



DEPARTMENT OF MANAGEMENT SCIENCE

PHD MANAGEMENT SCIENCE

Extensions to Newsvendor Problems

Congzheng Liu

Submitted for the degree of Doctor of Philosophy at Lancaster
University.

2022

Acknowledgement

The path to PhD has been the most exciting journey I have had in my life. Throughout this journey, I have received help from many people. Here, I would like to thank the people below in particular.

This path would not have been possible without the hard work of my dear supervisors, Adam Letchford and Ivan Svetunkov, who not only supported me academically but also cheered me up at my lowest points. I am very grateful to both of you for your guidance and support!

Additionally, I would like to thank all the supporting staff in the Department of Management Science, who were very professional at their jobs and provided me with all the help I could possibly require. Without their work, my PhD path would have been much harder than it was.

Last but not least, I am endlessly grateful to my wife Ming for her patience and unconditional support. We started the journey as boyfriend and girlfriend. Now we finish it as husband and wife. Special thanks go to all my family members and my friends. I would not have been where I am now without their encouragement.

Declaration

This thesis is my original work and has not been submitted, in whole or in part, for a degree at this or any other university. Nor does it contain, to the best of my knowledge and belief, any material published or written by another person, except as acknowledged in the text.

This thesis is constructed as a series of researches and thus Chapters 2 to 5 should be read as separate entities.

Chapter 2 has been published as Liu, C., Letchford, A.N. and Svetunkov, I., 2021. Newsvendor problems: An integrated method for estimation and optimisation. European Journal of Operational Research.

Chapter 3 has been submitted for publication as Liu, C., Letchford, A.N. and Svetunkov, I., 2022. On Sensitivity and Parametric Analysis for Multi-Item Newsvendor Problems. EURO Journal on Computational Optimisation.

Chapter 4 has been submitted for publication as Liu, C., Letchford, A.N. and Svetunkov, I., 2022. Naïve Newsvendor Adjustments: Are They Always Detrimental? Journal of the Operational Research Society.

Chapter 5 has been submitted for publication as Liu, C., Zhu, W., 2022. Newsvendor Conditional Value-at-Risk Minimisation on Non-parametric Approach. Operations Research.

Abstract

News vendor problems (NVPs) form an important and much-studied family of inventory control problems. Although the use of the term varies somewhat, in most situations the term NVP refers to a single-period stochastic inventory control problem involving a single product. Assuming that the demand comes from a known probability distribution, this classic problem can be solved easily with calculus (Arrow et al., 1951), and the solution appears in nearly all inventory management textbooks.

In this thesis, we expand the literature in four directions. In Chapter 2, we consider an integrated approach, in which the NVP order quantities are determined directly from the data. Though the topic of integrated approaches has already been studied in the literature, the idea of constructing a robust approach that deals with nonlinear NVPs is novel. In this chapter, we introduce such an approach, and we perform extensive simulation experiments to examine the performance of the approach in different settings, including situations when the true model is known and when the underlying model is mis-specified.

In Chapter 3, we consider the effect that small changes in NVP parameters would have on the optimal solution, which is commonly referred to as sensitivity analysis. We show that one can perform sensitivity analysis for NVP using techniques from stochastic programming and discrete approximation. Our method is very general and can handle changes in prices and costs, changes in demand distributions, and

cross-price elasticities of demand. Moreover, computational results show that our method yields accurate estimates with very reasonable computing effort.

In Chapter 4, we examine the effect of judgemental adjustments in an NVP context. Several attempts have been made to quantify the outcomes of such adjustments. However, much of this literature assumes that accurate demand forecasts are available. We consider the (more realistic) case in which the forecasts may be inaccurate, due for example to insufficient data or model mis-specification. Computational results indicate that, in some cases, judgemental adjustment can lead to an increase in profit rather than a decrease. We discuss conditions under which the adjustments are beneficial and the situations when they are not. We also propose a heuristic algorithm for “tuning” the adjustment parameters in practice.

In Chapter 5, we propose an alternative non-parametric approach to the variant of the NVP in which the goal is to minimise the conditional value at risk (CVaR). Given the difficulties with treating observations with extreme values, the existing parametric methods often underestimate the downside risk and lead to a significant loss in extreme cases. The existing non-parametric methods, on the other hand, are extremely computationally expensive with large instances and depend heavily on the form of the profit function. Using both simulation and real-life case studies, we show that our proposed method can be very useful in practice, allowing decision-makers to suffer far less downside loss in extreme cases while requiring reasonable computing effort.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Newsvendor Problems	4
1.3	Linear Programming	6
1.4	Stochastic Programming	8
1.5	Quantile Regression	9
1.6	Judgemental Adjustment	10
1.7	Structure of the Thesis	11
2	Newsvendor Problems: An Integrated Method for Estimation and Optimisation	13
2.1	Introduction	14
2.2	Literature Review	16
2.2.1	The classical newsvendor problem	16
2.2.2	More complex newsvendor problems	17

2.2.3	Quantile regression	19
2.2.4	The Ban and Rudin approach	20
2.3	Analysing and Extending the Ban and Rudin Method	20
2.3.1	Analysis	21
2.3.2	Extension to the nonlinear case	21
2.4	Computational Experiments	22
2.4.1	Experimental setup	23
2.4.2	When the true model is known	25
2.4.3	When the true model is not known	31
2.5	Real-life Case	36
2.6	Concluding Remarks	41
3	On Sensitivity and Parametric Analysis for Multi-Item Newsvendor Problems	42
3.1	Introduction	43
3.1.1	Notation and terminology	44
3.2	Literature Review	46
3.3	The Stochastic Program and Its Linear Approximation	48
3.4	Sensitivity and Parametric Analysis	52
3.4.1	Some elementary results	52

3.4.2	Changes in the demand distributions	53
3.4.3	Cross-price elasticities	54
3.5	Results with Artificial Data	56
3.5.1	Instance generation	56
3.5.2	Increasing the number of scenarios	56
3.5.3	Increasing the problem size	60
3.6	Results with Real Data	62
3.7	Concluding Remarks	66
4	Naïve Newsvendor Adjustments: Are They Always Detrimental?	68
4.1	Introduction	68
4.2	Literature Review	71
4.2.1	Demand chasing effect	71
4.2.2	Pull-to-centre effect	73
4.3	An Alternative Perspective and Model	74
4.3.1	In favour of naïve adjustment	74
4.3.2	An integrated adjustment model	75
4.4	Experiments on Simulated Data	77
4.4.1	Methodology	77
4.4.2	When the DGP is known	79

4.4.3	When the model is misspecified	83
4.5	Tuning Algorithm	86
4.5.1	Procedure	87
4.5.2	Real life example	88
4.5.3	Results	89
4.6	Concluding Remarks	91
5	Newsvendor Conditional Value-at-Risk Minimisation on Non-parametric Approach	93
5.1	Introduction	94
5.2	Single-Item Newsvendor Problems	97
5.3	NVPs Under CVaR Minimisation	98
5.4	Non-Parametric CVaR Minimisation	101
5.4.1	Empirical CVaR minimisation via NPC	102
5.4.2	Adaptive CVaR minimisation via NPC	103
5.5	Baseline Experiment	109
5.5.1	Experimental setup	110
5.5.2	Results of the baseline experiment	112
5.6	Experiments With Other Parameters	116
5.6.1	Varying the number of features	117

5.6.2	With other profit functions	119
5.6.3	With other forms of the error term	121
5.6.4	Sum up	123
5.7	Real-life Example	123
5.8	Concluding Remarks	128
6	Conclusion	130
6.1	Summary	130
6.2	Further Research	133
	Appendix	136
A	Proof of Equivalence to Quantile Regression	136
B	Proof of Maximising-Minimising Transformation	137
C	Small Artificial Instance	138
D	Constraints for Real-Life Example	139
E	Elasticities for Real-Life Example	139
F	Heatmaps for model misspecification	140
G	Expectation maximisation vs. CVaR minimisation	141
H	Baseline experiment parameters	142
I	Histogram example	143
J	Baseline experiment full results	144

K	Features for food preparation problem	146
Bibliography		147

Chapter 1

Introduction

1.1 Motivation

News vendor problems (NVPs), formerly called *Newsboy problems*, are a classic topic in inventory control (Arrow et al., 1951; Silver et al., 1998). In short, they are an example of a mathematical model that can be used to determine optimal inventory levels under demand uncertainty. They are relatively simple inventory control models since they are focused on a single planning period. Nevertheless, they have received considerable attention from Operational Researchers, due to their many applications. For example, they are used for making inventory decisions for perishable products in retail. They are also used in other areas, such as wholesaling, manufacturing, transportation and insurance.

The earliest known work to deal with an NVP is Edgeworth (1888). Edgeworth used the central limit theorem to determine the optimal cash reserves to satisfy random withdrawals from depositors. The term “newsboy” was first mentioned in Morse and Kimball (1951), while the modern formulation follows Arrow et al. (1951). Since then, a large number of works have appeared on the topic (see, e.g., Choi, 2012;

Porteus, 2002; Silver et al., 1998; Zipkin, 2000 for surveys).

Although NVPs have received a great deal of attention, many interesting questions remain unanswered, including some of the practical importance. We address several of these questions in this thesis. We consider not only the Newsvendor problems themselves but also the forecasting and marketing decisions associated with them. We explain our motivation from the following four directions.

In early work on NVPs, it is assumed that the demand for each product in any given time period comes from a known probability distribution, and the only decision to make is on the inventory levels. In practice, however, the real demand distribution is never known. In fact, estimating demand distributions is undoubtedly one of the main challenges in inventory management. The typical solution for this issue is to use a *disjoint* method; that is, estimating the demand distribution first, and determining the inventory levels second. However, this approach has one major drawback: in certain situations, it can lead to severe bias, resulting in severe profit loss (see e.g., Bertsimas and Thiele, 2005; Beutel and Minner, 2012; Karmarkar, 1994; Korpela and Tuominen, 1996). To get around these difficulties, we propose an *integrated* approach.

