard deviation from 250 units ($\frac{1}{4}$ of the mean) to 1 ($\frac{1}{1,000}$), with a 50 units step.

7.1.1 Expected profit behavior when standard deviation changes

Analyzing the profit behavior, at first glance, it is apparent that the profit curve gets spikier as the standard deviation decreases, becoming almost a triangle when the standard deviation reaches its minimum.
When the standard deviation is high, like in chart 1), the distribution is quite flat and presents a plateau in which the profit sensitivity is quite low, whereas in chart 6) the sensitivity is at its peak.
It is however very important to notice that between chart 1) and chart 6) there is a huge gap in the optimal profit of almost 700€.
The explanation is that, by increasing the standard deviation, the probability distribution becomes more sparse, flattening the profit curve and decreasing its maximum.
The behavior can be summarized by the following:

The expected profit associated to a quantity, in case of a distribution with lower standard deviation than the benchmark, is at any point better or equal to the profit of the baseline scenario.

Looking at the table 7.2, it is possible to observe this behavior more closely: fixing the sub-optimal quantity at 973 units, the profit associated monotonically increases as the standard deviation decreases, remaining constant after a certain value. This happens for values of standard deviation that make the cumulative distribution at the selected quantity being equal to 0%. At that point, the computation is totally deterministic and independent from the distribution: it becomes certain that all units purchased will be sold, and thus can be calculated as follows:

$$Profit = (up - uc) \cdot q = (8 - 5) \cdot 973 = 3€ \cdot 973 = 2,919€$$

The expected profit would be certain and identical to the actual final profit.

7.1.2 Optimal solution behavior when standard deviation changes

Also the behavior of the optimal solution is of great interest. As expected, the optimal quantity increases when the standard deviation decreases, passing from 955 units to 1,000 at standard deviations of 250 and 1 respectively. The optimal profit increases as well by a fixed amount of 13.74€ every 5 units (2.75€ per unit). This value is specific to this particular Newsvendor, but it offers a straightforward method to compute the optimal profit for any given standard deviation, when the distribution and the economic parameters remain constant.

As we can se from the figure 7.1, at higher values of standard deviation, the optimal quantity is function both of the distribution and of the economic parameters. However, as the standard deviation becomes smaller, the impact of the distribution expected value becomes gradually more significant, up to the point where the distribution is collapsed to the mean (when the standard deviation is close to 0).

The optimal profit in this scenario, both expected and actual, would be:

$$Profit = (up - uc) \cdot q = (8 - 5) \cdot 1,000 = 3€ \cdot 1,000 = 3,000€$$

From this starting point it is possible to compute the optimal expected profit for any value of standard deviation, using the formula

$$Profit_{sd=x} = Profit_{\sigma=0} - \text{coeff}_{Prof} \cdot \Delta\sigma = 3,000 - 2.75 \cdot \Delta\sigma$$

Simulating the original standard deviation of 150, the expected profit is:

$$Profit_{\sigma=150} = 3,000 - 2.75 \cdot 150 = 3,000 - 412.5 = 2,587.5€$$

exactly the value expected.

The same conclusions can be drawn for the optimal quantity as well. In fact, it also increases by a standard amount of 0.18 units for each unit of standard deviation lost.

7.1.3 Slope in the special case of no standard deviation

Once the decrease in the standard deviation reaches the limit scenario where it touches 0, the resulting Newsvendor represents the best possible case, becoming totally deterministic and independent from the distribution. In this scenario, the profit curve is composed by two incident lines that meet each other in the mean of the distribution (that is also the only possible demand). Each unit before this point is certain that will be sold, providing an increase in profit of $(up - uc) = (8 - 5) = 3$€ per unit.
Thus, the slope can be easily calculated as

$$m = \frac{\Delta Profit}{\Delta Qty} = \frac{3}{1} = 3$$

After the optimal point is reached, each additional quantity purchased will not be sold, resulting in an extra cost of $(uc - sv) = (5 - 1) = 4$€ per additional unit. In this case, the slope is

$$m = \frac{\Delta Profit}{\Delta Qty} = \frac{-4}{1} = -4$$

steeper than the previous because the optimal cumulative probability, by construction, is lower than 50%.
The quantity that makes the profit equal to 0€ is

$$q_0 = 1,000 + 3,000/(4) = 1,750 \text{ units}$$

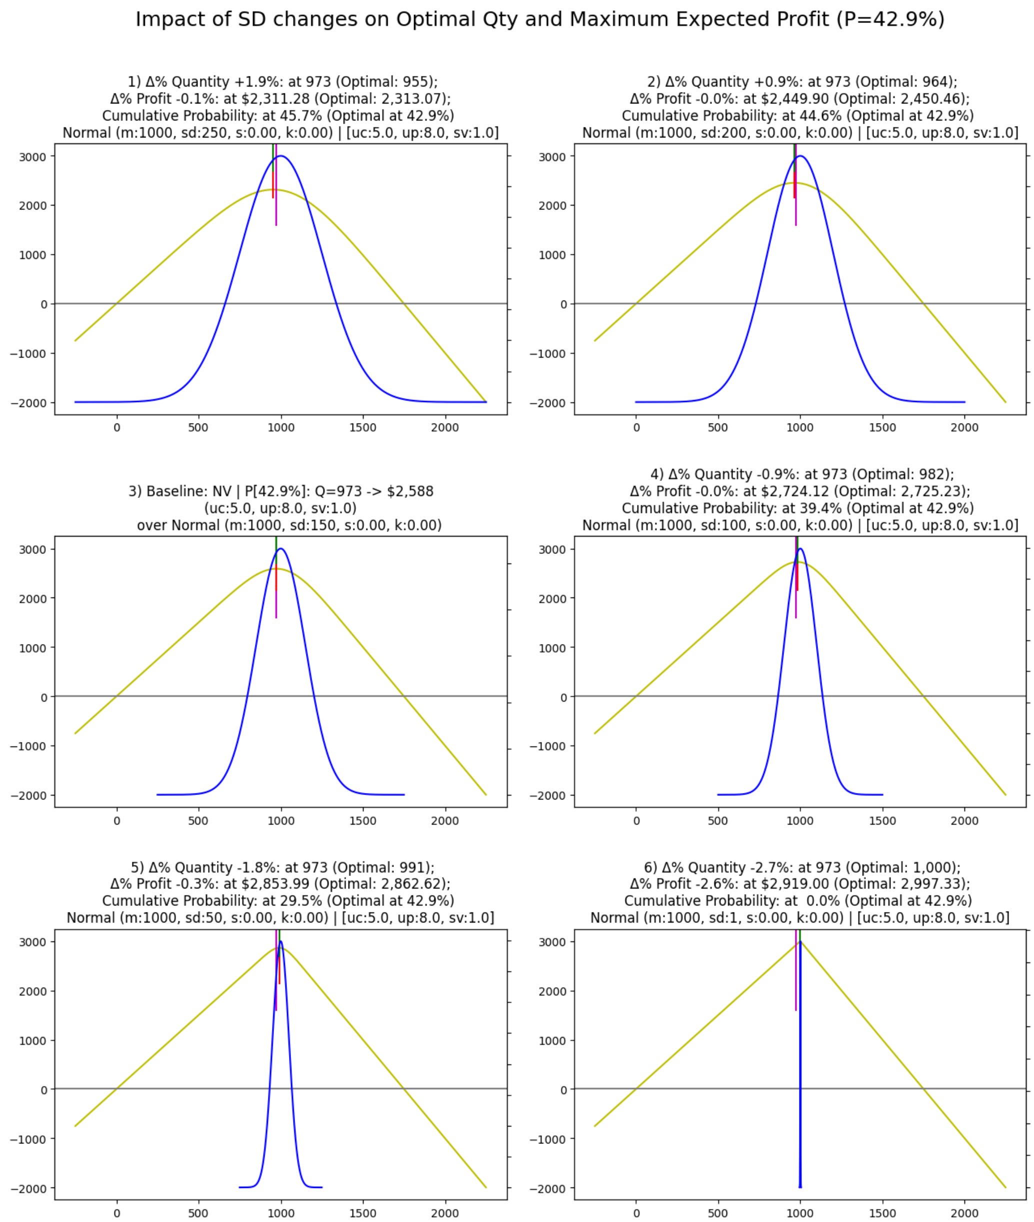

Figure 7.1: Impact of standard deviation changes in case of a Normal distribution at constant mean

sd	Opt. Qty	Opt. Profit	Opt. Prob.	SO Qty	SO Profit	SO Prob.	Δ Profit
250	955	€ 2,313.07	42.9%	973	€ 2,311.28	45.7%	€ 1.79
245	956	€ 2,326.81	42.9%	973	€ 2,325.16	45.6%	€ 1.65
240	957	€ 2,340.55	42.9%	973	€ 2,339.04	45.5%	€ 1.51
235	958	€ 2,354.29	42.9%	973	€ 2,352.91	45.4%	€ 1.37
230	959	€ 2,368.03	42.9%	973	€ 2,366.78	45.3%	€ 1.24
225	959	€ 2,381.76	42.9%	973	€ 2,380.65	45.2%	€ 1.12
220	960	€ 2,395.50	42.9%	973	€ 2,394.51	45.1%	€ 0.99
215	961	€ 2,409.24	42.9%	973	€ 2,408.36	45.0%	€ 0.88
210	962	€ 2,422.98	42.9%	973	€ 2,422.22	44.9%	€ 0.76
205	963	€ 2,436.72	42.9%	973	€ 2,436.06	44.8%	€ 0.66
200	964	€ 2,450.46	42.9%	973	€ 2,449.90	44.6%	€ 0.56
195	965	€ 2,464.20	42.9%	973	€ 2,463.73	44.5%	€ 0.46
190	966	€ 2,477.93	42.9%	973	€ 2,477.56	44.3%	€ 0.37
185	967	€ 2,491.67	42.9%	973	€ 2,491.38	44.2%	€ 0.29
180	968	€ 2,505.41	42.9%	973	€ 2,505.19	44.0%	€ 0.22
175	968	€ 2,519.15	42.9%	973	€ 2,518.99	43.9%	€ 0.16
170	969	€ 2,532.89	42.9%	973	€ 2,532.78	43.7%	€ 0.10
165	970	€ 2,546.63	42.9%	973	€ 2,546.57	43.5%	€ 0.06
160	971	€ 2,560.37	42.9%	973	€ 2,560.34	43.3%	€ 0.03
155	972	€ 2,574.10	42.9%	973	€ 2,574.10	43.1%	€ 0.01
150	973	€ 2,587.84	42.9%	973	€ 2,587.84	42.9%	€ -
145	974	€ 2,601.58	42.9%	973	€ 2,601.57	42.6%	€ 0.01
140	975	€ 2,615.32	42.9%	973	€ 2,615.29	42.4%	€ 0.03
135	976	€ 2,629.06	42.9%	973	€ 2,628.99	42.1%	€ 0.07
130	977	€ 2,642.80	42.9%	973	€ 2,642.66	41.8%	€ 0.13
125	977	€ 2,656.53	42.9%	973	€ 2,656.31	41.4%	€ 0.22
120	978	€ 2,670.27	42.9%	973	€ 2,669.94	41.1%	€ 0.33
115	979	€ 2,684.01	42.9%	973	€ 2,683.54	40.7%	€ 0.47
110	980	€ 2,697.75	42.9%	973	€ 2,697.11	40.3%	€ 0.64
105	981	€ 2,711.49	42.9%	973	€ 2,710.64	39.9%	€ 0.85
100	982	€ 2,725.23	42.9%	973	€ 2,724.12	39.4%	€ 1.11
95	983	€ 2,738.97	42.9%	973	€ 2,737.56	38.8%	€ 1.41
90	984	€ 2,752.70	42.9%	973	€ 2,750.94	38.2%	€ 1.76
85	985	€ 2,766.44	42.9%	973	€ 2,764.25	37.5%	€ 2.19
80	986	€ 2,780.18	42.9%	973	€ 2,777.49	36.8%	€ 2.69
75	986	€ 2,793.92	42.9%	973	€ 2,790.63	35.9%	€ 3.29
70	987	€ 2,807.66	42.9%	973	€ 2,803.66	35.0%	€ 4.00
65	988	€ 2,821.39	42.9%	973	€ 2,816.54	33.9%	€ 4.85
60	989	€ 2,835.13	42.9%	973	€ 2,829.26	32.6%	€ 5.87
55	990	€ 2,848.87	42.9%	973	€ 2,841.76	31.2%	€ 7.11
50	991	€ 2,862.62	42.9%	973	€ 2,853.99	29.5%	€ 8.62
45	992	€ 2,876.35	42.9%	973	€ 2,865.87	27.4%	€ 10.48
40	993	€ 2,890.09	42.9%	973	€ 2,877.27	25.0%	€ 12.81
35	994	€ 2,903.82	42.9%	973	€ 2,888.04	22.0%	€ 15.79
30	995	€ 2,917.56	42.9%	973	€ 2,897.91	18.4%	€ 19.65
25	995	€ 2,931.30	42.9%	973	€ 2,906.51	14.0%	€ 24.79
20	996	€ 2,945.03	42.9%	973	€ 2,913.28	8.9%	€ 31.76
15	997	€ 2,958.77	42.9%	973	€ 2,917.50	3.6%	€ 41.27
10	998	€ 2,972.51	42.9%	973	€ 2,918.92	0.3%	€ 53.59
5	999	€ 2,986.26	42.9%	973	€ 2,919.00	0.0%	€ 67.26
1	1,000	€ 2,997.18	42.9%	973	€ 2,919.00	0.0%	€ 78.18

Figure 7.2: Table that portrays the impact of standard deviation changes to better evaluate the sensitivity in case of a normal distribution and constant mean

7.2 Changes in the standard deviation in case of skewed Newsvendor

In the previous section, the optimal probability was slightly below 50%. We now proceed to analyze the behavior of a Newsvendor model with more extreme economic parameters, specifically examining the effects due to variations in the standard deviation. in figure 7.3 it is possible to analyze the behavior of two Newsvendor, NV1 and NV2, having optimal cumulative probabilities of respectively 20% and 83%.