Furthermore, in reality, the planning decisions are more complicated than in the classical NVPs, involving not only the ordering policy but also decisions about pricing and promotional activities. In turn, these decisions influence demand for products and may lead to changes in quantities ordered, creating a feedback loop. This is especially important for companies that have to make such decisions for thousands of units in real-time and for which changes in prices or promotional activities of some products might influence sales of other, related products. To address these issues, we provide a tool to help decision-makers understand the potential effects of their decisions. We show that *Sensitivity Analysis* (SA) and *Parametric Analysis* (PA) can be considered useful tools in this situation.

The third aspect of NVPs that we consider is the effect of *judgemental adjustments*. In the textbook formulation, it is assumed that a decision maker has a correct model of the demand distribution, with correct parameters. In real life, however, correctness is rarely assured. Moreover, even if the model is correct, the parameters may evolve over time (for example, due to market shocks or product innovations by competitors). For these and other reasons, decision-makers often make so-called “judgemental adjustments” to the theoretically “optimal” order quantities. To evaluate the effect of judgemental adjustments, we perform a numerical study. In particular, we focus on two specific kinds of adjustments which are normally considered to be particularly naïve: *demand chasing* and *pull-to-centre*. We discuss how these adjustments work in practice and what they imply in a variety of settings.

Finally, we consider the case in which the decision maker prefers to minimise the *conditional value-at-risk* (CVaR) rather than maximise the expected profit. The CVaR is currently a very popular measure in financial risk management. However, given that the CVaR concerns the observations with extreme values, which are often treated as outliers in traditional statistical approaches, the parametric methods often underestimate the downside risk and lead to a significant loss in extreme cases. The existing non-parametric methods, on the other hand, are extremely computationally expensive and depend heavily on the form of the profit function. We propose an alternative non-parametric approach, which sidesteps the aforementioned limitations.

To summarise, our study is divided into four parts:

1. Development of an integrated approach that determines the order quantities directly from past data;
2. Construction of a “feedback loop” which uses information gained in the optimisation phase to inform marketing decisions;

-
3. Evaluation of the effect of judgemental adjustments on inventory decisions under a variety of conditions;
 4. Development of an alternative non-parametric approach to CVaR minimisation in the Newsvendor context.

In the following five subsections, we give the background needed to understand the subsequent chapters. The subsections cover NVPs, linear programming, stochastic programming, quantile regression and judgemental adjustment.

1.2 Newsvendor Problems

Newsvendor problems involve the determination of the optimal ordering policy, given a forecast of demand distribution. In the simplest NVP, as defined, for example, by Arrow et al. in 1951, a company purchases goods at the beginning of a time period and aims to sell them by its end. The demand is a random variable with a known probability distribution. At the end of the period, any surplus goods will lead to a *disposal cost*. On the other hand, a shortage of goods will lead to a *shortage cost*. The goal is to determine an *order quantity* prior to the period, which would maximise the expected profit.

Since the introduction of the basic NVP, researchers have considered many extensions of the problem, including variants with multiple product types (Hadley and Whitin, 1963; Lau and Lau, 1996; Moon and Silver, 2000), quantity discounts (Khouja, 1995), different risk measures (Eeckhoudt et al., 1995), alternative objectives (Anvari, 1987; Eeckhoudt et al., 1995), product substitution (Bassok et al., 1999), nonlinear cost functions (Halman et al., 2012), non-stationary demand (Kim et al., 2015), multiple decision periods (Matsuyama, 2006), and price setting (Karlin and Carr, 1962; Mills, 1959; Petruzzi and Dada, 1999).

Nonlinear Newsvendor problems, which we call NNVPs, allow the disposal and shortage costs to be nonlinear functions of ordered quantity, instead of being linear. Using the NNVP enables one to model a larger variety of real-life problems. For example, in real life, a minor shortage may not cause large costs, but a major shortage could damage the reputation of the company. As another example, a small amount of excess stock can often be sold at a discount, but this may not be possible for a large amount of stock. Moreover, as a product stays on the shelf longer, the opportunity cost of the shelf space may increase over time. All these examples show that the costs can be non-linear in real-life situations, suggesting that the NNVP could be applied frequently in practice. Additionally, numerical studies have shown that the order quantities suggested by NNVP are usually different from the quantities from its NVP approximating model (Halman et al., 2012; Kyparisis and Koulamas, 2018).

Multi-item (a.k.a. multi-product) NVPs, which we call MNVPs for short, have n products instead of one (see Turken et al., 2012 for a survey). The company needs to determine the order quantity for each product simultaneously. There are also one or more side constraints, such as constraints on total storage space or the total budget available for purchasing. Several solution methods for MNVPs are available, see, e.g., Abdel-Malek and Areeratchakul, 2007; Ben-Daya and Raouf, 1993; Hadley and Whitin, 1963; Lau and Lau, 1995; Nahmias and Schmidt, 1984.

Another well-studied extension of NVP concerns the minimisation of the *conditional value-at-risk (CVaR)*, a most preferable measure in financial risk management (Rockafellar and Uryasev, 2002). This is due to the fact that some stakeholders focus on what they could lose in extreme situations rather than on what they could gain on average. In the work of Gotoh and Takano (2007), a closed-form solution was given for the CVaR variant of the NVP. Unlike the solution to the classical version, the CVaR solution takes the form of a weighted average of two critical quantiles.

Thus, depending on the cost parameters, the CVaR minimisation solution can be either greater than or less than the expectation maximisation solution.

1.3 Linear Programming

One of the tools that we will find useful in our study of NVPs is linear programming. A *Linear Programme* (LP) is a special kind of optimisation problem, in which one wishes to maximise or minimise a linear objective function of a set of decision variables, subject to linear constraints (Dantzig, 1955). Geometrically speaking, the set of feasible solutions to an LP is a convex polyhedron and the goal is to find a point in the polyhedron where the objective function has the largest (or smallest) value if such a point exists.

Following Dantzig (1955), an LP with n variables and m constraints can be written in the form

$$\max \{ \mathbf{c}^T \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \in \mathbb{R}_+^n \}, \quad (1.1)$$

where the components of \mathbf{x} are the variables whose values are to be determined, $\mathbf{c} \in \mathbb{Q}^n$ is the *objective function* vector, $\mathbf{b} \in \mathbb{Q}^m$ is the vector of *right-hand sides*, and $\mathbf{A} \in \mathbb{Q}^{m \times n}$ is the *constraint matrix*.

A variety of solution approaches (a.k.a. algorithms) have been discovered for LPs. The most well-known approaches are the Simplex method (Dantzig, 1963) and Interior-Point Methods or IPMs (Gondzio, 2012; Marsten et al., 1990). Nowadays, there exist many excellent commercial and academic software packages for solving LPs (and some more general problems) that implement these methods. Many of these solvers are capable of solving LPs with thousands of variables and/or constraints in a reasonable time. (Examples of such solvers include CPLEX, Gurobi and Xpress.)

Given an LP of the form (1.1), called the *primal* LP, one can form another LP, called the *dual*. The dual LP takes the form:

$$\min \{ \mathbf{b}^T \mathbf{y} : \mathbf{A}^T \mathbf{y} \geq \mathbf{c}, \mathbf{y} \in \mathbb{R}_+^m \}.$$

The *strong duality theorem* states that the optimal profit for the primal LP is equal to the optimal cost for the dual LP (Gale et al., 1951). The above-mentioned solution methods, Simplex and IPMs, compute optimal primal and dual solutions, both of which are essential when performing *sensitivity analysis*.

Sensitivity Analysis considers the effect that small changes in the LP input parameters would have on the optimal solution and objective value (e.g., Gal, 1995; Gal and Greenberg, 1997). In the NVP context, for instance, it informs us how the profit is affected as prices change. The basic idea is as follows. Let $\mathbf{x}^* \in \mathbb{Q}_+^n$ and $\mathbf{y}^* \in \mathbb{Q}_+^m$ be the optimal primal and dual solutions, respectively. Suppose we “perturb” the LP, by changing \mathbf{c} to $\mathbf{c} + \boldsymbol{\delta}$ and changing \mathbf{b} to $\mathbf{b} + \boldsymbol{\gamma}$. (Here, $\boldsymbol{\delta} \in \mathbb{Q}^n$ and $\boldsymbol{\gamma} \in \mathbb{Q}^m$.) The increase in the optimal profit will be $\boldsymbol{\delta}^T \mathbf{x}^* + \boldsymbol{\gamma}^T \mathbf{y}^*$, provided that the components of $\boldsymbol{\delta}$ and $\boldsymbol{\gamma}$ are sufficiently small. A precise definition of “sufficiently small” can be found in Wendell (1985).

Closely related to sensitivity analysis is *parametric analysis* (e.g., Dantzig, 1963; Vanderbei, 2020). Suppose we have selected a specific constraint in the system $\mathbf{A}\mathbf{x} \leq \mathbf{b}$. For any given real r , let $\phi(r)$ be the optimal profit when the right-hand side of the given constraint is increased by r . Then ϕ is a piecewise-linear concave function of r , and it can be computed using a modified version of the Simplex method (Dantzig, 1963). Note that r is not restricted to taking small values, as in sensitivity analysis. Although parametric analysis requires greater effort than sensitivity analysis, it indeed allows us to study the effect of a larger scale of “perturbation”.

1.4 Stochastic Programming

Many real-world decisions involve uncertainty. For example, imagine a company that sells products to households. Since the demand from the customers is uncertain, the problem of determining the best inventory level cannot be modelled as a simple LP.

This leads naturally to *stochastic programming* (SP), which is an important generalisation of linear programming. A stochastic programme is an optimisation problem in which some or all of the problem parameters are uncertain but follow known probability distributions. The goal of SP is to find a solution which both optimises some criterion chosen by the decision maker and appropriately accounts for the uncertainty of the problem parameters (Birge and Louveaux, 2011; Kall and Wallace, 1994). SP has found applications in a broad range of areas, including NVPs. Other applications of SP include, but are not limited to: portfolio selection, traffic management, and production management.

One of the most well-known SP models is the *2-stage stochastic linear program with recourse* (SLP for short) (Prékopa, 2013). In this model, the variables are classified as “first-stage” or “second-stage” variables. The decision maker begins by choosing values for the first-stage variables. Later on, after more information has become available, the values of the second-stage variables can be selected. The objective is to minimise the cost of the first-stage decision plus the *expected* cost of the second-stage decision.

A general SLP can be written in the form

$$\min_{\mathbf{x} \in \mathbb{R}^n} \left\{ \mathbf{c}^T \mathbf{x} + \mathbb{E}_{\xi} [Q(\mathbf{x}, \xi)] : \mathbf{A} \mathbf{x} = \mathbf{b} \right\}, \quad (1.2)$$

where

$$Q(\mathbf{x}, \xi) = \min_{\mathbf{y} \in \mathbb{R}^m} \left\{ \mathbf{q}(\xi)^T \mathbf{y} : \mathbf{T}(\xi) \mathbf{x} + \mathbf{W}(\xi) \mathbf{y} = \mathbf{h}(\xi) \right\}.$$

Here, \mathbf{x} and \mathbf{y} are the vectors of first- and second-stage variables, respectively. The

important thing to note is that the vectors \mathbf{q} and \mathbf{h} and the matrices \mathbf{T} and \mathbf{W} are functions of a random variable ξ . The realisation of ξ becomes known only after the values of the first-stage variables have already been chosen.

We will see in Chapter 3 that many NVPs of interest, including ones involving multiple product types, can be easily modelled as SLPs. A nice thing about SLPs is that they can be solved (to any desired accuracy) with the help of an LP solver.

1.5 Quantile Regression

Attempts have been made in the literature to explore integrated approaches to NVPs. The most famous one is using *Quantile Regression* (QR). QR is an extension of linear regression, used when the conditions of linear regression are not met (Koenker and Hallock, 2001). Whereas the method of least squares estimates the conditional mean of the response variable across values of the predictor variables, quantile regression estimates the conditional quantiles of the response variable. Moreover, the quantile regression estimates are more robust against outliers in the response measurements. Table 1.1 summarises some important differences between linear regression and quantile regression.

Table 1.1: Differences between linear regression and quantile regression

Linear Regression	Quantile regression
Predicts the conditional mean	Predicts conditional quantile
Works on small samples sizes	Needs sufficient data
Often assumes normality	Distribution agnostic
Sensitive to outliers	Robust to outliers
Computationally inexpensive	Computationally intensive

If we use q_τ to denote the τ th quantile of a known distribution Y , we have

$$Pr(Y \leq q_\tau) = \tau. \quad (1.3)$$

However, if the distribution is unknown, we need to use quantile regression on observations y_t to estimate any given quantile. Defining the loss function as

$$\rho_\tau(y) = y(\tau - \mathbb{I}(y < 0)), \quad (1.4)$$

where $\mathbb{I}(x)$ is an indicator function, taking value 1 if the condition x is met, taking value 0 otherwise, the τ th sample quantile can be obtained by

$$\hat{q}_\tau = \arg \min_{q \in \mathbb{R}} \sum_{t=1}^n \rho_\tau(y_t - q). \quad (1.5)$$

The use of quantile regression can be further extended to include exogenous variables, say \mathbf{x}_t , as features. Supposing that the τ th quantile takes the form $\mathbf{q}_\tau = \mathbf{x}_t^\top \boldsymbol{\beta}_\tau$, we can estimate the conditional quantile by first estimating $\hat{\boldsymbol{\beta}}_\tau$. This can be done in the following way:

$$\hat{\boldsymbol{\beta}}_\tau = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^k} \sum_{t=1}^n \rho_\tau(y_t - \mathbf{x}_t^\top \boldsymbol{\beta}). \quad (1.6)$$

Given that quantile regression is advantageous when conditional quantile functions are of interest, it can be very useful in the study of newsvendor systems (Ban and Rudin, 2019; Huber et al., 2019).

1.6 Judgemental Adjustment

Adjustments are often made by decision-makers after a statistical forecasting procedure has been applied. There exist many papers investigating judgemental adjustments in the context of NVP decision-making (see Lim and O'Connor, 1995; Webby and O'Connor, 1996). The advantage of judgemental adjustments is that they can take into account information which is not included in statistical models, such as

promotions, large sports events, holidays, or recent events that are not yet available in the data.

Many researchers have suggested, however, that practitioners adjust much more often than they should, and many times for the wrong reasons (Schweitzer and Cachon, 2000). Unlike statistical forecasts, which can be generated by the same mathematical formulae every time, judgemental adjustments depend heavily on human cognition and are vulnerable to its limitations (see, for example, Webby and O'Connor, 1996).

On the other hand, there is evidence that the behaviour of decision-makers can be improved through training. In Benzion et al. (2008), the authors recorded that in the first round of their experiment, decision-makers tended to be more “judgemental” than in the last round. They also found that the judgmental order decisions converge to the optimal level slowly as the experiment proceeds. This implies that, by receiving immediate feedback after each round, decision-makers are able to improve their decisions. In the works by Bostian et al. (2008), the effect of training is further studied. They conclude that experience can overwrite pre-conceived biases more effectively than knowledge gained from third-party sources (e.g., information about the demand distribution). Thus, it makes sense to consider the possibility of applying judgemental adjustments in an automated fashion, e.g., using a heuristic to estimate the “optimal” adjustment parameters.

1.7 Structure of the Thesis

The next four chapters of this thesis consist of four journal articles that have been submitted/accepted for publication. Each chapter contains theory, experiments and/or applications which are motivated by aspects of modelling Newsvendor problems.

Chapter 2 has been published as Liu, C., Letchford, A.N. and Svetunkov, I., 2022. Newsvendor problems: An integrated method for estimation and optimisation. *Eur. J. Oper. Res.*, 300(2), 590–601.

Chapter 3 has been submitted for publication as Liu, C., Letchford, A.N. and Svetunkov, I., 2022. On sensitivity and parametric analysis for multi-item newsvendor problems. (EURO Journal on Computational Optimisation.)

Chapter 4 has been submitted for publication as Liu, C., Letchford, A.N. and Svetunkov, I., 2022. Naïve newsvendor adjustments: Are they always detrimental? (Journal of the Operational Research Society.)

Chapter 5 has been submitted for publication as Liu, C. and Zhu, W., 2022. Newsvendor conditional value-at-risk minimisation on non-parametric approach. (Operations Research.)

Finally, *Chapter 6* gives concluding remarks and discusses some potential areas for future research.

Chapter 2

News vendor Problems: An Integrated Method for Estimation and Optimisation

In this chapter, we consider a data-driven method for the classical NVP, proposed recently by Ban and Rudin. We first examine it from a statistical viewpoint and establish a connection with quantile regression. We then extend the approach to the nonlinear NVP. Finally, we give extensive experimental results, on both simulated and real data. The results indicate that the approach performs as well as conventional ones when applied to the classical NVP, but performs better when applied to the nonlinear NVP. There is also evidence that the approach is more robust with respect to model misspecification.

2.1 Introduction

Inventory control is an important topic in Operations Research and Operations Management (see, e.g. Porteus, 2002; Silver et al., 1998; Zipkin, 2000). In this chapter, we focus on *Newsvendor Problems* (NVPs), by which we mean single-period inventory control problems with stochastic demand.

In early work on NVPs (Arrow et al., 1951; Morse and Kimball, 1951), it is assumed that the demand in each time period comes from a known probability distribution. Of course, in practice, this is not the case — a fact already noted by Scarf in 1958. Assuming that historical demand data is available, one can attempt to address this issue by decomposing the problem into an estimation/forecasting phase and an optimisation phase. In the first phase, one makes some assumptions (i.e. specific model form and distributional assumptions) regarding the underlying data-generating process and uses the past data to estimate the parameters of the model. In the second phase, one determines the order quantity (or quantities) based on the estimated parameter values. Throughout this chapter, we will call this two-phase approach the *disjoint* approach.

An advantage of the disjoint approach is that forecasting and optimisation experts can operate independently within an organisation. This makes things easier to manage. On the other hand, as noticed by several authors (Bertsimas and Thiele, 2005; Beutel and Minner, 2012; Karmarkar, 1994; Korpela and Tuominen, 1996), there are two disadvantages:

- The two phases use different objective functions. Indeed, in the first phase, the objective is to minimise a function of the forecasting errors, such as the root mean square error or mean absolute error. In the second phase, however, the goal is usually to maximise the expected profit.

-
- If the forecasting model is misspecified, and/or there is substantial noise in the data, then this might impact the optimisation phase in an unexpected way, possibly leading to sub-optimal solutions. In particular, upside and downside errors may have very different effects on expected profit, due to different costs associated with over- and under-stocking.

An alternative to the disjoint approach is to use a single, *integrated* approach, in which the order quantities are determined directly from the data based on an assumed model or filter. In this case, an adjusted loss function is used, as in *quantile regression* (Bruzda, 2016; Huber et al., 2019) and *SPO* loss (Elmachtoub and Grigas, 2017). The advantages of these approaches are that they do not make assumptions about the demand distribution while remaining explainable. Unfortunately, they can only be applied to relatively simple NVPs, for which the objective function is linear.

Another example of integrated approach involves *machine learning* algorithms (Bertsimas and Kallus, 2020; Bertsimas and Thiele, 2005; Liyanage and Shanthikumar, 2005). With sufficient data, they can build relationships between the optimal order quantity and exogenous variables, sidestepping the need to have two phases. This approach is called *Feature-Based NVP* in He et al. (2012).

Most Feature-Based NVP approaches are “black box” approaches, which are hard to interpret or explain. A more transparent Feature-Based NVP approach was proposed recently by Ban and Rudin in 2019. In their approach, statistical parameters are selected in a way that directly attempts to minimise the expected opportunity costs.

In this chapter, we consider the approach in Ban and Rudin in more detail. We begin by examining it from a statistical viewpoint, and establish a connection with quantile regression. We then extend the approach to nonlinear NVP, leading to

what we call an *Integrated Method for Estimation and Optimisation* (IMEO). We also provide extensive simulation experiments to examine the performance of IMEO in different settings, including situations when the true model is known and when the underlying model is mis-specified.

The rest of the chapter is organised as follows. Section 2.2 is a literature review. Section 2.3 presents the IMEO framework and the theoretical analysis. Section 2.4 reports experimental results on simulated demand data. In Section 2.5, we apply our approach to some real-life data. Some concluding remarks are made in Section 2.6.

2.2 Literature Review

Since the literature on NVPs is vast, we mention here only works of direct relevance. In Subsections 2.2.1 and 2.2.2, we review the classical NVP and its extensions, respectively. Subsection 2.2.3 and 2.2.4 review quantile regression and the Ban and Rudin method, respectively.