For NV1, the profit associated to each unit sold is $(up-uc) = (6-5) = 1€$, the mean is fixed at 1,000 units and therefore the profit in case of no sd is $1€ \cdot 1,000 = 1000€$. The benchmark for the column is chart 5), set to a standard deviation of 150 units and a profit of 790€. With such information, it is possible to compute the coefficient of profit decrease per additional unit of standard deviation, like in the subsection 7.1.2, as

$$\text{coeff}_{Prof} = \frac{1,000 - 790}{150} = \frac{210}{150} = 1.4$$

This coefficient is coherent with all the scenarios based on the first Newsvendor.

The profit coefficient for NV2 is calculated in the same way and it corresponds to

$$\text{coeff}_{Prof} = \frac{(10-5) \cdot 1,000 - 4,775}{150} = \frac{5,000 - 4,775}{150} = \frac{225}{150} = 1.5$$

also in this case, perfectly coherent with all the charts related to NV2.

The two coefficients are similar and, therefore, so is their impact to the optimal profit decrease per additional unit of standard deviation. However, in percentage over their benchmark optimal profit, the difference is much more significant for NV1, because of its substantially lower reference expected profit. Once again, the negative effects of the variability have a much larger impact for the Newsvendor with lower critical fractile and profitability margin.

Such profit coefficient provides an exact solution in case of symmetrical distribution, where a change of standard deviation has the same impact in both tails, but can be calculated also for skewed ones (e.g., gamma and

exponential) remaining quite reliable over reasonable changes in standard deviations.
The quantities follows the same pattern: an increase of 0.84 units for NV1 and a decrease of 0.97 units for NV2 per unit of standard deviation lost Both values are significantly higher than the 0.18 observed in in the subsection 7.1.2. The substantially higher value can be easily explainable by the fact that all three distribution converge to the same optimal quantity (1,000 units) when the standard deviation is close to 0, but because the optimal quantity is farther away from the mean, the step needs to be higher. NV2, having its optimal quantity occurring at 83% of the distribution, is the farthest away from the mean and as a result needs the highest step.

In addition, all the other findings found in the previous sections are still valid here:

- The optimal quantity tends towards the mean of the distribution as the standard deviation decreases. For NV1, it increases from 790 to 1,000 units; in NV3 it starts from 1,242 units and decreases;
- Fixing a quantity, the expected profit associated always increases as the standard deviation decreases;
- Before the distribution begins and after it ends, the profit behaves as a line having a slope function of the economic parameters;
- As the standard deviation increases, the distance between the optimal and the sub-optimal profit keeps increasing, but at a significant slower pace if compared to an equivalent decrease in standard deviation. This is easily explainable by the shape of the distribution that, flattening, becomes less sensible to a change in quantity around an increasing interval around the optimal.

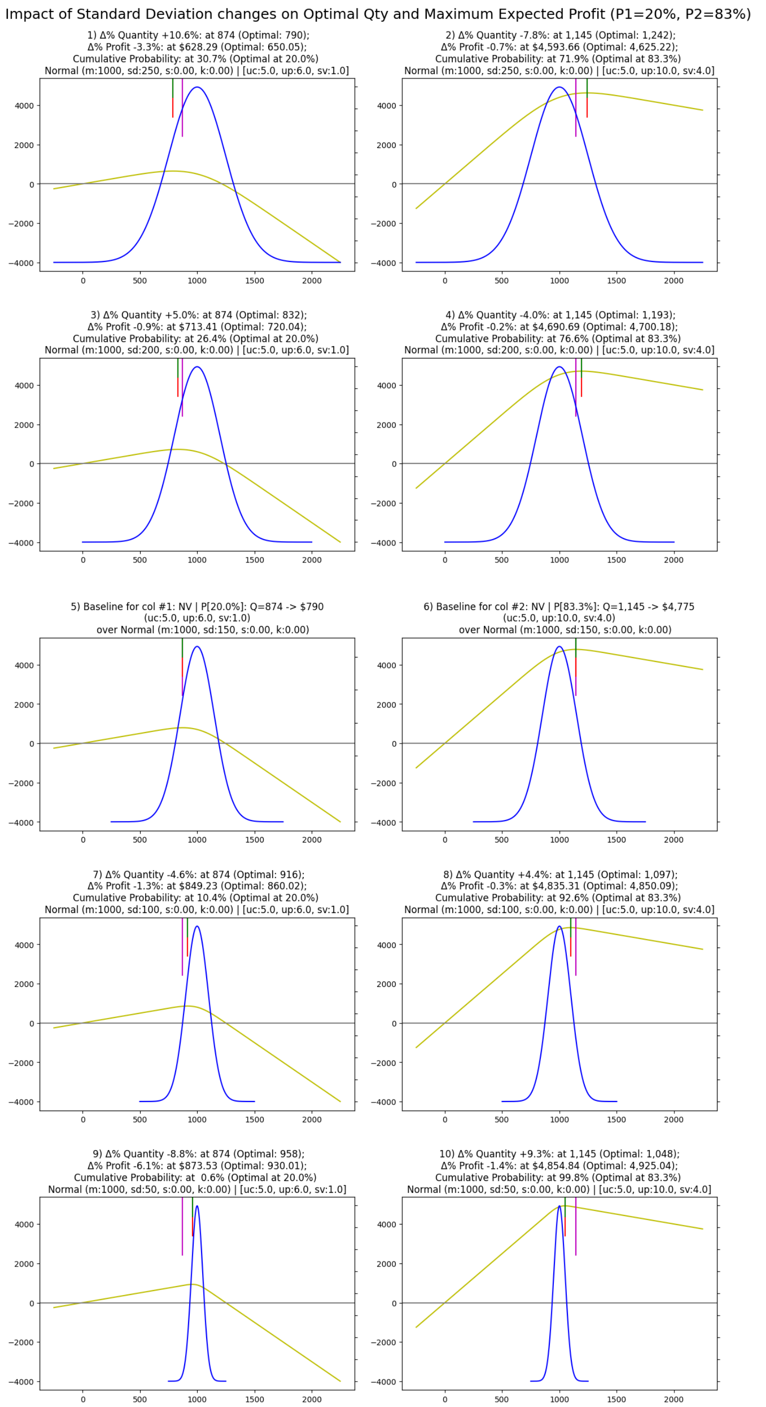

Figure 7.3: Impact of standard deviation changes in case of a normal distribution on Newsvendors with extreme optimal probabilities

Chapter 8

Impact on optimal solution in case of mean changes and constant coefficient of variation

After analyzing the impact of changes to the mean and standard deviation separately, this section aims to assess the effects of simultaneous changes to both parameters. The analysis involves varying the mean, as done in chapter 6, but instead of keeping the standard deviation fixed, the coefficient of variation is held constant at 15%.

8.1 Impact when the optimal probability is near 50%

As analyzed in chapter 7, if both the distribution remains a the same (only shifting its mean) and the economic parameters do not vary, the coefficient of optimal profit decrease per additional unit of standard deviation remains the same.
Calculating such coefficient from the charts of figure 8.1 it is possible to obtain:

$$\text{coeff}_{1)} = \frac{(8-5)\cdot 1,250 - 3,235}{188} = 2.75$$
$$\text{coeff}_{3)} = \frac{(8-5)\cdot 1,000 - 2,588}{150} = 2.75$$
$$\text{coeff}_{6)} = \frac{(8-5)\cdot 850 - 2,200}{128} = 2.75$$

As expected, such value is the same found in subsection 7.1.2 and the same applies to the coefficient of optimal quantity shift per additional unit of standard deviation, 0.18.

Therefore, in case of normal distribution, it is possible to exactly compute analytically the optimal expected profit, by combining the effects of the two different changes:

$$impact_{mean} = \Delta Q \cdot (up - uc)$$

$$impact_{sd} = -CV \cdot \Delta Q \cdot \text{coeff}_{Prof}$$

where CV is the coefficient of variation. The combined impact can be calculated as:

$$impact_{both} = \Delta Q \cdot [(up - uc) - CV \cdot \text{coeff}_{Prof}]$$

that, applied to the scenario of chart 6), enables the calculation of the optimal profit from the baseline as:

$$Profit_{6)} = Profit_{3)} + \Delta Q \cdot [(up - uc) - CV \cdot \text{coeff}_{Prof}] =$$

$$= 2,588 + 250 \cdot [(8 - 5) - 2.75 \cdot 0.15] =$$

$$= 2,588 + 250 \cdot [3 - 0.41] = 2,588 + 647 = 3,235€$$

which is the same result found through simulation.

Also the optimal quantity can be found analytically :

$$impact_{mean} = \Delta Q$$

$$impact_{sd} = -CV * \Delta Q \cdot \text{coeff}_{Qty}$$

resulting in a combined effect of

$$impact_{both} = \Delta Q \cdot (1 - CV \cdot \text{coeff}_{Qty})$$

Applying once again this formula to chart 6), it is possible to compute analytically the new optimal quantity as:

$$Qty_{6)} = Qty_{3)} + \Delta Q \cdot (1 - CV \cdot \text{coeff}_{Qty}) =$$

$$= 973 + 250 \cdot (1 - 0.15 \cdot 0.18) = 973 + 250 \cdot 0.973 = 973 + 243 = 1,216$$

that, again, perfectly matches the value found in the simulation.
In case of non-symmetrical distribution, the results found would be an approximation of the actual quantity and profit.

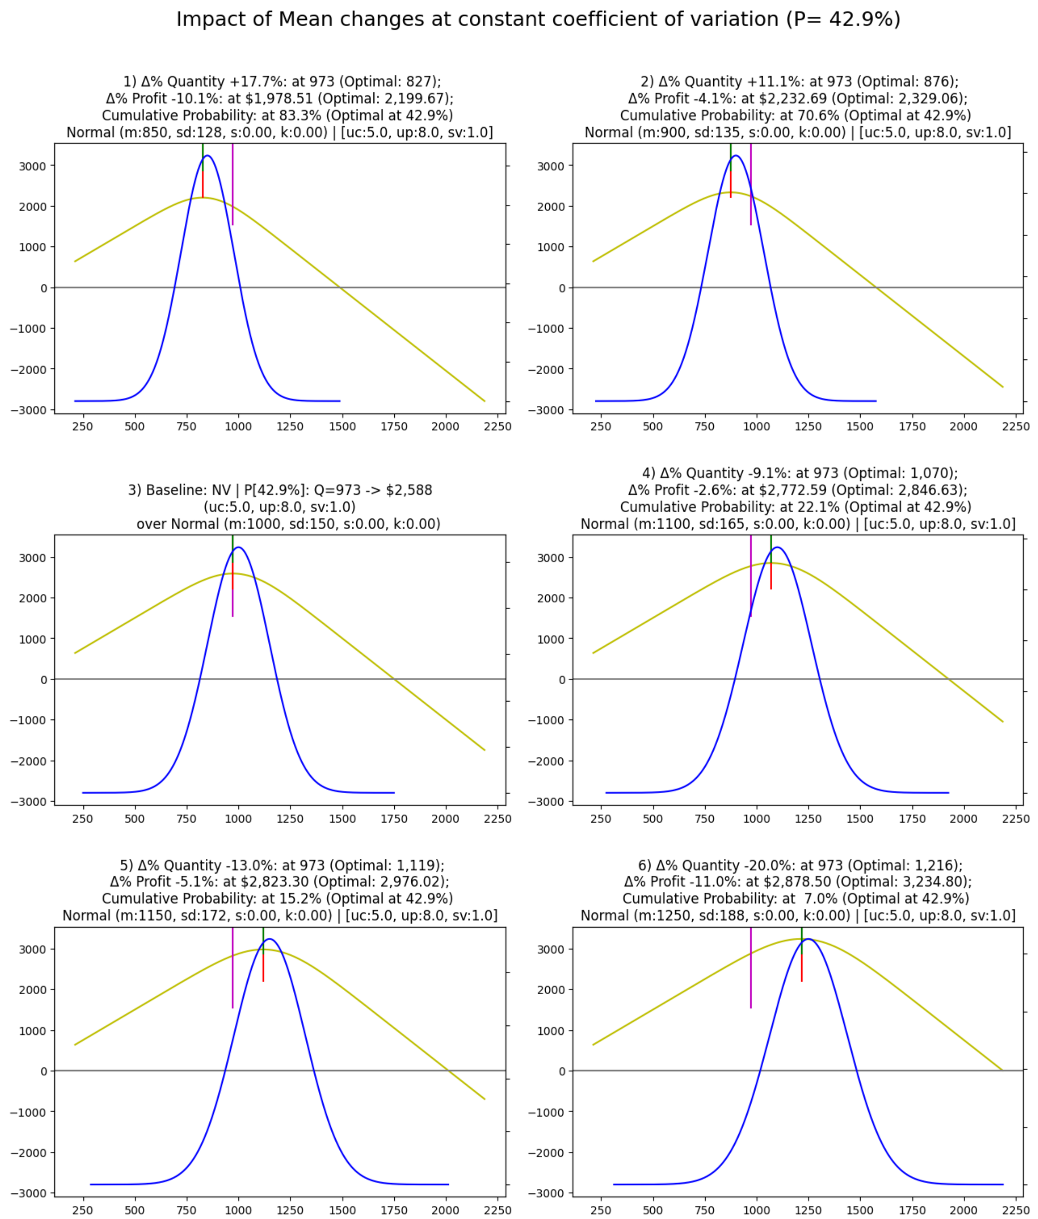

Figure 8.1: Impact of mean changes in case of a Normal distribution keeping the coefficient of variation constant

In conclusion, the consolidated impact of the two changes is additive and the impact of the mean is significantly higher due to the damping factor of the coefficient of variation: considering a CV of 15%, a 100 units increase in the mean corresponds to only a +15 units change in standard deviation. Even though additive, the two effects are opposite in respect to their impact on the expected profit: an increase in the mean has a positive effect, whereas the linked rise in standard deviation is detrimental.