2.2.1 The classical newsvendor problem

In the simplest NVP, as defined, for example, by Choi in 2012, a company purchases goods at the beginning of a time period at a cost of v per unit, and aims to sell them by the end of the period at a price p per unit. The demand during the period is a random variable Y with known probability density function f and cumulative distribution function F . At the end of the period, any surplus goods will lead to a *disposal cost* of c_h per unit. On the other hand, a shortage of goods during the period will lead to a *shortage cost* of c_s per unit. The goal is to determine an *order quantity* Q , prior to the period, that maximises the expected profit.

For a given Q and a given realisation y of Y , the profit over the period is:

$$\pi(Q, y) = \begin{cases} py - vQ - c_h(Q - y), & \text{if } Q \geq y \\ pQ - vQ - c_s(y - Q), & \text{if } Q < y. \end{cases} \quad (2.1)$$

The expected value of $\pi(Q, y)$ is:

$$\Pi(Q) = \int_0^Q [py - vQ - c_h(Q - y)] f(y) dy + \int_Q^\infty [pQ - vQ - c_s(y - Q)] f(y) dy. \quad (2.2)$$

It is common to call $c_u = p - v + c_s$ the ‘underage’ cost and $c_o = v + c_h$ the ‘overage’ cost. Some calculus then shows that the order quantity that maximises $\Pi(Q)$ is (Choi, 2012):

$$Q^* = F^{-1} \left(\frac{c_u}{c_o + c_u} \right), \quad (2.3)$$

where F^{-1} is the inverse function of F . Thus, Q^* is the τ^{th} quantile of f , with $\tau = c_u/(c_o + c_u)$. One can think of the quantity τ as a “target service level”, since aiming for this target will bring the company maximised expected profit.

2.2.2 More complex newsvendor problems

Since the introduction of the NVP by Arrow et al. in 1951, researchers have considered several extensions of the problem, including variants with multiple product types (Hadley and Whitin, 1963; Lau and Lau, 1996; Moon and Silver, 2000), quantity discounts (Khouja, 1995), different risk measures (Eeckhoudt et al., 1995), product substitution (Bassok et al., 1999), nonlinear cost functions (Halman et al., 2012), non-stationary demand (Kim et al., 2015), and price setting (Karlin and Carr, 1962; Mills, 1959; Petruzzi and Dada, 1999).

For the purpose of what follows, we now explain one variant, the ‘Nonlinear Newsvendor Problem’ (NNVP), which can be found in works of Halman et al. (2012), Khouja (1995), Kyparisis and Koulamas (2018) and Pal et al. (2015) and many

others. In the general NNVP, the profit function takes the form:

$$\pi(Q, y) = \begin{cases} P(Q, y) - V(Q) - C_h(Q, y), & \text{for } Q \geq y \\ P(Q, y) - V(Q) - C_s(Q, y), & \text{for } Q < y, \end{cases} \quad (2.4)$$

where V , P , C_h and C_s are now *functions* rather than constants.

Using the NNVP enables one to model more real-life problems. Indeed, as noticed by Pantumsinchai and Knowles in 1991 and Khouja in 1995, non-linear costs arise frequently in practice. For example, in real life, a minor shortage may not cause large costs, but a major shortage could damage the reputation of the company. As another example, a small amount of excess stock can often be sold at a discount, but this may not be possible for a large amount, not to mention the loss of goodwill. Moreover, as a product stays on the shelf longer, the opportunity cost of the shelf space may increase over time. All these examples show that overage and underage costs can be non-linear in real-life situations.

If $\pi(Q, y)$ has a particularly simple form (e.g., if it is piecewise-linear as in the classical NVP), then it may be possible to use calculus to express the optimal order quantity as a quantile (Choi, 2012). In general, however, a closed-form expression as a quantile is unlikely to exist (Halman et al., 2012; Porteus, 2002). In such cases, one must resort to numerical integration and search techniques to solve the NVP (Solis and Wets, 1981).

We now recall one specific NNVP, taken from Kyparisis and Koulamas in 2018, that we are going to use in our experiments later in this chapter. As second-round sales in salvage markets and proportional shortage penalties are very common in real life (Kashefi, 2016; Liberopoulos et al., 2010), it is particularly necessary to check the performance of the proposed method on this NNVP. The purchase cost v and selling price p are constants, but C_h and C_s are functions. Overstock items incur a constant unit penalty $\alpha > 0$, but they can be sold in a salvage market with a fixed unit sales price β , with $0 < \beta < v$. The demand in the salvage market is itself a random

variable, with a known distribution, which we denote by u . That is, we have:

$$C_h(Q, y) = \alpha[Q - y]^+ - \beta \mathbb{E} \left[\min \{ [Q - y]^+, u \} \right]. \quad (2.5)$$

Moreover, the shortage penalty is proportional to the shortage quantity. That is:

$$C_s(Q, y) = \zeta ([y - Q]^+)^2 \quad (2.6)$$

for some constant $\zeta > 0$. This problem cannot be solved analytically, but there are known approximation methods that give adequate solutions (Kyparisis and Koullamas, 2018).

2.2.3 Quantile regression

Returning to the classical NVP, we now consider the (more realistic) case in which the demand distribution is unknown, but we have historical demands y_1, y_2, \dots, y_s . For this case, *quantile regression* has proven to perform well, and the basic idea is as follows (Bertsimas and Thiele, 2005):

1. Compute the value of τ that maximises expected profit;
2. Use quantile regression to compute an estimate of the τ^{th} quantile of the demand in the next time period, which we denote by $\hat{y}_{s+1}^{(\tau)}$;
3. Set the order quantity \hat{Q}_{s+1} equal to $\hat{y}_{s+1}^{(\tau)}$.

The biggest advantage of quantile regression is that it does not assume a specific demand distribution (Huber et al., 2019). However, it is efficient only on large samples (Ban and Rudin, 2019; Huber et al., 2019). Another drawback is that its performance depends crucially on the underlying target service level. Ban and Rudin in 2019 demonstrated that the benefit of using quantile regression is limited to target service levels smaller than 0.8. If the target service level is higher, then much

more data is needed in order to correctly estimate a specific quantile. Moreover, as mentioned in the previous subsection, for the more complex NVPs, there is no easy way to express the optimal order quantity as a quantile.

2.2.4 The Ban and Rudin approach

In the approach of Ban and Rudin, a statistical model is built, in which exogenous variables are regressed against the order quantity. In more detail, we have historical data $[(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_s, y_s)]$, where $\mathbf{x}_t = [x_t^1, \dots, x_t^p]$ represents features related to the demand, such as seasonality, price, promotions and so on. The problem now becomes that of finding the optimal function $q(\cdot)$ that maps the observed features $\mathbf{x}_{s+1} \in \mathcal{X}$ to an order $q(\mathbf{x}_{s+1}) \in \mathbb{R}$. Ban and Rudin propose to select $q(\cdot)$ from a pre-specified family of functions, in a way that minimises the empirical opportunity cost:

$$\min \sum_{t=1}^s (c_u[y_t - q(\mathbf{x}_t)]^+ + c_o[q(\mathbf{x}_t) - y_t]^+), \quad (2.7)$$

The simplest version of their method uses linear functions, of the form:

$$q(\mathbf{x}_t) = \mathbf{x}_t^\top \boldsymbol{\beta} = \sum_{j=1}^p x_t^j \beta^j, \quad (2.8)$$

With this choice, one can determine the β^j easily with linear programming. Ban and Rudin also mention more complex variants, that use polynomials of the original features, and/or a quadratic regularisation term.

2.3 Analysing and Extending the Ban and Rudin Method

In this section, we analyse the approach of Ban and Rudin in more depth, and extend it to a more general family of problems.

2.3.1 Analysis

To begin, we attempt to provide an intuition behind the method. Suppose that, for each historical period $t \in [1, s]$, the observed demand y_t is a realisation of a random variable Y_t . Then, in principle, there exists an order quantity, say Q_t^* , that maximises the expected profit given Y_t and the function Π . Thus, if we had set Q_t to Q_t^* prior to observing the true demand y_t , we would have maximised our expected profit in period t . Putting it another way, if we could somehow uncover the structure of the unobservable time series of orders $\{Q_1^*, \dots, Q_s^*\}$, we would be able to estimate Q_{s+1}^* directly.

Of course, in practice, the distribution of demand Y_t is unknown, and the values Q_t^* are not observable. So the method approximates the Q_t^* values with the $q(\mathbf{x}_t)$ values. One, therefore, solves an optimisation problem to find the choice of $q(\cdot)$ that minimises the expected opportunity cost. Once that is done, the order quantity for the next time period can be set to $\hat{Q}_{s+1} = q(\mathbf{x}_{s+1})$.

Ban and Rudin did not give a statistical analysis of the behaviour of the estimates of the model parameters that arise from their method. A partial answer is that, in the case of a linear profit function, their method is equivalent to quantile regression (see Appendix A for proof). Therefore, in the linear case, their method immediately inherits all of the desirable properties of quantile regression, such as consistency (Koenker, 2005), efficiency (Koenker and Machado, 1999) and asymptotic normality (Kocherginsky et al., 2005) of the estimates of parameters.

2.3.2 Extension to the nonlinear case

We now consider how to extend the approach in Ban and Rudin to the NNVP. The key issue here is that there is no simple formula for the opportunity cost in the

nonlinear case. Indeed, in the literature on the NNVP, authors work directly with the function (2.4), rather than attempting to derive explicit functions for c_o and c_u .

To get around this difficulty, we propose to maximise the expected profit instead of minimising the expected opportunity cost (2.7). More precisely, we propose to compute the function $q(\cdot)$ that maximises the function

$$\max \sum_{t=1}^s \pi(q(\mathbf{x}_t), y_t), \quad (2.9)$$

where π can be a profit function of any level of complexity.

Maximising (2.9) is a continuous nonlinear optimisation problem. Under reasonable assumptions on the functions P , V , C_h and C_s in (2.4), and the function $q(\cdot)$ itself, the profit function (2.9) will be concave. Unfortunately, it is unlikely to be everywhere differentiable. As a result, general-purpose algorithms for nonlinear optimisation are not guaranteed to converge to global maxima. Fortunately, the experiments in the next section indicate that this does not cause serious problems.

In what follows, we call our approach an *Integrated Method for Estimation and Optimisation* or IMEO. We remark that, in the case of a linear profit function, IMEO is equivalent to the method of Ban and Rudin (see Appendix B for proof). Thus, in that case, it is again equivalent to quantile regression.

2.4 Computational Experiments

In order to assess the performance of IMEO and to understand its strengths and weakness, we conduct a simulation experiment. We start the discussion with Subsection 2.4.2, where the simplest case is studied, in which the profit function is linear and the underlying data generation process (DGP) is known. The case in which the profit function is nonlinear is discussed in Subsection 2.4.2. Finally, we discuss more complicated scenarios, with misspecified models, in Subsections 2.4.3

and 2.4.3. Given the discussion in Section 2.3, the analysis in Subsections 2.4.2 and 2.4.3 will also tell us how the “Feature-based NVP” approach by Ban and Rudin compares with the conventional approaches in different situations.

2.4.1 Experimental setup

In our experiments, we consider NVPs with quarterly demand data, and we generate data from a seasonal ARIMA(1, 0, 0)(1, 0, 0)₄ process with $\theta = 0.3$, $\Theta = 0.5$, and constant level 500. We also assume that the error term of the DGP follows the normal distribution $\mathcal{N}(0, 200^2)$. We choose a seasonal ARIMA model since it has been shown working well in the field of inventory management, and has been properly studied in the literature (for example, see the review paper on supply chain forecasting by Syntetos et al., 2016). We have also experimented with other models and parameters for the DGP. The scripts of extended experiments have been made available on Github (Liu, 2020). The proposed method showed strong robustness and the results were very similar to the ones presented below.

For each of the cases, we simulate 20,000 sets of demand data, each consisting of 4800 observations. From each set, we extract sub-sequences of lengths 40, 120, 480 and 1200. These are used to explore how the amount of available data affects the performance of each method. The first three data lengths can help us to simulate real-life circumstances, while the data lengths of 1200 and 4800 allow us to explore the asymptotic behaviour of the approach. For each set of demand data and each method, we compute \hat{Q}_{s+1} , the 1-step ahead forecast of orders. We then compute the following three quantities (and aggregate them using a mean over 20,000 sets):

1. Percentage Profit Loss: $PPL = \frac{\pi(y_{s+1}, y_{s+1}) - \pi(\hat{Q}_{s+1}, y_{s+1})}{\pi(y_{s+1}, y_{s+1})}$, which shows the percentage of profit that would be lost due to using each method instead of knowing the true demand. In the ideal situation, $MPPL$ (mean PPL) should be

equal to zero.

2. Service Level: $SL = \frac{\sum \mathbb{I}(\hat{Q}_{s+1} > y_{s+1})}{20,000}$, where $\mathbb{I}(\cdot)$ is the indicator function, equal to 1 if the condition inside it is satisfied. This measure shows the achieved service level. In the ideal situation, SL should correspond to the target service level.
3. Fill Rate: $FR = \frac{\min\{\hat{Q}_{s+1}, y_{s+1}\}}{y_{s+1}}$, which shows how the demand is serviced. In the ideal situation, MFR (mean FR) should be equal to one.

In the proposed method (“IMEO”), we use the Limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm (L-BFGS) for the estimation of the parameters of the model. The L-BFGS algorithm is very popular in the nonlinear programming community (see Liu and Nocedal, 1989), and is commonly used for ARIMA fitting in software packages, for example in R (R Core Team, 2020).

Besides the proposed method, three benchmark methods are also considered:

- A disjoint method (“DJ”) that estimates the parameters of the model in the first phase, and then determines the optimal order quantity in the second phase.
- An integrated method that uses Quantile Regression (“QR”) in order to determine the order quantity. We choose it as a rival since it is one of the most widely used statistical approaches, which has proved to work well for the NVP.
- Finally, for the purposes of benchmarking, we use DJ with the exact model and parameters from the DGP to perform a forecast and determine the order quantity. This last method is “idealised”, since, in real life, one would not know the true model or true parameters of demand. But it allows us to see, how far approaches are from the ideal one. We call this last method “DGP”.

2.4.2 When the true model is known

In this scenario, we assume that it is known that the true model is ARIMA(1, 0, 0)-(1, 0, 0)₄ with constant, but its parameters need to be estimated. In all the methods, it is also assumed that the error term follows a normal distribution with unknown variance.

Linear case

We conduct several experiments with a combination of costs that give service levels of 0.3, 0.5, 0.63 and 0.9. It is important to recall that the optimal service level is a critical solution to the optimisation problem. In most cases, the most profitable service level is not a “high” stocking level. Other levels can also be considered, but in additional experiments that we have conducted, we have not found any significant changes in the performance of the approaches.


Initially, we choose the following parameters for our (linear) profit function:

- $p = 20$, $v = 10$, $c_h = -3$, $c_s = -7$.

One can check that the corresponding pair (c_u, c_o) is $(3, 7)$, and the target service level evaluates to 0.3. Applying the methods to this scenario, we get values of PPL, SL and FR.

Figure 2.1 shows the mean percentage profit loss obtained when the target service level was set to 0.3. As one would hope, the losses for all methods converge as the sample size increases. It is also apparent that IMEO and QR have very similar performances. This is to be expected, given that IMEO can be considered equivalent to QR in the linear NVP case. (The small gap between those two methods may be due to different optimisation algorithms that are applied, as mentioned in Section

Figure 2.1: Mean percentage profit loss vs. data size at 0.3 target service level



2.3.) Another thing to note is that the integrated approaches perform very similarly to the disjoint one, which is also expected, given the knowledge of the true model and the simplicity of the NVP.

Figure 2.2: Service level vs. sample size at 0.3 target service level




Figure 2.2 represents the service level, i.e., the proportion of iterations in which the demand was satisfied in the simulation. We see that all four methods converge to the desired target of 0.3 as the sample size grows. Interestingly, DJ approaches the target from below, while IMEO and QR slightly overestimate the level on small samples. A possible explanation of this phenomenon is that in the first phase of

DJ, the estimated parameters tend to be closer to zero when the sample size is small. This causes the estimated order quantities to be further away from the mean than they should be and leads to lower-than-needed SL. On the other hand, the integrated methods take the underage and overage costs into account. For the given cost parameters, over-stocking is less costly than under-stocking, which leads to a higher SL than needed.

Figure 2.3: Mean fill rate vs. sample size at 0.3 target service level




Figure 2.3 shows the mean fill rate. As one might expect, the performance of all methods improves with the growth of the sample size. The thing to note is that the DJ requires a much larger sample size than the other methods in order to achieve the same MFR, although the differences between the DJ and other methods are not very big (only a couple of percentage points). The possible explanation for this behaviour of DJ is probably similar to the situation with the service levels.

Next, we explore whether the target service level has a significant effect on the performance of the methods. Specifically, we consider the following three alternative parameter settings:

- $p = 20, v = 8, c_h = -3, c_s = -7$

- $p = 20, v = 8, c_h = 3, c_s = 7$
- $p = 20, v = 8, c_h = -7, c_s = -3.$

These correspond to pairs (c_u, c_o) are $(5, 5)$, $(19, 11)$ and $(9, 1)$, respectively, leading to the target service levels of 0.5, 0.63 and 0.9, respectively. It is important to recall that the target service level refers to the most profitable service level for given parameters.

The relevant data is given in Tables 2.1 and 2.2, for the cases of sample sizes $s = 40$ and $s = 4800$, respectively. The best cases (excluding DGP, which is expected to be the best by construct) are marked in boldface for each of the error measures and each of the service levels.

Table 2.1: Target service level effect with $s = 40$

Target service level	Mean percentage profit loss				Service level				Mean fill rate			
	DGP	DJ	QR	IMEO	DGP	DJ	QR	IMEO	DGP	DJ	QR	IMEO
0.3	5.2%	5.6%	5.7%	5.6%	0.30	0.26	0.32	0.32	91.1%	88.2%	90.6%	90.8%
0.5	4.9%	5.3%	5.3%	5.2%	0.50	0.44	0.50	0.50	95.2%	93.4%	94.7%	94.8%
0.63	13.8%	15.1%	15.0%	14.8%	0.63	0.58	0.62	0.62	97.1%	95.8%	96.6%	96.6%
0.9	2.1%	2.4%	2.4%	2.3%	0.90	0.91	0.91	0.90	99.6%	99.1%	99.2%	99.4%

We can see from Table 2.1 that IMEO performs consistently better than the other approaches on the small sample, regardless of the target service level, although the difference in performance between the methods is not substantial. QR is the second-best approach in this scenario.

When it comes to large samples (Table 2.2), IMEO and QR perform slightly worse than DJ, although the difference between the methods is not substantial. This holds for all target service levels, irrespective of the error measures. A thing to note is that the MPPL is particularly high for all methods when the target service level is 0.63. This is probably due to the relatively large overage and underage costs in

Table 2.2: Target service level effect with $s = 4800$

Target service level	Mean percentage profit loss				Service level				Mean fill rate			
	DGP	DJ	QR	IMEO	DGP	DJ	QR	IMEO	DGP	DJ	QR	IMEO
0.3	5.1%	5.1%	5.2%	5.2%	0.30	0.30	0.30	0.30	91.1%	91.0%	91.0%	91.1%
0.5	4.9%	4.9%	5.0%	5.0%	0.50	0.50	0.50	0.50	95.3%	95.2%	95.2%	95.2%
0.63	13.8%	13.8%	14.0%	14.0%	0.63	0.63	0.63	0.63	97.0%	97.0%	97.0%	96.9%
0.9	2.1%	2.1%	2.2%	2.1%	0.90	0.90	0.90	0.90	99.5%	99.2%	99.1%	99.4%

that case. In addition, note that, asymptotically, all methods approach the target service level. This makes IMEO and QR more desirable than DJ, because they do at least as well as DJ both on small and large samples.

Summarising, we can see that IMEO does at least as well as the classical disjoint method and quantile regression in different scenarios of the linear NVP when the model is correctly specified.