8.2 Considerations in case of skewed Newsvendor

Comparing the simulations in figure 8.2 to those in figure 6.3 on page 61, analyzed when assessing the impact of the pure change in mean, it is possible to observe two distinct behaviors for the optimal:

- a higher value of optimal quantity and profit for charts 1) to 4);
- a lower value of optimal quantity and profit for charts 7) to 10).

This is consistent with the findings of the previous section. Specifically, the lower standard deviation partially offsets the negative effect of the mean decrease in the first four charts; conversely, the opposite occurs in the last four. The mean continues to have a more significant impact than the standard deviation, as evidenced by the increase in both the optimal quantity and the optimal expected profit across the rows in figure 8.2.
All previous considerations and calculations from other sections of this chapter remain applicable here.

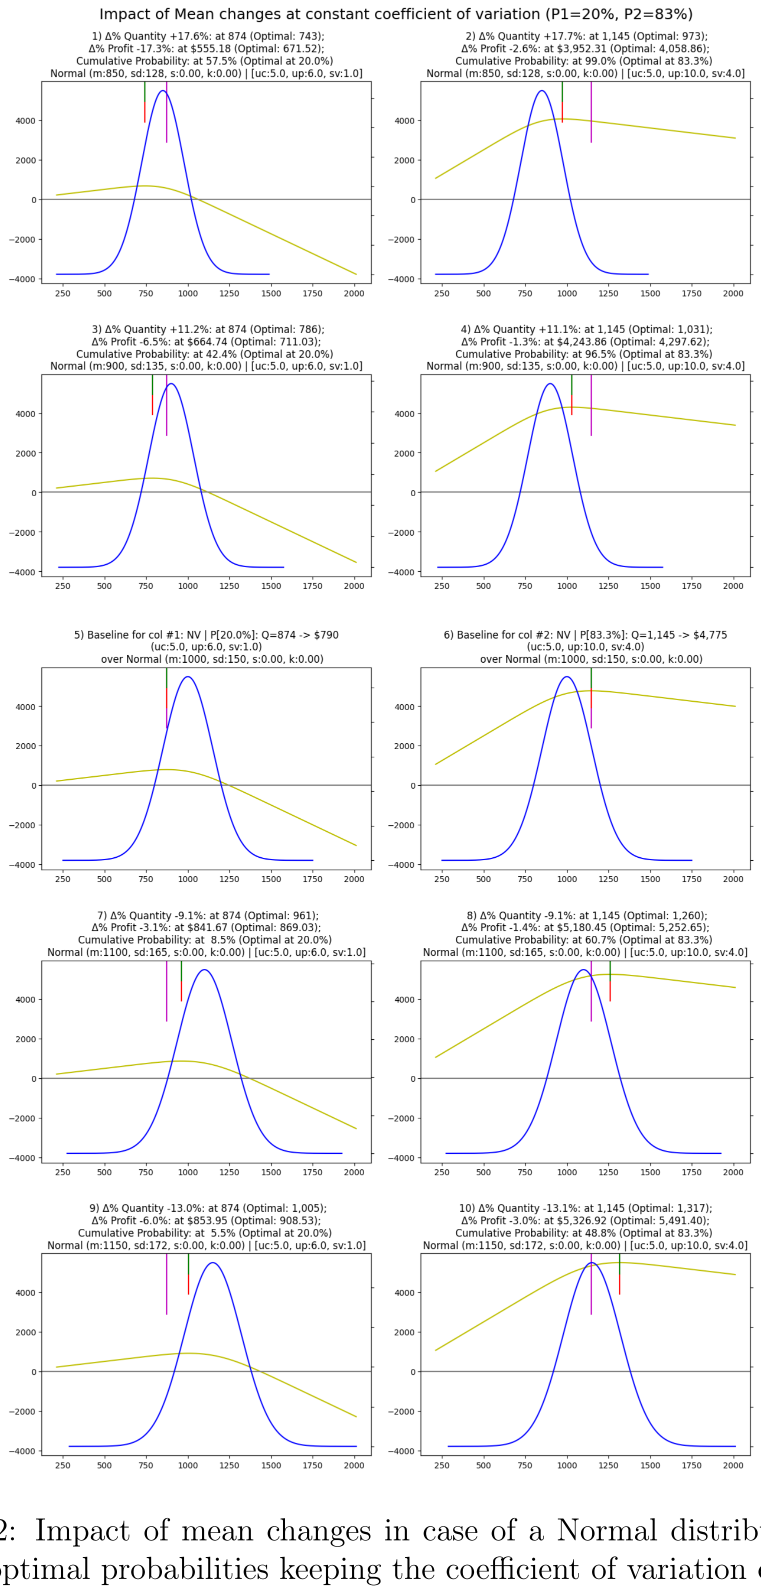

Figure 8.2: Impact of mean changes in case of a Normal distribution with extreme optimal probabilities keeping the coefficient of variation constant

Chapter 9

Impact on optimal solution in case of skewness change, based on known distributions

Unlike previous sensitivity analyses where it was possible to change one parameter at a time, assessing skewness necessitates varying the kurtosis as well. Therefore, the analysis will be conducted using three different distributions:

- **Exponential:** having a skewness of 2 (or -2, if reversed), it is the extreme case;
- **Normal:** perfectly symmetrical distribution having skewness equal to 0, it is the central case, and benchmarking basis of the analysis;
- **Gamma:** bridge between the distributions above, it allows skewness between 0 and 2, acting as a generalized version that encompasses both the exponential and normal distributions at its respective extremes. The skewness is function of the parameter α, proper of the distribution, and can be calculated as

$$\gamma = \frac{2}{\sqrt{\alpha}}$$

9.1 Impact of skewness when the optimal probability is near 50%

Figure 9.1 illustrates the complete range of skewness from 2 to -2, with increments of 0.5. Both the mean and standard deviation are held constant

at 1,000 and 200 units, respectively, to ensure full comparability between all the charts.
The first notable observation is that, when the optimal quantity is fixed at 964, the associated cumulative probability is 56% with a skewness of 2. This probability consistently decreases to 43% in the case of the uniform distribution (the benchmark of this analysis) and continues to decrease for negative skewness values, reaching its minimum of less than 31% with the reversed exponential distribution.
Opposite behavior for the optimal quantity, which increases from 912 to 1,031, passing for the benchmark of 964.
However, both the expected optimal and sub-optimal profit do not reflect this consistent behavior. In fact, the minimum expected profits does not occur neither in the exponential distributions nor in the normal one, but in the gamma plotted in chart 8).
At first glance, distributions with a higher $|\gamma|$ (absolute value of the skewness) show a sharper peak at their optimum, whereas for the normal distribution it is flatter. Both the exponential distributions have an optimal profit higher than the baseline case, but for different reasons:

- **Chart 1):** Up to a quantity of 800, the profit increases at the optimal rate of 3€, as the cumulative probability up to that point is equal to 0%. Because the optimal point is so close to the peak of the distribution, this is the best scenario.

- **Chart 2):** The distribution is reversed and the least effective side is on the left. However, it is still more concentrated than the normal distribution, even in case of constant standard deviation, as explained in chapter 10. In fact, at a quantity of 964 units the cumulative probability is 30%, while around 1,200 it is 100%. This means that 70% of the distribution is concentrated in just about 250 units. By comparison, in a normal distribution 68% of the probability is included in the range $[\mu - \sigma; \mu + \sigma]$, which in this case is equal to 400 units.

Chart 1) has the best optimal profit, benefiting from both the positive effect of the steeper profit curve, typical of the negative skewness, and the higher concentration of probability.
Conversely, chart 8) has the lowest expected profit, because the positive contribution of the higher concentration is not enough to offset the negative effect of the negative skewness.
However, the effects of skewness seem to be very limited across all the charts analyzed.

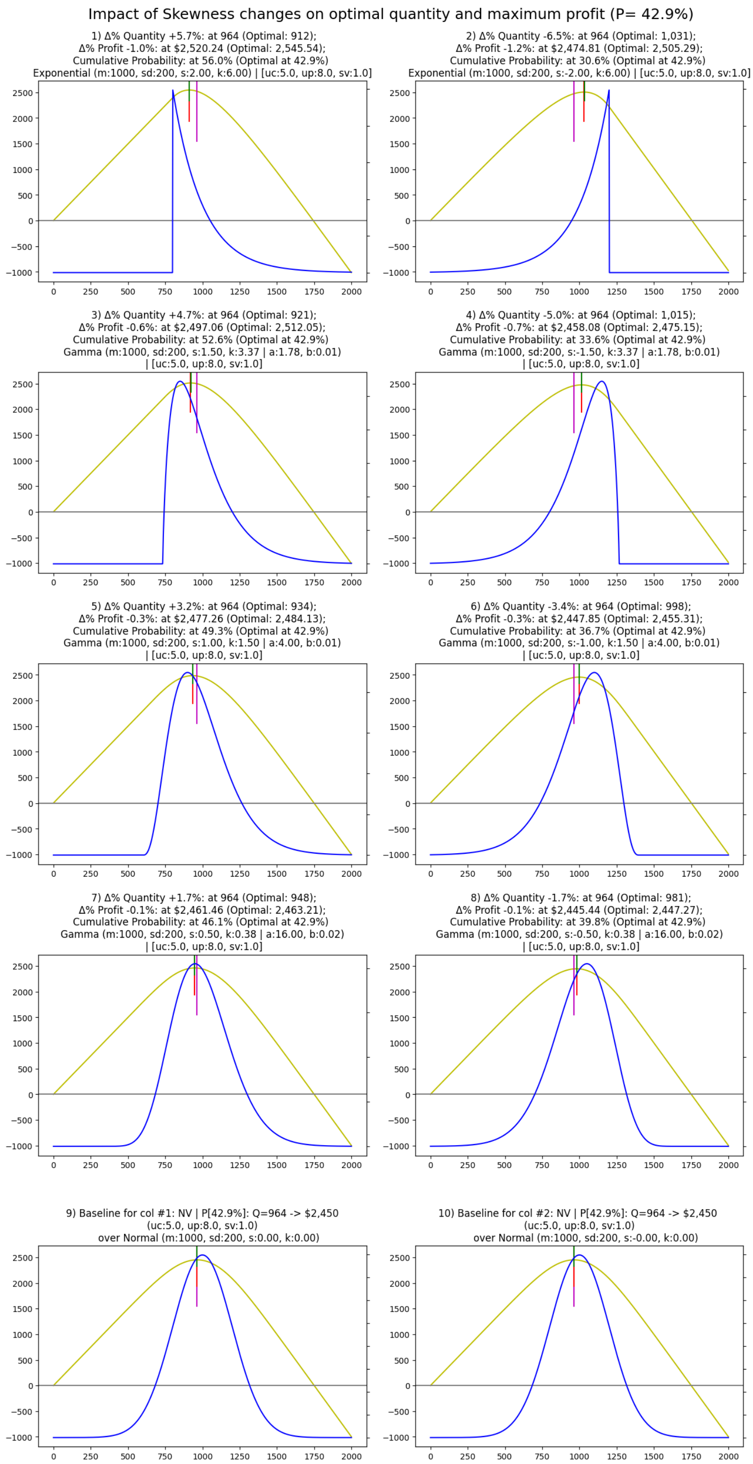

Figure 9.1: Impact of skewness changes in case of a normal distribution and constant mean and standard deviation

9.2 Impact of skewness in case of skewed Newsvendor

In the first column of figure 9.2, where the economic parameters result in a critical fractile of 20%, the behavior is not consistent with the findings of section 9.1: the minimum quantity occurs on the second row, whereas the optimal profit is constantly decreasing.
However, also in this case, the differences between optimal and sub-optimal profit within a chart are quite minimal. The best profits are associated to the first chart which is extremely condensed in the beginning of the distribution, where both the optimal and sub-optimal quantities fall. The two quantities are really close to each other and to the peak of the distribution (i.e., the left limit). As a result, the profit curve reaches the optimal point almost following a line with slope equal to

$$m = (up - uc) = (6 - 5) = 1$$

In fact, the profit is only slightly lower than the quantity (both for optimal and sub-optimal), as the profit obtained is almost deterministic due to the low critical fractile and significant positive skewness.

On the other hand, the second column showcases a very different the scenario: The best optimal profit can be found in chart 10), where the skewness is equal to -2; whilst the worst occurs in chart 2), having a skewness of 2. The profit consistently decreases as the skewness does the same.
Conversely, the maximum optimal quantity is set in chart 6) (the baseline, where the skewness is 0) and decreases as the skewness moves in both directions. This is logical, given the very high critical fractile and the fact that the normal is the only distribution in figure 9.2 having infinite tails.
The optimal profit can be found in chart 10). Because the optimal quantity occurs in the end of the distribution, the most concentrated part of it, the average rate of increase is the highest among all the charts; an opposite behavior if compared to the first column, due to the fact that the two optimal quantities lay on different sides of the mean.
Again, the distance between optimal and sub-optimal profit is minimal within a chart and also the difference in sub-optimal profits between charts is almost negligible.