Nonlinear profit function

In this subsection, we examine the relative performance of the methods when applied to the NNVP. As before, we assume that the true model for the demand is known. We use the following nonlinear profit function (Kyparisis and Koulamas, 2018):

$$\pi(Q, y) = \begin{cases} 20y - 8Q - 4(Q - y) + 5 \mathbb{E}[\min\{(Q - y), u\}], & \text{if } Q \geq y \\ 20Q - 8Q - 0.01(y - Q)^2, & \text{if } Q < y, \end{cases} \quad (2.10)$$


where $u \sim \mathcal{N}(30, 5^2)$. Given that the target service level does not have a closed form in the NNVP, one can use the technique proposed by Kyparisis and Koulamas, or other numerical approaches, to verify that it is approximately equal to 0.56.

Since the nonlinearity makes it impossible to apply QR (as discussed in Subsection 2.2.3), we present results only for DJ and IMEO, together with the “idealised” method based on the DGP. Moreover, we found that the MFR plots were similar in

all scenarios and did not give any additional important information. Therefore, we do not present them in what follows.

Figure 2.4 shows the MPPL and SL of each method, for each sample size. It is apparent that both DJ and IMEO perform well in terms of MPPL, with very similar values for all sample sizes. IMEO is slightly better on small samples (similar to what we have seen in Subsection 2.4.2), but the differences in performance between the methods are not substantial. When it comes to SL, the picture is similar to what we have observed in Subsection 2.4.2: the SL of IMEO is very close to the target service level even when the sample size is small, while the disjoint method needs more data to reach the target service level, and it approaches it from below. The reason for such performance is similar to the one discussed in Subsection 2.4.2. As expected, both methods converge to the DGP values in both MPPL and SL with the increase of the sample size.

Figure 2.4: Performance vs. sample size with nonlinear profit function



(a) percentage profit loss vs. data size

(b) service level vs. data size

We would like to stress that, unlike the disjoint method, IMEO does not need any complicated numerical optimisation or simulation methods to estimate the optimal order quantity – it does that directly. In addition, we have found in our experiments that it typically required less computational time than the disjoint method: approximately only one-tenth of the time spent on DJ was needed for IMEO when $s = 40$.

Overall, we see that IMEO performs at least as well as DJ when the true model is known. However, the true model is typically not known in practice, so next, we investigate situations, when the model is misspecified.

2.4.3 When the true model is not known

We examine the effect of *model misspecification* on the relative performance of the various methods. We consider three scenarios:

1. Model omits important variables, which typically leads to biased estimates of parameters;
2. Model has redundant variables, which usually leads to inefficient estimates of parameters;
3. The assumed distribution of the error term is wrong.

While in reality there can be other scenarios, the proposed three scenarios cover the main possible issues with the model misspecification.


We don't include the QR method in further discussions, because it performs similarly to IMEO in the linear case, and cannot be used in the nonlinear case.

Linear case

We first consider the situation in which **the model omits important variables** (i.e., is under-parameterised): the seasonal order of ARIMA is dropped in the estimation and the solution of the NVP. In other words, when estimating the optimal order quantity, we apply the AR(1) model with a constant.

The MPPL and SL for this case are shown in Figure 2.5. By comparing with Figure

Figure 2.5: Performance vs. sample size with under-parameterised linear model



(a) percentage profit loss vs. data size


(b) service level vs. data size

2.1, we see that the use of an under-parameterised model causes both DJ and IMEO to incur small additional losses in MPPL, irrespective of the sample size. This is expected, as the seasonal pattern in the data is ignored by the models. These losses do not significantly improve as the number of data points increases, because seasonality plays a key role when making one-step-ahead forecasts. When it comes to the SL, both approaches perform similarly, reaching the target on larger samples. But when it comes to small samples, they both reach higher than needed levels. Interestingly, the performance of DJ and IMEO is similar, even though they are very different approaches.

Next, we consider the situation in which **the model contains redundant variables** (i.e., is over-parameterised): an unnecessary lag term is included. Specifically, we apply an $ARIMA(2, 0, 0)(1, 0, 0)_4$ model to the data.

The MPPL and SL for this case can be seen in Figure 2.6. Both DJ and IMEO incur a loss in profit, as before. Now, however, the profit loss decreases with the increase in the sample size. This can be explained by a well-known statistical phenomenon when estimating models with redundant variables (Farrar and Glauber, 1967): they tend to lead to less efficient estimates of parameters on small samples but do not lead to systematic bias (as omitted variables typically do). As for the service level, it can be seen that both methods asymptotically converge to the target, but tend to

Figure 2.6: Performance vs. sample size with over-parameterised linear model



(a) percentage profit loss vs. data size

(b) service level vs. data size

be less precise on small samples, with IMEO reaching higher levels than needed and DJ leading to lower values. The performance is similar to what we saw in Subsection 2.4.2.


We do not present results for other target service levels in this part, since they were very similar to what we observed above.

Finally, we explore the effect on the performance of different methods, when **the assumed distribution is incorrect**. In particular, we generate data using a modified version of our seasonal $ARIMA(1, 0, 0)(1, 0, 0)_4$ model, in which the *error term* follows the Laplace distribution with mean 0 and scale 141 (which will have a standard deviation of 200). We remark that the Laplace distribution has “fatter tails”, or, more formally, a higher kurtosis than the normal distribution. When estimating the optimal order quantity, however, we use the incorrect assumption that the error term follows the normal distribution.

The results for this scenario are presented in Figure 2.7. We can see that the incorrect distributional assumption leads only to a very small loss in profit for both DJ and IMEO. They both converge to DGP as the sample size increases. However, the analysis of SL shows that while IMEO rapidly converges to the target service level from above, the DJ achieves a lower-than-needed service level, producing a biased value – and this drawback is not remedied by an increase in the sample

size. A possible explanation is that the integrated approach works directly with the data, and does not rely on the assumption of normality, being in this sense “non-parametric”, while the DJ relies on normality and consistently underestimates the uncertainty in the data. This example suggests that IMEO may be more robust.

Figure 2.7: Performance vs. sample size with Laplace distributed error term




(a) percentage profit loss vs. data size

(b) service level vs. data size

Nonlinear case

Finally, we explore the case in which there is a nonlinear profit function and misspecification simultaneously. The same three scenarios of misspecification are considered.

Figure 2.8: Performance vs. sample size with under-parameterised nonlinear model



(a) percentage profit loss vs. data size

(b) service level vs. data size

Figure 2.8 represents the scenario in which applied models are under-parameterised. We can see that the performance of the methods in terms of MPPL and SL are

similar to the case of linear NVP discussed in Subsection 2.4.3. The percentage profit loss for both DJ and IMEO stabilises around 14.75% and never converges to the level of DGP, due to the absence of the important variable in the model. This is compensated by SL, for which DJ performs similarly to DGP, with IMEO reaching a slightly higher-than-needed level.

Figure 2.9: Performance vs. sample size with over-parameterised nonlinear model

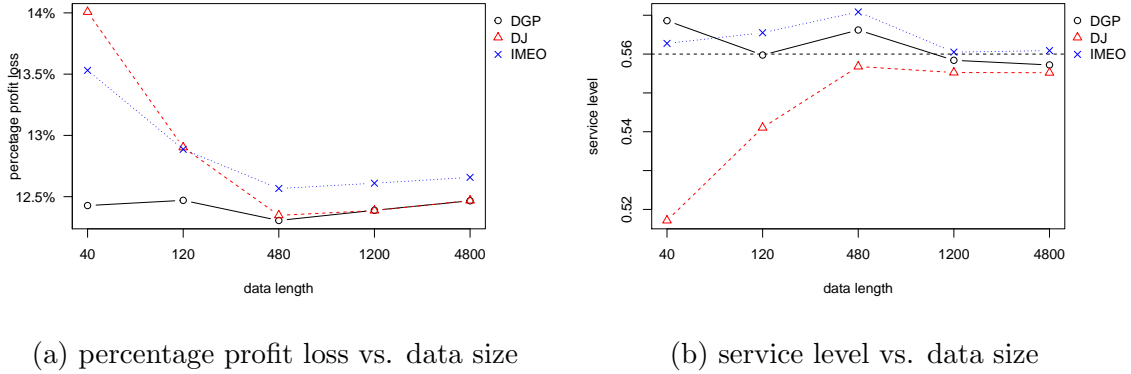



Figure 2.9 demonstrates the over-parameterised scenario. We can see that IMEO has lower MPPL than DJ on small samples. With an increase in sample size, the latter method converges to the DGP, while IMEO still has a slight bias, producing around 0.1% higher loss than the DGP. The good performance of IMEO on small samples could be because it does not need to estimate the variance of the error term. As for the SL, IMEO reaches the target level much faster than DJ, performing especially well on small samples, where the latter method reaches a much lower service level than needed. As the sample size increases, both methods converge to the target level.

Finally, we consider the scenario in which an incorrect distribution of error term is assumed. The results are presented in Figure 2.10. It becomes apparent that IMEO outperforms DJ in terms of MPPL on small samples and converges to DGP together with DJ. IMEO also seems more stable than DJ in terms of SL across all sample sizes. While IMEO does not converge to the target level on larger samples, it is consistent and less biased than DJ, which converges to a value much higher than

needed. One interesting thing we can see from both Figure 2.10 and Figure 2.7 is that DJ always gives a very low SL when the sample size is small and converges to a higher than needed target as the sample size grows. This “wrong” target is obtained in both situations, not surprisingly, because DJ cannot overcome the incorrectness in the distributional assumption.

Summarising this subsection, we see that IMEO is robust and does not fail as badly as the classical disjoint method does in severe cases of misspecification. At worst, IMEO performs similarly to DJ. In addition, it looks like IMEO does consistently better than the DJ in terms of service level, so if this is more important for a company than profit loss, then we would recommend using IMEO.

Figure 2.10: Performance vs. sample size with Laplace distributed error term, nonlinear model



(a) percentage profit loss vs. data size

(b) service level vs. data size

2.5 Real-life Case

In this section, we examine the performance of IMEO on a real-life nurse staffing problem in a hospital, which must determine the next-day staffing level Q . The hospital incurs an underage cost (unexpectedly high death rate, reputation damage, etc.) if there are not enough nurses, and an overage cost (unnecessary exposure risk, big salary payment, etc.) if there are too many nurses. Both types of costs

are considered to be nonlinear in this case (Al Thobaity and Alshammari, 2020; Fernandez et al., 2020). The objective is to minimise the expected daily cost.

The data we use comes from the NHS open data set (*NHS Statistics on COVID-19 hospital activity* 2020). It includes the total bed occupancies for a large UK general hospital from April to October 2020 on a daily basis.

We assume a fixed 1 to 3 nurse-to-bed ratio (similar to how it was done in Ban and Rudin, 2019), hence the demand y is the total number of beds occupied divided by 3. In addition, we do not require the number of nurses to be an integer, due to the possibility of them working part-time. Based on the study of Chen et al. in 2020 and Liu et al. in 2020, we can approximate this problem as:


$$\pi(Q, y) = \alpha[Q - y]^+ - \beta \mathbb{E} \left[\min \{ [Q - y]^+, u \} \right] + \zeta ([y - Q]^+)^2. \quad (2.11)$$

Here, we assume that each over-scheduled nurse incurs a fixed cost $\alpha > 0$, but can be reassigned to help other departments and reduce the cost by the amount β ; whereas demand of other under-staffed departments can be seen as a random variable u . Moreover, we assume that the shortage penalty of nurse staffing is proportional to the shortage quantity with rate ζ , since that shortage may be covered by reassignment or overcharging, but a significant shortage could be lethal to patients. The parameter values are chosen based on the studies of Chen et al. (2020), Coronini-Cronberg et al. (2020) and Liu et al. (2020):

- $\alpha = 10$, $\beta = 4$, $\zeta = 1$, and $u \sim U(0, 15)$.

To get some sense of the data, we provide a time-series plot in Figure 2.11. It can be seen that from mid-April to October, the number of beds occupied exhibits seasonality and that from late August, the number becomes stable. According to the public information in the NHS data set, the hospital nearly reaches its maximum capacity.

Figure 2.11: Time-series plot of bed occupancy



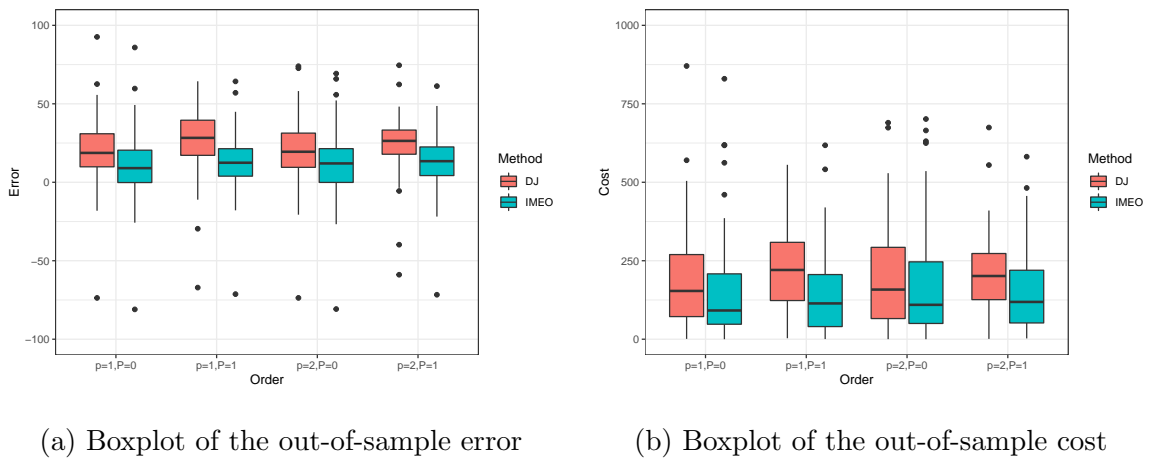
To perform a fair comparison, we apply both the disjoint method (“DJ”) and the proposed method (“IMEO”) to this NNVP, with several ARIMA models with orders $p = 1, 2$ and $P = 0, 1$. We use the ARIMA model here to maintain consistency with previous sections. It is, of course, possible that other models would fit better than the chosen ARIMA. However, our intention here is not to find the best-fitting model, but to compare the performance of IMEO and the disjoint method fairly. After all, it is more interesting to see how the two approaches compare when the applied model is wrong.

For the purpose of generality, we include four pairs of orders in ARIMA ($p = 1, P = 0$; $p = 2, P = 0$; $p = 1, P = 1$ and $p = 2, P = 1$). To compare the performance of the methods, we obtain their 1-step ahead forecasts with rolling horizon (Tashman, 2000), where origin length is $s = 100$ and the origin is shifted $n = 80$ times. For each forecasted value, we compute the over-scheduled/under-scheduled nurse number, the service level, and the daily cost.

The boxplots are shown in Figures 2.12a and 2.12b. The black lines in the boxes represent mean values rather than medians. From the plots, we can see that IMEO


outperforms DJ with all ARIMA models, in terms of both the mean cost and the mean scheduling error. Moreover, DJ produces a larger variance of both errors and costs than IMEO. Thus, IMEO not only has lower mean costs but also works more efficiently overall. This finding is in line with the MPPL result at data length 120 in Subsection 2.4.2.

Figure 2.12: Boxplot of the out-of-sample performance. The black lines in the boxes represent mean values



We also used the `auto.arima()` function from forecast package (Hyndman et al., 2020) for R in order to select the most appropriate ARIMA model for the DJ method and then we used the same model for the IMEO (Figures 2.13a and 2.13b). In this case, we are favouring DJ approach, producing the model closer to the true one based on Akaike’s information criterion. At the same time, the appropriate order selection mechanism for IMEO is not yet developed, so the same model will not be optimal for it. Still, analysing the plots in Figure 2.13, we can see that the difference between the two methods is not substantial, with boxplots being very close to each other. This additional experiment shows that even when the DJ approach is done properly, using state-of-the-art forecasting techniques, IMEO does not fail substantially and can be considered a decent alternative to DJ. Together with the results from Figures 2.12, this experiment shows that IMEO works well in a wide variety of cases.

Figure 2.13: Boxplot of the out-of-sample performance with auto-fitting. The black lines in the boxes represent mean values



(a) Boxplot of the out-of-sample error with auto-fitting (b) Boxplot of the out-of-sample cost with auto-fitting

The service level achieved by each method is summarised in Table 2.3. One can check by simulation or numerical methods that the “target service level” that minimises cost for the given parameters is around 0.8. In that case, IMEO achieves a much closer service level to the target than the disjoint method, no matter what ARIMA model is used, while the disjoint method provides a higher-than-needed service level in all four cases.

Table 2.3: Service level of each method for real-life case

Order	$p = 1, P = 0$	$p = 2, P = 0$	$p = 1, P = 1$	$p = 2, P = 1$
DJ	0.950	0.950	0.975	0.975
IMEO	0.750	0.750	0.825	0.850

Note: The service level achieved by an auto-fitting algorithm is 0.950.

This example shows that IMEO is a robust approach that results in lower costs and a service level closer to the target, even when the model is specified incorrectly. The disjoint method is much more sensitive to the specification of the model, performing poorly when the model is misspecified.

2.6 Concluding Remarks

In this chapter, we extended the method of Ban and Rudin (2019) to a broader framework, called “Integrated Method for Estimation and Optimisation” or IMEO. IMEO attempts to maximise the expected profit instead of minimising the expected opportunity cost, which turns out to be an important distinction in the nonlinear case. We showed that IMEO reduces to the method of Ban and Rudin (2019) in the linear case when both methods turn out to be equivalent to quantile regression. Our experiments indicate that IMEO performs at least as well as the benchmark methods, in terms of both mean percentage profit loss and service level. It also appears that IMEO is more robust with regard to model misspecification. It also does well in terms of service level, which could be attractive in real-life applications.

While the focus of the experiments in this chapter was on ARIMA models, the proposed approach could be applied to other models as well, such as linear regression with explanatory variables, nonlinear regression and ETS.

There are several interesting topics for further research. First, it would be interesting to study the performance of IMEO with other demand models, such as ETS. Second, it would be desirable to develop a variable selection mechanism in IMEO. The conventional disjoint method allows one to do this in the first phase, for example by using cross-validation or a stepwise technique based on information criteria, while Ban and Rudin (2019) uses regularisation for the selection and estimation. While these are good approaches, they require large samples and are computationally expensive. Our hope is that a more efficient feature selection method can be developed. Third, we focused our research on NVP, but IMEO could be potentially extended to multi-period inventory problems. Finally, it would be interesting to extend IMEO to multi-item problems, either with or without substitution effects between products.

Chapter 3

On Sensitivity and Parametric Analysis for Multi-Item Newsvendor Problems

The *Multi-item Newsvendor Problem* (MNP) is the extension of the NVP in which there are several product types, sharing one or more resources (such as warehouse space). Although the MNP has received much attention from the Operational Research community, not much work has been done on algorithms for performing sensitivity and/or parametric analysis for it. Yet, such analyses can be of interest in practice, for example, to inform decisions about pricing and promotional activities. In this chapter, we present a simple method for performing such analyses, based on discrete approximation and linear programming. The method can easily handle changes in costs, prices and resource availabilities. Under certain conditions, it can also handle changes in the demand distribution itself. Computational results show that the method yields accurate results quickly.

3.1 Introduction

The *Multi-item Newsvendor Problem* (MNP) is a classic stochastic optimisation problem that arises in the context of inventory control (Arrow et al., 1951). Although it is a relatively simple inventory model, it has received much attention from the Operational Research community (e.g., Choi, 2012; Khouja, 1999; Porteus, 2002; Silver et al., 1998; Turken et al., 2012; Zipkin, 2000). This is no doubt due to the wide range of applications of the model, including for example retailing, wholesaling, fast fashion, train and airline bookings, and the insurance sector.

In early work on the MNP, it is assumed that the demand in each time period comes from a fixed and known distribution, all other parameters are both fixed and known, and one needs only select the order quantities. Of course, in reality, the planning process is more complicated, involving not only the ordering policy but also decisions about pricing and promotional activities. These decisions in turn may influence the demand for products and thereby lead to changes in order quantities. There is therefore often a need to consider the relations between various decisions.

We believe that, in this regard, *Sensitivity Analysis* (SA) and *Parametric Analysis* (PA) are attractive tools. SA and PA are of course standard techniques in the optimisation literature (see, e.g., Dantzig, 1963; Gal, 1995; Gal and Greenberg, 1997). Surprisingly, however, not much has been written about SA for MNPs, and even less has appeared on PA for MNPs (see Section 3.2 for details). This chapter attempts to address this gap in the literature.