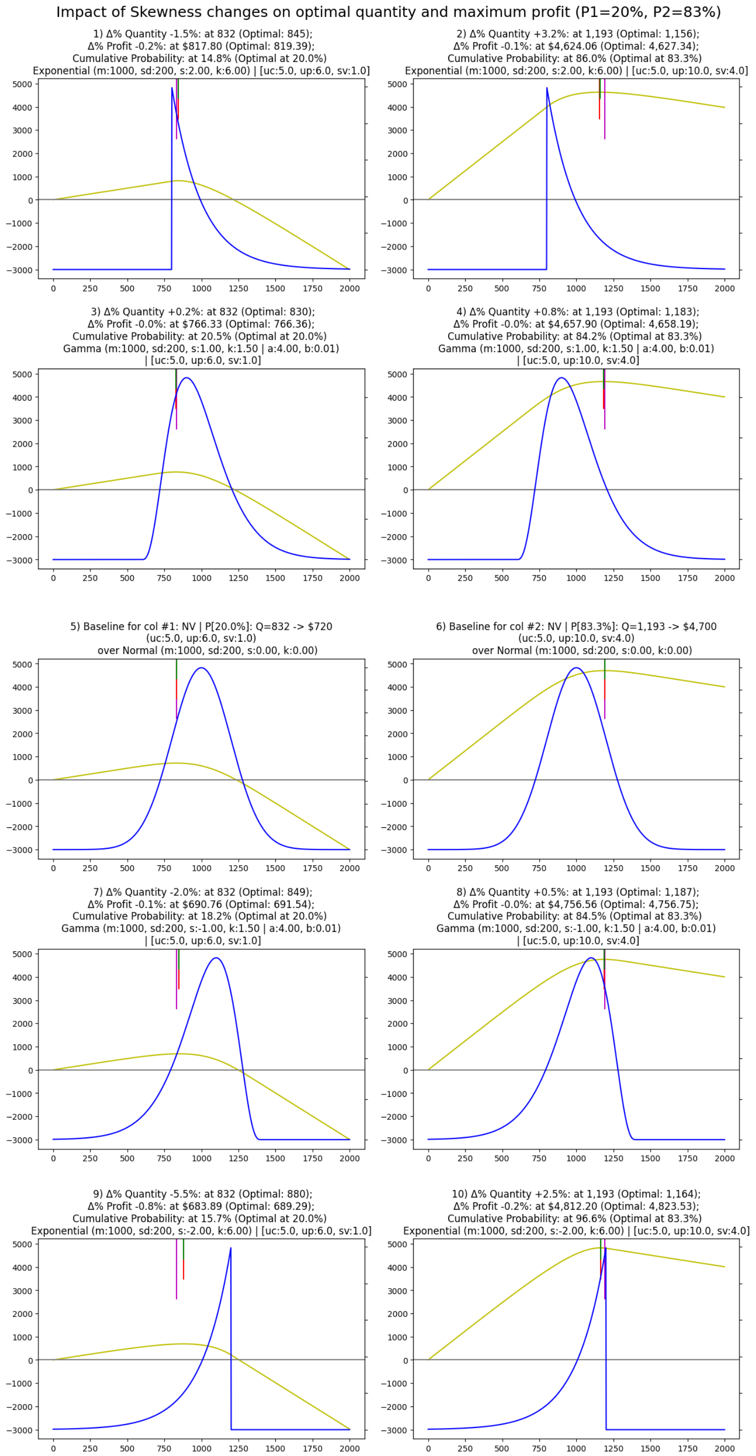

Figure 9.2: Impact of skewness changes in case of a normal distribution on Newsvendors with extreme optimal probabilities

From a decision-making perspective, the difference between optimal and sub-optimal profits is minimal across all scenarios. Therefore, in the absence of sufficient data to reliably estimate the actual demand distribution, it is advisable to make estimates based on a normal distribution.
In summary, Newsvendors with low critical fractile typically yield slightly better profits when the distribution is positively skewed, as in chart 1) of figure 9.2. Conversely, those with critical fractile close to 1 tend to perform best with negative skewness, as shown in chart 10).
Surprisingly, Newsvendors having optimal quantity close to the median tend to perform better when the distribution is skewed as well, regardless of the side. However, this behavior is not caused by the skewness per se, but by the other factors that are introduced when changing the distribution. In fact, such behavior will be explained on chapter 10, studying the impact of kurtosis changes on the model. The impact on the expected profit remains minimal in all scenarios.

To summarize, the impact of skewness is the least significant among those analyzed so far, favoring Newsvendors that have optimal quantity in proximity to the mode of the demand distribution, but not providing a solid general rule due to the impossibility to only vary it without changing other factors as well.

Chapter 10

Impact on optimal solution in case of pure kurtosis change

To study the effects of changes in kurtosis, it is possible to properly perform a sensitivity analysis without impacting any additional measure, as done for both mean and standard deviation in the previous chapters.
Using a Beta distribution and varying its parameters α and β, it is possible to obtain a large portfolio of distributions: forcing both parameters to the same value, the resulting curve will be a symmetrical distribution having its kurtosis as a function of α and β

$$\kappa = \frac{6\left[(\alpha-\beta)^2(\alpha+\beta+1)-\alpha\beta(\alpha+\beta+2)\right]}{\alpha\beta(\alpha+\beta+2)(\alpha+\beta+3)}$$

given that $\alpha = \beta$, this simplifies to

$$\kappa = \frac{-6}{2\alpha+3}$$

Using this workaround, the impact of kurtosis can be studied independently from the other measures and its effects assessed separately.

10.1 Impact of kurtosis when the optimal probability is near 50%

In figure 10.1 it is possible to observe the impact on the Newsvendor model when the value of the kurtosis is changed from 0 (normal distribution) to -1.2 (uniform) until reaching the extreme case of -2.
Analyzing the first four charts, it is obvious that the distributions becomes

more and more ‘squared’, converging into a uniform distribution.
The optimal quantity decreases, starting from 964 units in the normal case, to 951 in the uniform one. The impact of the distribution slowly gets less significant as it turns into a uniform and the optimal quantity is shifted towards the tails. Because the critical fractile is $<$50% in all the charts of figure 10.1, the optimal quantity in this case shifts to the left side.
The impact on the sub-optimal profit is negligible: the difference between optimal and sub-optimal profit in the uniform case is just 1€ out of more than 2,400€.

The last two charts are very extremes situations, unlikely to occur in real-world scenarios, in which the probability is mainly concentrated in the tails.
In chart 6) three different slopes can be identified:

$$m_1 = (100\% - 0\%) \cdot (8 - 1) - (5 - 1) = 8 - 5 + 1 - 1 = +3\text{€ per unit}$$

$$m_2 = (100\% - 50\%) \cdot (8 - 1) - (5 - 1) = 3.5 - 5 = -1.5\text{€ per unit}$$

$$m_3 = (100\% - 100\%) \cdot (8 - 1) - (5 - 1) = 0 - 4 = -4\text{€ per unit}$$

and the optimal quantity can be only in one of the two spikes. All Newsvendors with optimal probability $<$50% would have optimal quantity at 798 and an optimal profit of

$$Profit_{Opt} = Qty_{Opt} \cdot (up - uc) = 798 \cdot 3 = 2,394 \text{ units}$$

the best profit possible for that given quantity in any scenario, but having a very unfavorable sub-optimal profit when the purchased quantity is farther away.

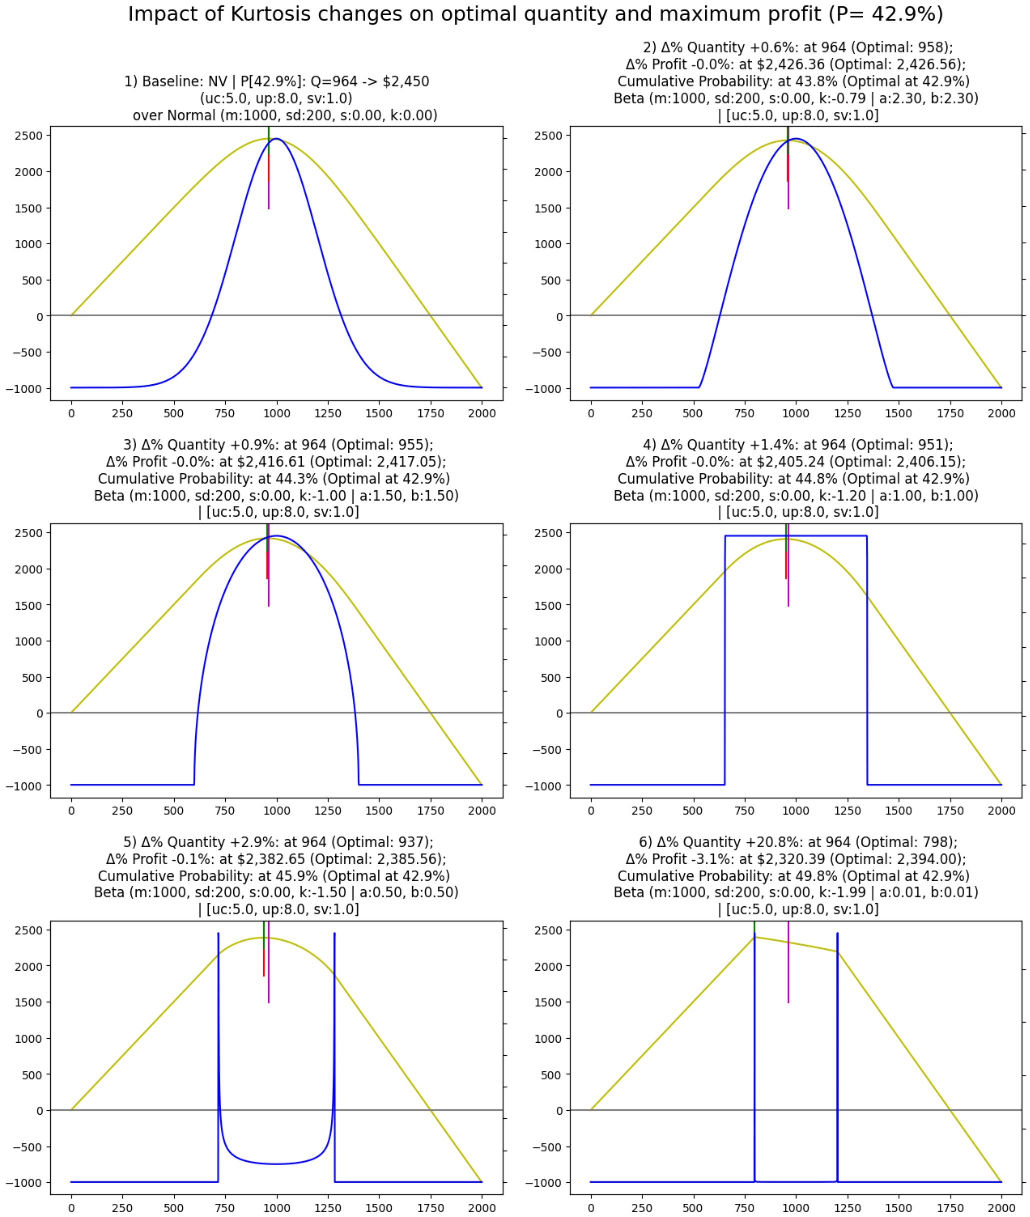

Figure 10.1: Impact of kurtosis changes in case of a normal distribution and constant mean and standard deviation

10.2 Impact of kurtosis in case of skewed Newsvendor

In this section it is analyzed whether the model behaves differently when the critical fractile is either close to 0 or 1. Focusing the attention to the first four charts on the left column of figure 10.2, the behavior of the optimal quantity is identical to what observed in section 10.1: it decreases as the kurtosis decreases. Conversely, the second column showcases the opposite behavior: the quantity is increasing, coherently with the expectation of its shift towards the external of the distribution.

The profit however is not easy to be disclosed. In fact, the highest optimal profit for the left column can be found in chart 7) and not in chart 1). The optimal expected profit in chart 7) increases at the maximum rate up the lower bound of the distribution. Consequently, if the critical fractile is close to 0 the profit is almost deterministic. However, for slightly higher values of the critical fractile, the flat distribution quickly reduces the contribution of the unit price, making the profit to increase slower. After a threshold of the critical fractile, the less disperse distribution having higher kurtosis becomes more advantageous, as in section 10.1, where the optimal quantity falls at a cumulative probability of 43% and the best distribution for the profit was the normal.

The right column shows instead a steady behavior of profit increase, starting from 4,700€ for the normal distribution up reaching 4,711€ in chart 8) with the uniform. A flatter distribution, at constant critical fractile, makes the optimal quantity occur at higher quantities than the benchmark distribution having higher kurtosis. After a given critical fractile, were such positive effect totally offsets the profit gains obtained in the left side of the curve, the distribution having lower kurtosis provides a slightly better profit.

In fact, a decrease in kurtosis makes the extremes of the demand distribution closer, but significantly flattening the curve between those. This leads to different behaviors:

- In case of an extremely low optimal cumulative probability, it results in an increase of the optimal quantity, which falls close to the portion of the left tail that is cut-off. The decrease in kurtosis makes the sub-optimal quantity end up outside of the distribution, on its left. The result is best possible deterministic profit and an optimal quantity occurring just after the starting point of the distribution. In this range, the distribution having the lowest kurtosis shows better performance, as observed in the left column of figure 10.2

- Higher critical fractiles (but still lower than 50%) occur at higher quantities for the distributions with higher kurtosis. This results in a decrease in the optimal quantity in case of the distributions with low kurtosis. After the previous gains in profit are offset by the new contribution, the flatter distribution results to be less profitable (as in figure 10.1)

- The remaining 50% showcases a similar but opposite behavior: The flatter distribution increases the value at which optimal quantities occur, providing a better contribution to the expected profit. After a given threshold, the offset is complete and the expected profits start to become better again in case of distribution with lower kurtosis, as observed in the right column of figure 10.2.
  However, the very last bit of the distribution has an even different behavior. After the uniform distribution ends, the normal one has not still reached 100%, having still a tail. This reverses once again the effect of the kurtosis, leading to higher quantities in the case of a more tailed distribution. Even though the effect is positive also for the expected profit contribution, such a small remaining portion of the distribution is not enough to compensate for the gains up to that point. Thus, the profit would likely remain slightly better in the case of the distribution with lower kurtosis, but with an even lower gap.

The Beta distribution does not cover, for its symmetrical instances, the case of positive kurtosis, which instead can be found in other symmetrical distributions like the logistic or the Laplace. However, these cases can be derived by generalizing the behavior portrayed in the analyzed range of kurtosis [-1.2, 0].

In conclusion, also in the case of an extreme critical fractile, the impact of the kurtosis is very limited, both in terms of ideal optimal expected profit (the optimal profits are very close between them) and regarding the gap between optimal and sub-optimal profit (it is below 1% in all scenarios and close to 0% in most), making its complex impact almost negligible when trying to assess ex-ante the risk level of a Newsvendor.

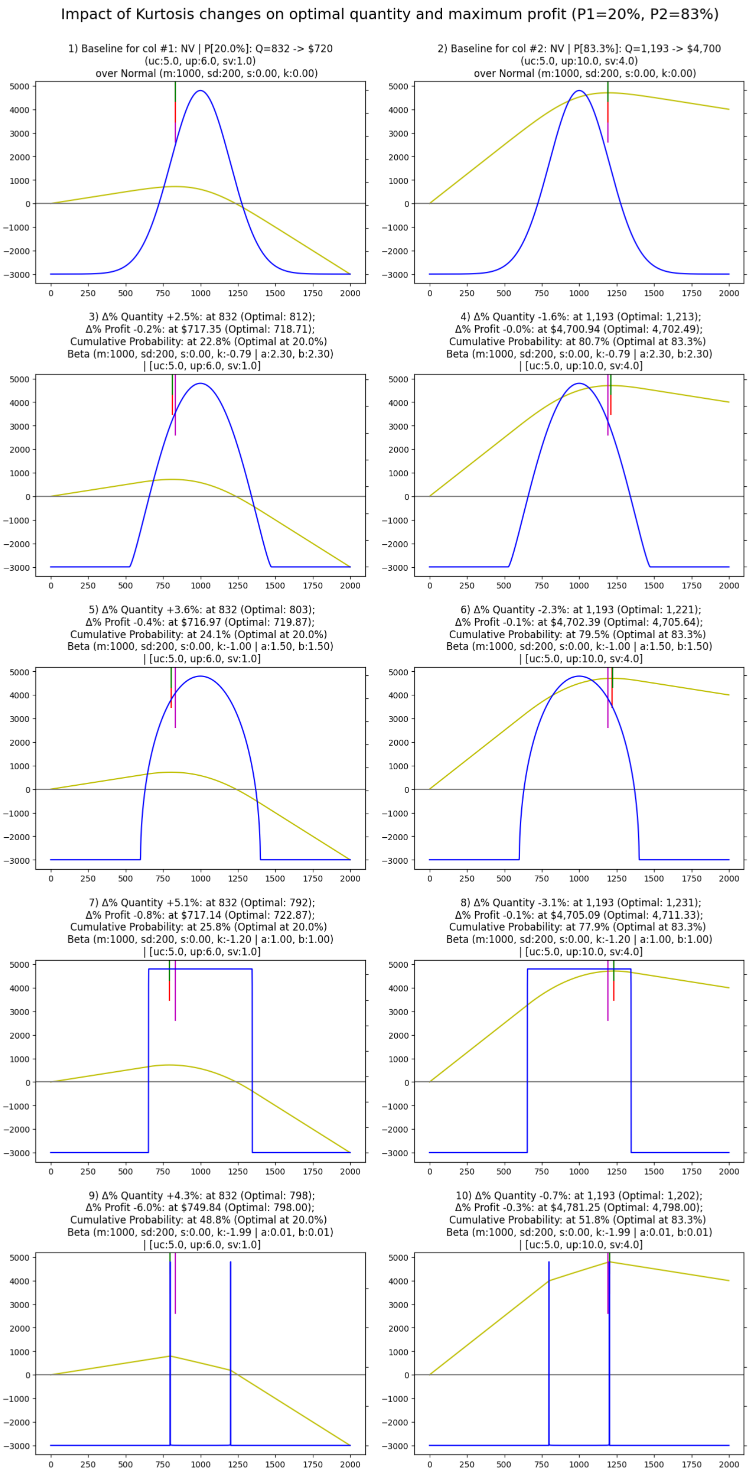

Figure 10.2: Impact of kurtosis changes in case of a normal distribution on Newsvendors with extreme optimal probabilities

10.3 Additional considerations regarding the skewness

This analysis adds material of discussion on chapter 9, where the skewness was not studied by itself, but the change in distribution also forced the kurtosis to different values.
The reason why the most skewed distribution were also the most concentrated, can be explained by the very high kurtosis associated to the exponential distribution ($\kappa = 6$) and therefore all the effects of the skewness were merged to those of the kurtosis.
The key takeaway is that while both effects are somewhat significant for the final outcome of the model, they are challenging to assess analytically and to evaluate independently.
In addition, the impact of skewness and kurtosis is significantly minor compared to that of the mean and standard deviation. Therefore, accurate estimation of the mean and standard deviation should remain the primary focus for any decision maker in the Newsvendor business.

Chapter 11

Impact on optimal solution in case of economic parameters change

After having measured the impact of the distribution on optimal quantity and profitability, the focus of the analysis has to be shifted on the other possible source of variability: the economic parameters.

These parameters are crucial as they determine the value of the critical fractile, which in turn is essential for estimating the optimal quantity and the resulting expected profit.

In figure 11.1 it is introduced the benchmark that will be used to assess the impact of a change of 1€ in each of the economic parameters. The baseline model has unit cost of 5€, unit price of 8€ and resale value of 1€. The distribution chosen is a Gamma having skewness of 1.6, in order to make the findings as general as possible, The mean and the standard deviation are respectively 1,000 and 200 units.

The benchmark critical fractile, key of this analysis, is fixed to 42.9%. The related optimal quantity and profit are respectively 919 units and 2,518€, as shown in figure 11.1.

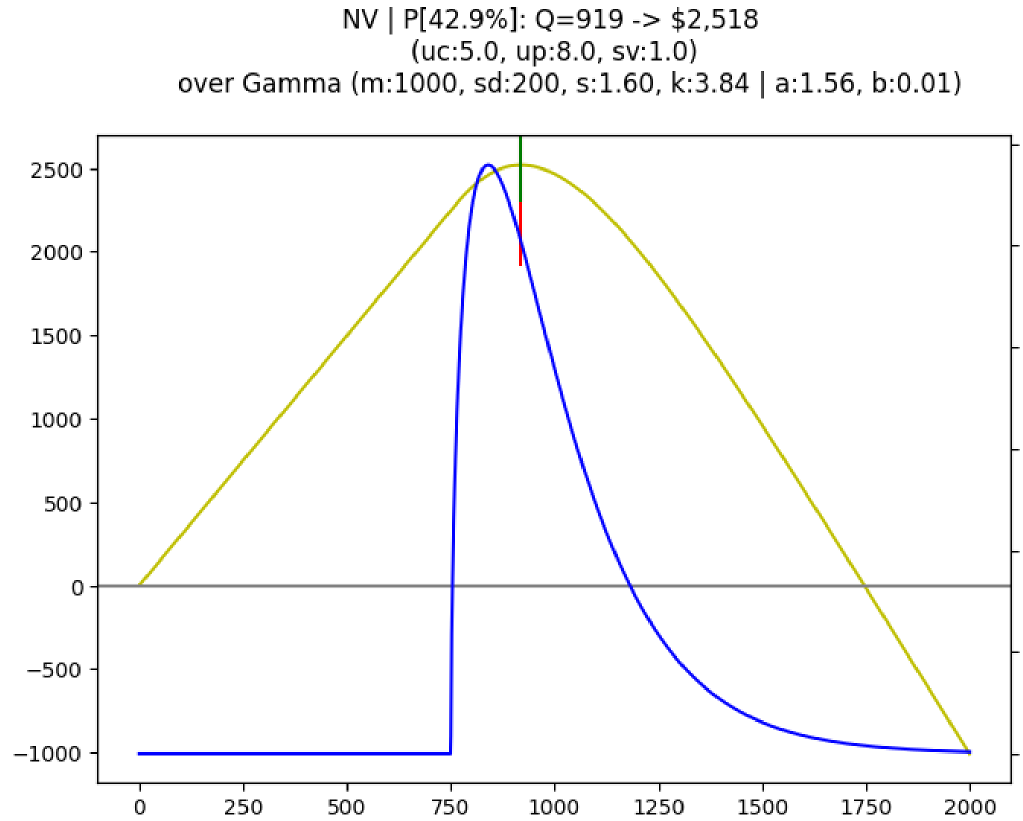

Figure 11.1: Newsvendor distribution used as benchmark for the sensitivity analysis on economic parameters

11.1 Empirical evidences on the variation of the economic parameters by a fixed amount

The purpose of this section is to find a satisfactory answer to an extremely important business question regarding the Newsvendor:

If I could focus my resources to either increase by 1€ the unit price / resale value or decrease by the same amount the unit cost, where should I focus my effort?

The first parameter to be the focus of the analysis is the unit price: as shown in the first row of figure 11.2, the impact of a change in unit price from 7€ to 9€, passing through the benchmark of 8€, results in a significant increase in the critical fractile. The change in optimal quantity is not as substantial, ranging from a minimum of 883 to a maximum of 949 units. However, the profit changes significantly, more than doubling from the worst to the best case. The range for the optimal probability is [33.3%, 50%], -9.6 percentage points (pp) and +7.1pp, respectively, from the benchmark.

The next important factor is the unit cost, shown in the second row of figure 11.2. At first glance, its impact is even greater than the one related to the unit price, both in terms of critical fractile and profit.
From a baseline of 5€, the charts show the impact of both decreasing it by 1€ to 4€ and increasing it by the same amount to 6€.
The changes in the optimal quantity are more significant than what was possible to observe in the previous paragraph with the unit price. This behavior can be explained by the higher probability range [28.6%, 57.1%], -14.2pp and +14.2pp respectively from the benchmark. Unlike the previous case, where the impact had a different effect in the two directions, in this specific scenario it is totally symmetrical and more significant in magnitude.
Coherently, also the optimal profit varies slightly more than it did in the unit price's scenario, even though the gap is not huge.

Lastly, in the third row of figure 11.2, it is showcased the impact of the remaining parameter, the residual value: the parameter varies between the values 0€ and 2€. The resulting optimal probability is [37.5%, 50%], -5.4pp and +7.1pp respectively.
The impact on the probability is the smallest between the economic parameters, as can be observed by comparing the right chart on the third row to the right one in the first row. In fact, even though this is not a general rule, in this specific case they share the same critical fractile and optimal quantity: 50% and 949 units.
However, the profit is not comparable between the two cases: 3,411€ (+893€ units) when the unit price is changed and just 2,558€ (+40€) for the change in resale value.
Its impact is the smallest among all the economic parameters, especially in terms of profitability.

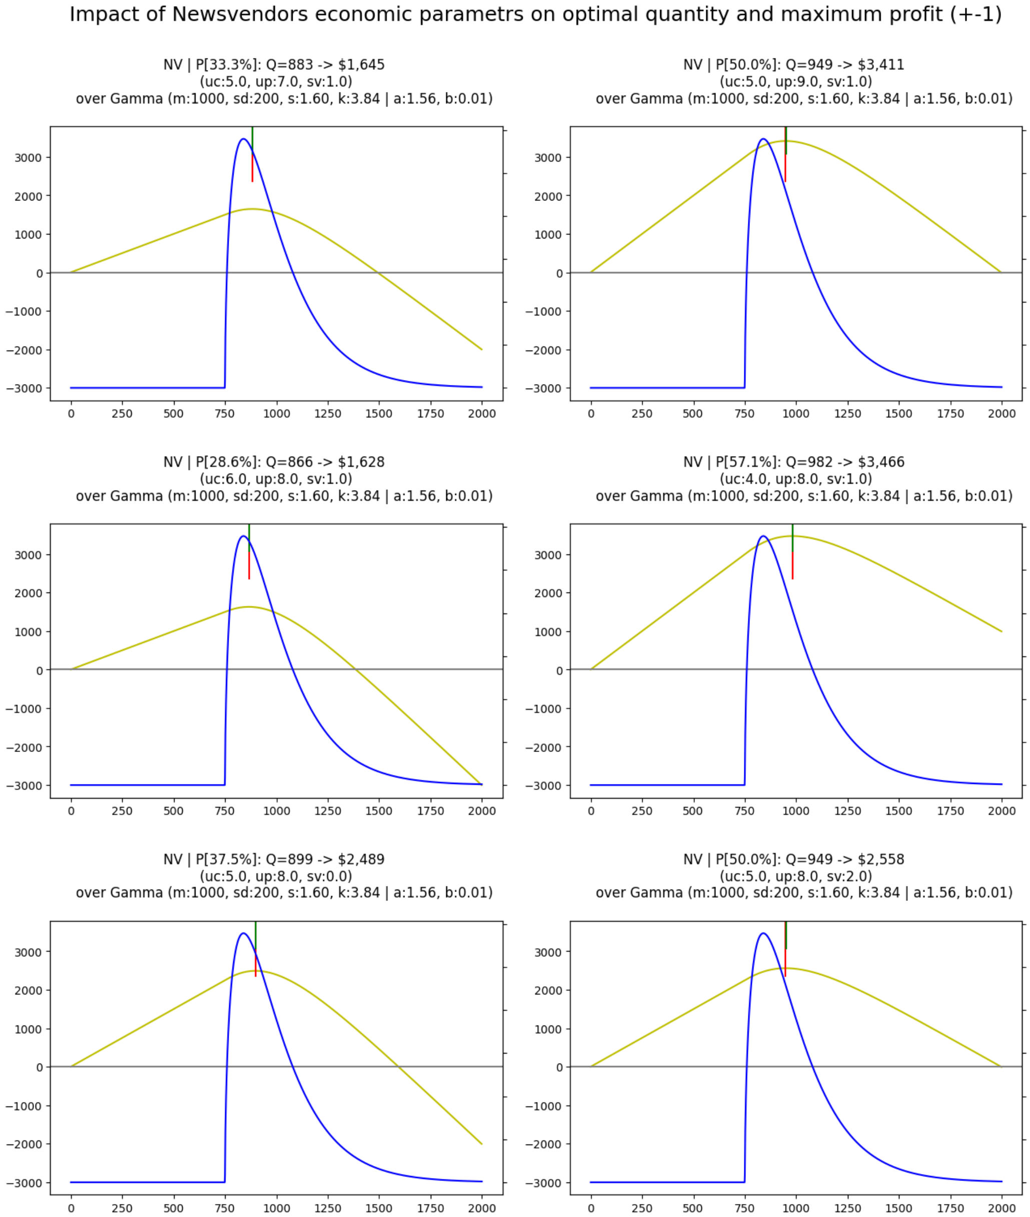

Figure 11.2: Sensitivity analysis performed by varying by +-1€ each economic parameter of the model

11.2 Empirical evidences on the variation of the economic parameters by a percentual amount

In the previous section it has been showed that the unit cost is, for the same change amount, the most impactful of the economic parameters. However, it should be considered that most companies do not perform pricing changes in absolute terms, but rather percentual.
Therefore, the follow-up business question of this section is:

If I could focus my resources to either increase by 10% the unit price / resale value or decrease by the same percentage the unit cost, where should I focus my effort?

In figure 11.3 it is portrayed the effect of a 10% change in all parameters. In this specific instance, the impact on the quantity is very similar between unit price and unit cost in both directions, but the profit varies significantly more changing the unit price rather than the unit cost: the changes in the cumulative probability, associated to the critical fractile, is comprised between -7.4pp and +5.8pp for the unit price and between -7.2pp and +7.2pp for the unit cost, confirming the symmetrical nature of the latter. Interestingly, even though an increase in unit price has a less significant effect on the optimal quantity than a decrease in unit cost, respectively +5.8% and +7.2%, the expected profit it generates is about 250€ higher (3,231€ vs. 2,985€).
Intuitively, this behavior can be justified because the unit price is forcefully higher than the unit cost. Thus, a 10% increase in price per unit is always better than the same decrease in costs. Comparing the profit contribution per unit sold $up - uc$ we obtain

$$8.8 - 5 = 3.8€$$

in case of the variation in unit price and

$$8 - 4.5 = 3.5€$$

when the change is related to the unit cost. Lower than the previous.

Being the resale value the smallest parameter in absolute value, and the one having the smallest impact (as exhibited in section 11.1), a 10% change is almost negligible both in terms of critical fractile change (ranging between -0.6pp and +0.6pp), optimal quantity and profit, with just 7€ difference between the best and the worst case scenario.

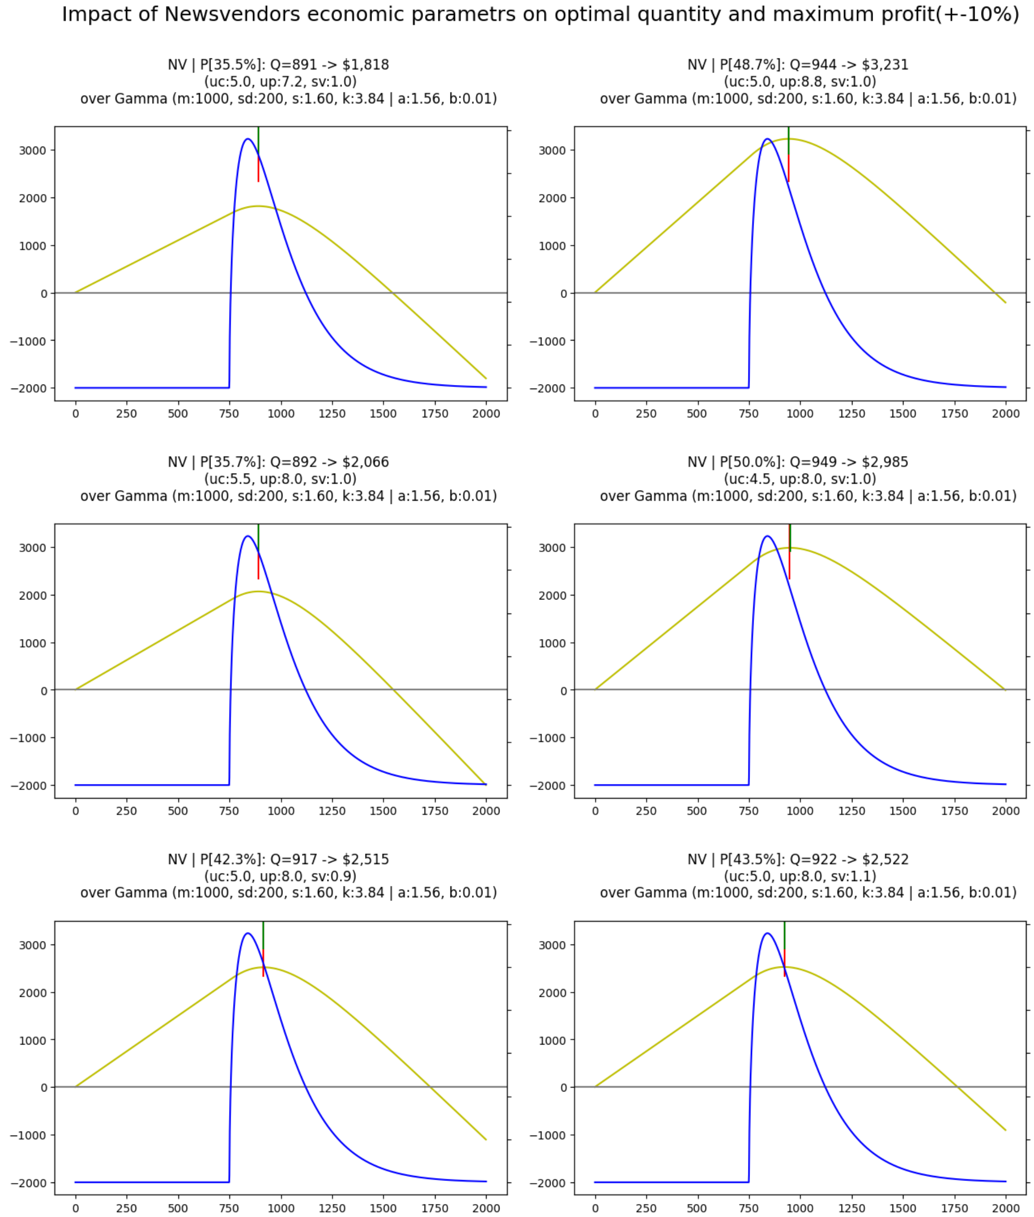

Figure 11.3: Sensitivity analysis performed by varying by +-10% from the baseline each economic parameter of the model

11.3 Mathematical explanation of parameters importance

The relationships between the three economic parameters are

- $up > uc > 0$: otherwise the business would be pursued at a loss, being the best case price of any unit sold lower than the cost sustained to produce it;
- $up > sv$: otherwise the Newsvendor would not make sense at all, given a seasonal price in the optimal period lower than the residual value of the object when it is disposed;
- $uc > sv$: otherwise it would mean unlimited profit, being the residual value at any moment higher than the cost.

Therefore, their ranking is $up > uc > sv$ in any scenario.

The instantaneous slope of the profit function at a quantity q, where $P(q)$ is the value of the cumulative probability of the demand distribution in q, can be computed as

$$m_q = (up - sv)(1 - P(q)) - (uc - sv)$$

that can be rewritten as

$$\begin{aligned} m_q &= up \cdot (1 - P(q)) - sv \cdot (1 - P(q)) - uc + sv \\ &= up \cdot (1 - P(q)) + sv \cdot P(q) - uc \\ &= up - uc - (up - sv) \cdot P(q) \end{aligned}$$

The resulting upper limit of profit per unit can be found when $P(q)$ is equal to 0 and it corresponds to $up - uc$, independent from the resale value; whereas the lower limit is the net cost $-(uc - sv)$ when $P(q)$ equals 100%.

The critical fractile $P(q')$, and therefore the optimal quantity (q'), are dependent on the ratio

$$P(q') = \frac{up - uc}{up - sv}$$

Keeping the distribution constant, the higher the ratio, the higher the expected quantity sold and therefore the higher the expected profit.

Regarding the critical fractile, the impact of increasing the unit price by an amount x_1, over a domain defined as $x_1 \in [uc - up, +\infty)$ is:

$$\frac{up - uc + x_1}{up - sv + x_1}$$

When $x_1 = uc - up$, the lowest accepted value for x_1, the price equals the cost and would make no sense to do business, and in fact the final profit would be 0. On the other end, as $x_1 \to +\infty$, the critical fractile tends to 100% and the expected profits increase as well. Even though the increase in price x_1 needed to make the critical fractile tend to 100% is very high, the final unit price is still an extremely important parameter for the per unit profit generation (m_q), leading to a potentially unlimited increase in profit as it increases.

The impact due to a change in unit cost by an amount x_2 would instead be

$$\frac{up - uc - x_2}{up - sv}$$

opposite in sign than the previous and defined by the domain $x_2 \in [sv - uc, up - uc]$, where both the limits are explained at the very beginning of the section. Unlike before, here the impact is only on the numerator (reason why the impact in percentage points is symmetrical) and each change to the costs could significantly impact the model, being the domain of x_2 limited in a range of $up - uc - (sv - uc) = up - sv$; x_1, by comparison, has unlimited domain. In fact, by changing the unit cost it is possible make the the critical fractile vary across the full spectrum of the distribution $[0\%, 100\%]$, but being the domain of x_2 limited, any decrease in unit cost will have a higher impact on the optimal quantity than the equivalent increase in unit price.
In addition, the unit cost is also by far the most impactful parameter among all, looking at the instantaneous slope of the profit function.

$$m_q = up \cdot (1 - P(q)) + sv \cdot P(q) - uc$$

In fact, unlike the unit price that is slowly weighted out as the cumulative probability increases, the unit cost has a constant impact regardless of the point reached in the distribution.

Finally, when the resale value is changed by an amount x_3, the function of the critical fractile becomes

$$\frac{up - uc}{up - sv - x_3}$$

in the domain $x_3 \in (-\infty, uc - sv]$. Usually, the lower bound of the domain should be $x_3 = -sv$, in order to keep the residual value non-negative. However, this is not a necessary condition: in fact, some products may have a disposal cost if they are not sold, that could additionally decrease both the optimal quantity and the related expected profit of the model.
As $x_3 \to -\infty$, the critical fractile, asymptotically goes to 0%, whereas getting closer to the other extreme ($x_3 = uc - sv$) it rapidly goes to 100%.
Regarding the profit per unit m_q, its impact is similar to the one given by the unit price: its contribution is reduced by a function of the the cumulative probability reached in q. However, such probability ($P(q)$) is the complementary of the one used by the unit price and, consequently, all the models having medium/low critical fractile are almost unaffected by any change in such parameter. It slowly becomes more significant as the value of the critical fractile increases.

Chapter 12

Conclusion

This thesis has explored the sensitivity and robustness of the Newsvendor, a fundamental model in inventory management.
Reviewing the literature, we found out that a broad and comprehensive analysis, regarding the robustness of the model to changes in the distribution or in the economic parameters, was still unseen. In fact, most articles just focus on very specific situations, rather than providing a general view on the sensitivity of the model to such sources of variability.

Consequently, the analysis has been based on an experimental methodology, employing a simulative approach to explore the Newsvendor behavior in the most general way.
As a first result, the reliability of the simulation has been verified in case of some known distributions.
Afterwards, the theoretical model has been tested on some particular counter-intuitive scenarios and it was demonstrated valid for any distribution. It was therefore confirmed that the cumulative probability of the demand distribution related the optimal quantity is only function of the economic parameters and not of the distribution. In addition, we proved that the optimal expected profit depends only on the portion of the distribution up to the purchased quantity, regardless of the shape of the curve afterwards.

Finally, we performed various sensitivity analysis on the most important measures associated to the demand distribution. According to our findings, mean and standard deviation are the two most important factors impacting the profitability of a specific Newsvendor. In fact, a higher mean results in significantly higher expected profit at any quantity; whereas the higher the standard deviation, the smaller the expected profit, because of the increased risk due to variability.

The specific impacts related to changes in the skewness and kurtosis are instead significantly less prominent and can be ignored in first approximation.

As a general rule, the distribution favorably impacts the profitability when the optimal quantity occurs at a quantity that is slightly higher than the mode of the distribution. In addition, if the probability density function is highly concentrated in a close range around the optimal, this also positively impacts the expected profit.

Among all the sources of variability, the economic parameters are the only which could be partially controlled by the decision maker. Their impact on the Newsvendor is of primary importance, especially in case of a change in unit price or cost. In addition, Newsvendors defined by a high critical fractile have higher profitability margins and are less likely to be disrupted by an estimation error in the underlying distribution; whereas those described by a lower critical fractile not only have a lower marginal profit per unit sold, but are also riskier and more subject to variability.

12.1 Limitations of current analysis and suggestions for future research

The performed analysis has been conceived to be as comprehensive and general as possible. However, many additional in-depth studies can be executed starting from the basis created with this work.
An insightful follow-up could involve a more formal and rigorous approach in explaining the impact of standard deviation changes on both optimal profit and quantity for symmetrical distributions, as empirically observed in chapter 7.
In addition, by enhancing the simulation capabilities, it is possible to include additional dimensions to the analysis, that would be impossible to tackle using a pure analytical approach. For example, price and demand are extremely correlated between them, but the impact of this relationship has not been investigated throughout this thesis. Below, some suggestions for additional in-depth analysis on the topic:

- A study of a scenario where the demand distribution is function of the unit price, that is fixed and defined ex-ante. The goal of such analysis would be to find the optimal price that maximize the total expected profit, by balancing the positive effect of an increase in unit profit to the negative impact of a decrease in volume.

- An analysis encompassing the concept of dynamic pricing: such strategies, like price skimming (where the price slowly reduces over time, in order to maximize the profitability of the early adopters, but afterwards optimizing the quantity when the price decreases), make the price vary over time. This would enable the analysis of the scenario where the price is function of the realized demand.

In conclusion, many different analysis could be performed by incorporating more complex real-world factors, in order to fill the gaps of the literature on such a valuable tool for inventory management and providing significant value to managers and decision makers.

Bibliography

[Andrew Butters(2019)] R. Andrew Butters. On demand uncertainty in the newsvendor model. *Economics Letters*, 185:108746, 2019. ISSN 0165-1765. doi: https://doi.org/10.1016/j.econlet.2019.108746. URL https://www.sciencedirect.com/science/article/pii/S016517651930374X.

[Borgonovo(2010)] E. Borgonovo. Sensitivity analysis with finite changes: An application to modified eoq models. *European Journal of Operational Research*, 200(1):127–138, 2010. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2008.12.025. URL https://www.sciencedirect.com/science/article/pii/S0377221708010631.

[DeMarle(2019)] Pat DeMarle. The basics of the newsvendor model. *Medium*, 2019. URL https://medium.com/@pdemarle/the-basics-of-the-newsvendor-model-ef756f203433.

[Hedayatinia et al.(2020)Hedayatinia, Lemoine, Massonnet, and Viviani] Pooya Hedayatinia, David Lemoine, Guillaume Massonnet, and Jean-Laurent Viviani. Analysis of the impact of demand volatility and return policies on a price-setting newsvendor. *IFAC-PapersOnLine*, 53(2):10843–10848, 2020. ISSN 2405-8963. doi: https://doi.org/10.1016/j.ifacol.2020.12.2871. URL https://www.sciencedirect.com/science/article/pii/S2405896320336375. 21st IFAC World Congress.

[Jammernegg et al.(2022)Jammernegg, Kischka, and Silbermayr] Werner Jammernegg, Peter Kischka, and Lena Silbermayr. Heterogeneity, asymmetry and applicability of behavioral newsvendor models. *European Journal of Operational Research*, 301(2):638–646, 2022. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2021.11.014. URL https://www.sciencedirect.com/science/article/pii/S0377221721009607.

[Khanra et al.(2014)Khanra, Soman, and Bandyopadhyay] Avijit Khanra, Chetan Soman, and Tathagata Bandyopadhyay. Sensitivity analysis of the newsvendor model. *European Journal of Operational Research*, 239(2):403–412, 2014. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2014.05.039. URL https://www.sciencedirect.com/science/article/pii/S0377221714004627.

[Krishnendu Adhikary and Kar(2018)] Jagannath Roy Krishnendu Adhikary and Samarjit Kar. A distribution-free newsboy problem with fuzzy-random demand. *International Journal of Management Science and Engineering Management*, 13(3):200–208, 2018. doi: 10.1080/17509653.2017.1381051. URL https://doi.org/10.1080/17509653.2017.1381051.

[Liu et al.(2022)Liu, Letchford, and Svetunkov] Congzheng Liu, Adam N. Letchford, and Ivan Svetunkov. Newsvendor problems: An integrated method for estimation and optimisation. *European Journal of Operational Research*, 300(2):590–601, 2022. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2021.08.013. URL https://www.sciencedirect.com/science/article/pii/S0377221721006810.

[Qin et al.(2011)Qin, Wang, Vakharia, Chen, and Seref] Yan Qin, Ruoxuan Wang, Asoo J. Vakharia, Yuwen Chen, and Michelle M.H. Seref. The newsvendor problem: Review and directions for future research. *European Journal of Operational Research*, 213(2):361–374, 2011. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2010.11.024. URL https://www.sciencedirect.com/science/article/pii/S0377221710008040.

[Yeoman and McMahon-Beattie(2017)] Ian Seymour Yeoman and Una McMahon-Beattie. The turning points of revenue management: a brief history of future evolution. *JOURNAL OF TOURISM FUTURES*, 3(1):66–72, 2017. ISSN 2055-5911. URL https://pdfs.semanticscholar.org/716d/3fa46c3dd8a5ee0cd08534cde403dd163766.pdf.

Appendix A

Code Implementation in Python

Figure 3.1 in chapter 3 showed the relationships between classes and in section 3.2 was explained the high level behavior of the components of the simulation. More detailed information on the technical implementation of the different classes is available in the following sections:

A.1 Distribution interface and sub-classes

The Distribution class is the perfect example of inheritance: in fact, it is created as an interface with all the basic attributes and methods. The constructor expects as parameters the 4 main measure (mean, standard deviation, skewness and kurtosis) and a flag that, if true, reverses the distribution on its y-axis. This super-class is then specified into 5 sub-classes:

- **Uniform:** Based on SciPy's uniform distribution, it equally spreads the probability between the minimum and maximum points provided, centered in the mean.
- **Normal and Exponential:** Obtained from SciPy and built using the mean and standard deviation provided.
- **Gamma:** Probability function based on 2 parameters (α and β). These are obtained directly starting from the main three moments. Gamma distribution is dependent on skewness and can represent all the curves between exponential ($\gamma = 2$) and normal ($\gamma = 0$) functions.
- **Beta:** Generalization of the gamma function, it needs 2 parameters as well. Differently from the latter, where it is possible to calculate α and β analytically starting from the three measures provided, the addition

of the kurtosis make this impossible.
Therefore, this distribution is handled in a totally custom way: an utility function creates almost 100,000 distribution for different values of α and β, varying with uneven steps in the range [0.01, 1000] and saving them into a Pandas dataframe. Then, whenever a new beta distribution is initialized, it looks for the closest values of skewness and kurtosis in the curves inside the dataframe and scales them to match also the mean and standard deviation provided.
This workaround allows to precisely obtain the desired beta starting from the four moments,

- **Custom (Boxed):** This is not an out of the box distribution, but a custom built curve that is used in chapter 5 to tests the general validity of the Newsvendor model. It can be explained as a n-uniform distribution, where n is the number of different distributions in the total range.
  Each uniform has at its disposal $\frac{1}{n}$ of the space between min and max and a probability equal to $\frac{1}{n}$. The uniform distribution is not forced to occupy the totality of the space available: in fact, it can also be condensed at the beginning, middle or end of its 'box', creating zones with different densities.

All these sub-classes share the main methods and attributes, but some of them have specific implementations, thanks to the polymorphism property. Thus, even though the behaviors within the sub-classes are different, they interact with the external world in the exact same way.
The logic and purpose of the main methods is exemplified below:

- **Print:** Two functions used to explain the distribution type and moments in two different formats (one minimal and the other more verbose)

- **Create demand vector:** According to the average and the standard deviations provided, it creates a vector of all the demand quantities possible.

- **Create probability vector:** For each value in the demand vector, it calculates the probability that each quantity will be the realized one. It is calculated as:

[fn.cdf(d_i + 0.5) - fn.cdf(d_i - 0.5) for d_i in self.d_vector]

Using words, because the quantity is an integer but the functions are continuous, the probability needs to be discretized by subtracting to the cumulative at Q+0.5 the cumulative at Q-0.5, for every value of demand (Q).

- **Get quantity:** method to be called from outside, it returns the quantity associated to the cumulative probability provided.
- **Get cumulative probability:** Reverse of the previous function, it returns the cumulative probability associated to the quantity provided.
- **Plot:** Function used to easily show the distribution in a graphical way.

A.2 Newsvendor class

This class contains in its attributes the values of all the features proper of a Newsvendor Model, including:

- **Economic parameters:** The three parameters that describe any Newsvendor Model. They are provided during the instance creation and saved as key parameters within the object data.
- **Distribution:** Second parameter of the constructor, this object can be any instance of the Distribution sub-classes we analyzed in the previous section. The class will manage by itself any manipulation needed on the distribution, as well as any data retrieval related to quantity and probability.
- **Optimal solution:** The best solution of the model, calculated both iterating and analytically in terms of cumulative probability, quantity associated and derived expected profit.
- **Profit vector:** Output of the iterative method, it is the expected profit (based on the demand distribution) associated to each possible quantity that can be purchased.

Below are listed the main methods in which all these logic processes and capabilities are described:

- **Print:** Two functions, having different detail levels, used to describe the Newsvendor explicating its economic parameters and the underlying demand curve (by calling the same method of the Distribution class).

- **Calculate optimal solutions:** The default run function of the class, compute the optimal value in both the approaches listed above.
- **Calculate sub-optimal solution:** Method called for benchmarking purposes. It is used to compute the profit at a given quantity and compare it to the optimal one.
- **Plot:** Function used to plot on a line chart both the profit curve and demand distribution, highlighting the optimal value.

A.3 Model class

The parent class used to compare different simulations, takes as main parameter a list of Newsvendors. Its most important use is to compare the different simulations in a significant way, both graphically and analytically. It plots all the models in a matrix (of specified size) by forcing them to share both the x and y axis, improving readability.
It also enables direct benchmarking by comparing them in three different ways:

- **Standard:** Calculate all the models at the optimal quantity specific to the selected benchmark, immediately showing sensitivity to the change.
- **By columns:** The benchmarking is no longer a single occurrence, but a whole row of the matrix. All the models of each column will be compared to the Newsvendor in the benchmark row of the same column.
- **By rows:** Same exact behavior of the previous, but swapping the comparison from columns to rows.

Appendix B

Probability density around the mean for different distribution types

As analyzed during part III, the Newsvendor model provides better results (at constant economic parameters) when the portion of the demand distribution just before the optimal quantity is more concentrated.
In figures B.1 - B.5 are illustrated the shape of the most known distributions, highlighting their mean (red line) and the probability range between $[\mu - \sigma, \mu + \sigma]$, identified by the green lines. In addition, the dashed gray lines provides information on the quartiles, showing the quantities related to a cumulative probability of 25%, 50% and 75%.
These figure are provided to act as a reference to assess whether the distribution associated with the Newsvendor has a positive impact on the expected profit (when the optimal quantity is slightly above the mode and the distribution is highly concentrated in a close range around it) or when its effect is negative (in case the optimal quantity falls at a point with low probability density).

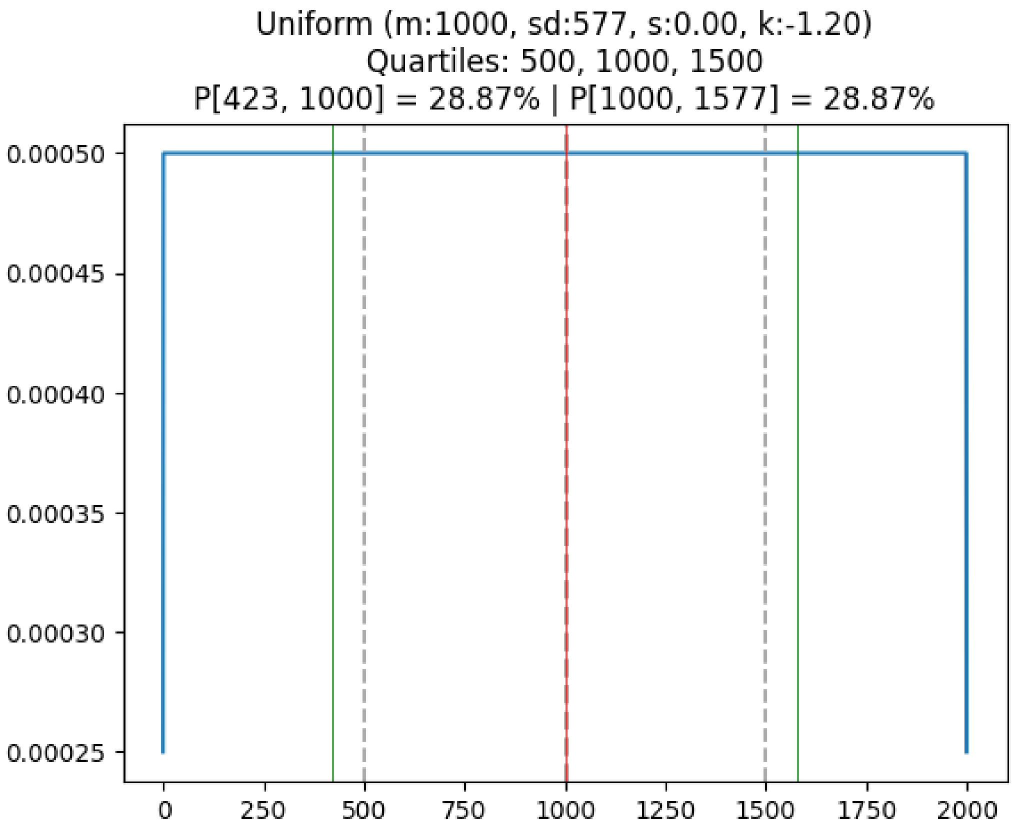

Figure B.1: Probability density function for the uniform distribution, highlighting mean, median and other quartiles

The Uniform distribution, depicted in figure B.1 , is the most sparse curve possible for a given tuple of mean and standard deviation. In fact, in the range $[\mu-\sigma, \mu+\sigma]$ it is included less than the 58% of the total probability.

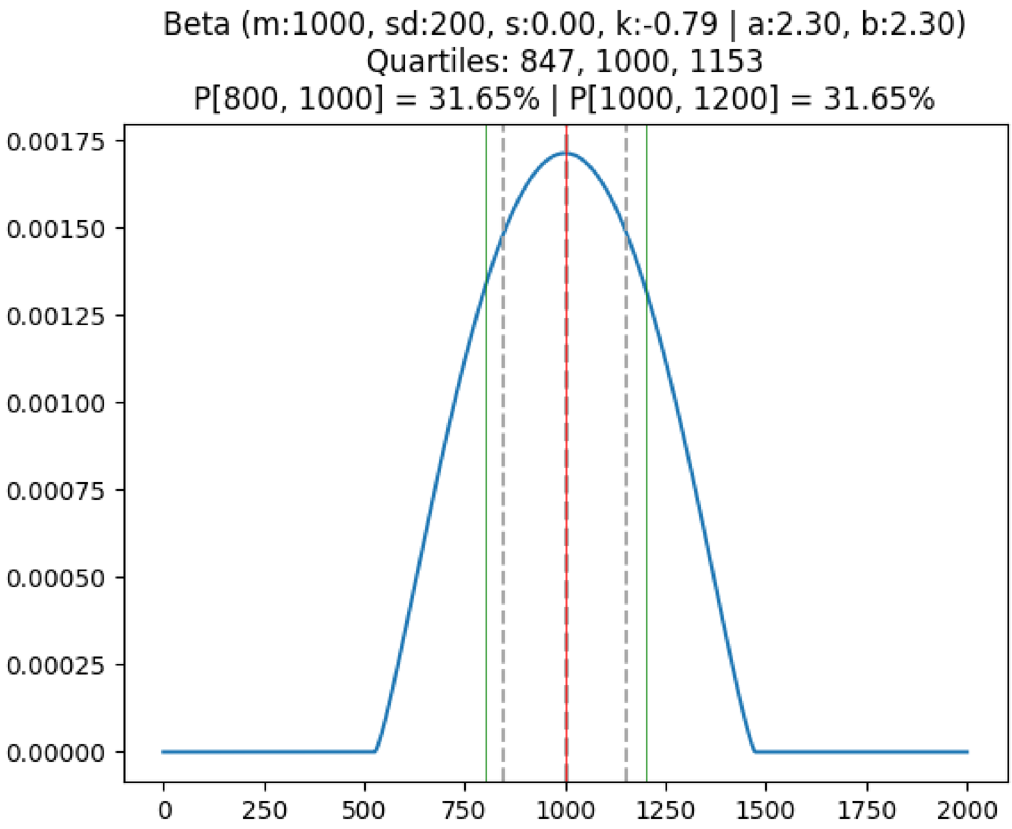

Figure B.2: Probability density function for the beta distribution, highlighting mean, median and other quartiles

A symmetrical Beta distribution, mid-where between the uniform and the normal, is slightly more concentrated towards the mean. In fact, in figure B.2 it is possible to notice that the range $[\mu - \sigma, \mu + \sigma]$ contains approximately 63% of the probability.

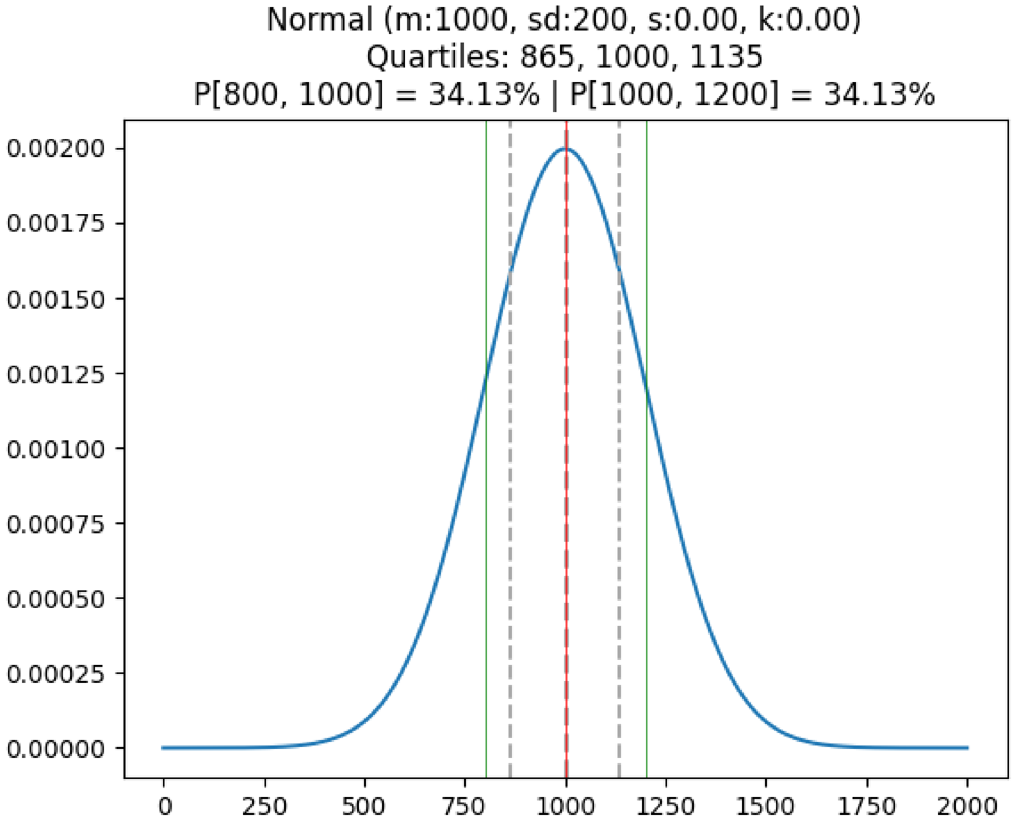

Figure B.3: Probability density function for the normal distribution, highlighting mean, median and other quartiles

The Normal distribution (figure B.3) is still symmetrical, but having a higher kurtosis it is more concentrated in the range $[\mu - \sigma, \mu + \sigma]$, surpassing 68%.

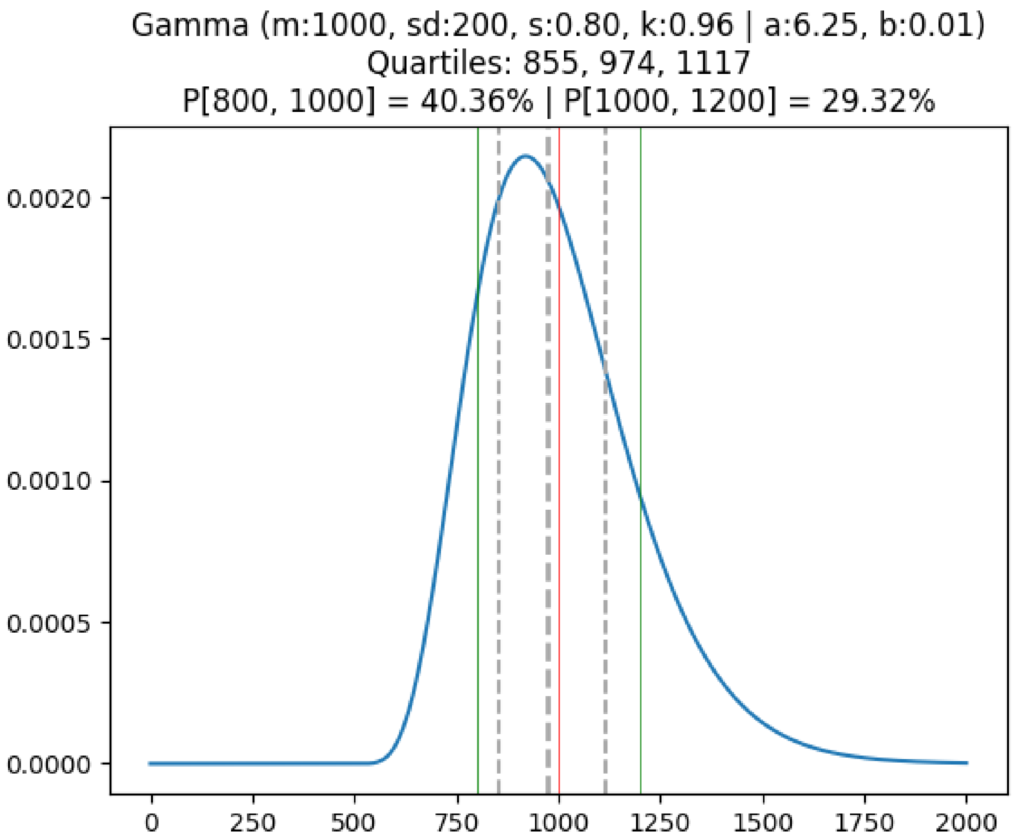

Figure B.4: Probability density function for the gamma distribution, highlighting mean, median and other quartiles

The Gamma distribution, whose example can be found in figure B.4, is an asymmetrical distribution defined by a positive kurtosis. In fact, in can range from $\kappa = 0$ (normal distribution) to $\kappa = 6$ (exponential distribution). Therefore, it is even more concentrated than the previous, reaching almost 70% of the cumulative probability in the range $[\mu - \sigma, \mu + \sigma]$. Such probability is not fixed, but depends on the value of α chosen to create the distribution. Like the kurtosis, it ranges from the probability of the normal distribution to the one obtainable with the exponential, which are both constant.

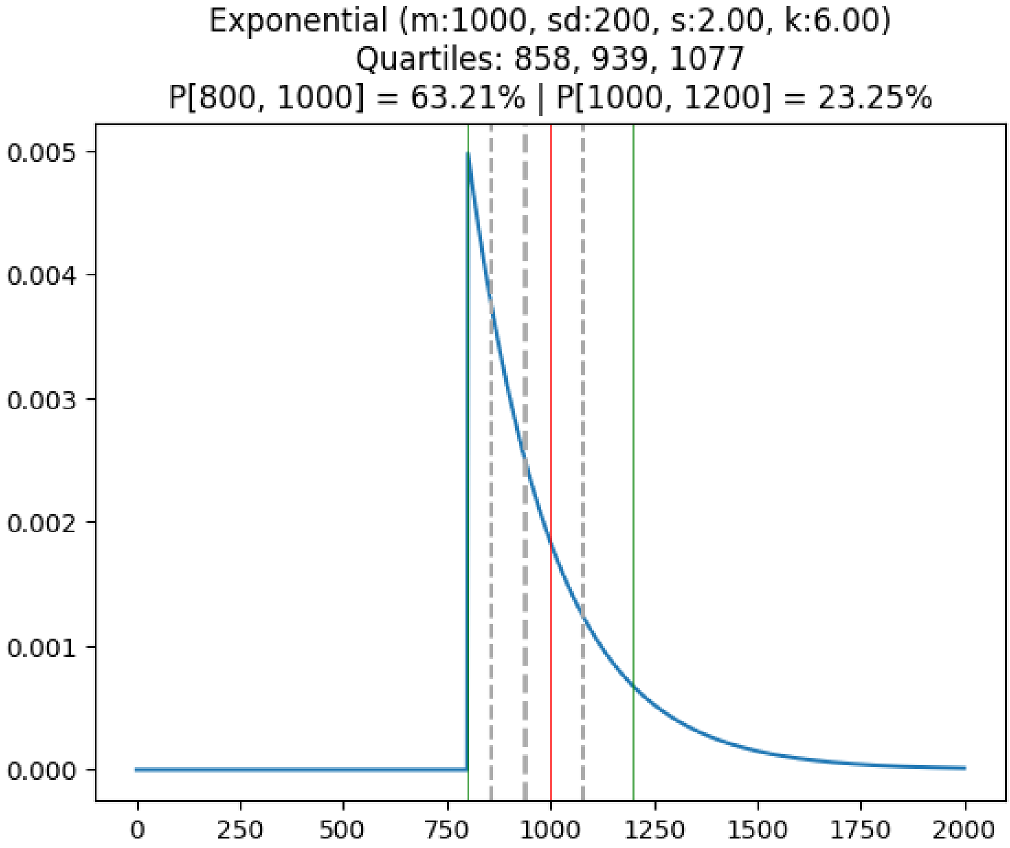

Figure B.5: Probability density function for the exponential distribution, highlighting mean, median and other quartiles

Lastly, the Exponential distribution, as shown in figure B.5, is the most concentrated among those analyzed in this thesis ($\kappa = 6$), including more than 63% of the total probability in half of the range of the previous distributions: $[\mu - \sigma, \mu]$ (or $[\mu, \mu + \sigma]$ if reversed).