We make the following contributions in this chapter. First, we briefly review the existing literature on the MNP, as well as the existing SA methods for it. After that, we present a very natural method for performing SA and PA for MNPs, based on a combination of discrete approximation and linear programming. Despite its simplicity, the approach is remarkably general, being able to handle easily changes

in costs, prices, resource availabilities and (under certain conditions) changes in the demand distribution itself. Extensive computational results, on both artificial and real examples, show that the method yields accurate results quickly.

The chapter is organised as follows. Subsection 3.1.1 introduces some notation and terminology. Section 3.2 is a literature review. In Section 3.3, we model a broad family of MNPs as a stochastic linear program (SLP) and then construct a linear program (LP) that closely approximates the SLP. In Section 3.4, we show how to use the LP approximation to perform various kinds of SA and PA. Sections 3.5 and 3.6 give computational results for artificial and real data respectively. Finally, Section 3.7 contains some concluding remarks.

3.1.1 Notation and terminology

Consider the case with n items (i.e., product types) and a single sales period. The demand for item j over the sales period is a random variable \tilde{d}_j , with mean μ_j , standard deviation σ_j , and known cumulative density function F_j . We are also given vectors $\mathbf{c}, \mathbf{r}, \mathbf{v}, \mathbf{g} \in \mathbb{Q}^n$, where

- c_j is the cost of purchasing one unit of item j ;
- r_j is the revenue gained by selling one unit of item j ;
- v_j is the disposal cost of each unsold unit of item j ;
- g_j is the shortage cost of each unit of unsatisfied demand for item j .

We assume without loss of generality that $r_j > c_j \geq 0$, $c_j > -v_j$ and $g_j \geq 0$ for all j . We permit v_j to be positive or negative. (A negative value could indicate that excess items can be sold at a discounted price.) We remark that g_j may be used to represent the “loss of customer goodwill” incurred by stockouts.

In addition to the above, there are also m resources. (These could represent, e.g., purchasing budget or warehouse space.) We are given a matrix $\mathbf{A} = \{a_{ij}\} \in \mathbb{Q}_+^{m \times n}$ and a vector $\mathbf{b} \in \mathbb{Q}_+^m$, where

- a_{ij} is the amount of resource i used by one unit of item j ;
- b_i is the amount of resource i available.

We assume without loss of generality that $b_i > 0$ for all i .

The retailer must decide how many units of each item j to order before the start of the sales period. We let x_j denote the number of units of item j ordered. The resource constraints can then be written as $\mathbf{Ax} \leq \mathbf{b}$. We assume for simplicity that the x_j are continuous.

For a given item j , a given value of x_j , and a given realisation d_j of \tilde{d}_j , the profit over the period is:

$$\pi_j(x_j, d_j) = \begin{cases} r_j d_j - c_j x_j - v_j(x_j - d_j) & \text{if } x_j \geq d_j \\ (r_j - c_j)x_j - g_j(d_j - x_j) & \text{if } x_j < d_j. \end{cases}$$

The goal is to find values for the x_j that maximise the total expected profit.

We remark that many authors prefer to work with the *opportunity cost* rather than the profit. It takes the form:

$$\begin{cases} c_j^o(x_j - d_j) & \text{if } x_j \geq d_j \\ c_j^u(d_j - x_j) & \text{if } x_j < d_j, \end{cases}$$

where $c_j^o = c_j + v_j$ and $c_j^u = r_j - c_j + g_j$ are the *overage* and *underage* costs, respectively. We prefer however to work with the original data since it makes the SA and PA output easier to interpret.

3.2 Literature Review

In this section, we review the relevant literature. For brevity, we focus on the “standard” MNP, in which the goal is to maximise the total expected profit and no product substitution occurs in the event of a stock-out.

To our knowledge, Hadley and Whitin (1963) were the first to consider the NVP with a single side constraint (i.e., the case $m = 1$). To solve the problem, they proposed to relax the side constraint in a Lagrangian fashion, and then perform a search for the optimal value of the Lagrangian multiplier. They also pointed out that the variant in which the x_j must be integers can be converted into a separable concave integer knapsack problem, which can be solved by dynamic programming.

Hodges and Moore (1970) considered the general MNP. They proposed an iterative approximate solution method, based on a stochastic programming formulation. We remark that their method yields approximate dual prices for the side constraints as a by-product.

Nahmias and Schmidt (1984) pointed out that the Lagrangian method in Hadley and Whitin (1963) runs into difficulties if any of the assumed demand distributions take negative values. (This happens, for example, if the demands have been modelled using the normal distribution.) They then presented four fast and simple heuristics for this particular case.

Ben-Daya and Raouf (1993) gave a closed-form solution for the case in which $m = 2$ and all demands are uniformly distributed. They then used this as the basis of a heuristic for the case of other demand distributions. They also briefly discuss SA for the uniform case.

Lau and Lau (1995, 1996) considered the general MNP. They pointed out another disadvantage of Lagrangian methods, beyond the one mentioned in Nahmias and

Schmidt (1984): the methods can break down when some products have “strictly positive” demands (that is, the probability of the demand taking a non-positive value is zero). To deal with this, they propose an “active-set” method. We remark that their method, like the one of Hodges and Moore (1970), yields approximate dual prices.

Lau and Lau (1997) dealt with the issue of how best to model real-life demand distributions in MNPs. In Section 4 of their paper, they proved some results concerned with SA. Some of these results are intuitively obvious (e.g., the optimal value of x_j increases as c_j^u increases), but some are much less so (e.g., an increase in σ_j can cause the optimal value of x_j to increase).

There was then a series of papers focusing on the case $m = 1$. Moon and Silver (2000) showed that the dynamic programming algorithm in Hadley and Whitin (1963) can be modified to deal with the added complication of a fixed ordering charge for each item type. Erlebacher (2000) gave closed-form solutions for two special cases, and then used them as the basis for heuristics for the general case. (We remark that one of the two special cases was already solved in Ben-Daya and Raouf (1993)). Abdel-Malek et al. (2004) gave a closed-form solution for another special case, in which all demand distributions are negative exponential. They also gave an iterative procedure for the case of general distributions.

Staying with the case $m = 1$, Abdel-Malek and Montanari (2005a) gave a more detailed analysis of the conditions under which the Lagrangian method in Hadley and Whitin (1963) breaks down. This analysis includes some PA concerning resource availability. An extension to the case $m = 2$ was given in Abdel-Malek and Montanari (2005b).

Abdel-Malek and Areeratchakul (2007) proposed an effective heuristic for the general MNP, based on convex quadratic programming (CQP). They also proposed to use

known results on SA for CQP to perform approximate SA for the MNP itself. The accuracy of their approach depends heavily on the demand distributions.

Niederhoff (2007) formulated the general MNP as a nonlinear program with a separable convex objective and linear constraints. She then approximated each term in the objective with a piecewise-linear convex function, which enabled her to obtain an approximate solution via LP. She mentioned that one can use the SA output from the LP solver to perform approximate SA for the original MNP.

Zhang et al. (2009) showed that MNPs with one side constraint can be solved exactly, to arbitrary fixed precision, in linear time. We remark that this result is actually a special case of a result of Hochbaum (1995) on knapsack problems with a separable concave objective. Zhang (2012) extended the algorithm in Zhang et al. (2009), to show that MNPs can be solved to arbitrary precision in time that is linear in n (but exponential in m).

3.3 The Stochastic Program and Its Linear Approximation

We considered several different approaches to SA and PA for MNPs. In the end, we chose to model the MNP as a stochastic program, and then convert it into an LP via a discrete approximation of the marginal demand distributions. Our motivations for this choice are as follows:

- The approach is (relatively) simple to understand and implement.
- The approach is very general, needing no assumptions on the demand distribution (such as continuity or independence across items).
- A wide range of excellent LP software packages are now available (see Fourer

(2021)), all of which are capable of (a) generating SA output, and (b) re-optimising quickly after small changes have been made to the problem.

Moreover, as we will see in Section 3.5 and 3.6, our approach works very well in practice.

The first step is to model the MNP as a *2-stage stochastic linear programme with simple recourse* (Beale, 1955; Dantzig, 1955), or SLP for short. The first-stage variables are the x_j variables mentioned above. Then, for $j = 1, \dots, n$, we have second-stage variables y_j and z_j , representing the amount of over- and under-stocking of item j , respectively. The SLP is then

$$\begin{aligned} \max \quad & \sum_{j=1}^n (r_j - c_j) \mu_j - \mathbb{E}_{\tilde{\mathbf{d}}} [f(\mathbf{x}, \tilde{\mathbf{d}})] \\ \text{s.t.} \quad & \sum_{j=1}^n a_{ij} x_j \leq b_i \quad (i = 1, \dots, m) \end{aligned} \quad (3.1)$$

$$x_j \geq 0 \quad (j = 1, \dots, n), \quad (3.2)$$

where $f(\mathbf{x}, \tilde{\mathbf{d}})$ is a random quantity, found by solving the following second-stage problem:

$$\begin{aligned} \min \quad & \sum_{j=1}^n (c_j + v_j) y_j + \sum_{j=1}^n (r_j - c_j + g_j) z_j \\ \text{s.t.} \quad & x_j - y_j + z_j = \tilde{d}_j \quad (j = 1, \dots, n) \end{aligned} \quad (3.3)$$

$$y_j, z_j \geq 0 \quad (j = 1, \dots, n). \quad (3.4)$$

Note that the term $\sum_{j=1}^n (r_j - c_j) \mu_j$ in the first-stage profit function is a constant, and can therefore be ignored for optimisation. Later on, however, we will need to take it into account, since we may wish to do a parametric analysis on an individual r_j or c_j .

To reduce the size of the SLP, we use the fact that each y and z variable appears in exactly one of the equations (3.3). We can therefore eliminate the y variables (or z if preferred) from the problem. One can check that the resulting simplified SLP

takes the form:

$$\begin{aligned} \max \quad & \sum_{j=1}^n (r_j + v_j) \mu_j - \sum_{j=1}^n (c_j + v_j) x_j - \mathbb{E}_{\tilde{\mathbf{d}}} [f(\mathbf{x}, \tilde{\mathbf{d}})] \\ \text{s.t.} \quad & \quad \quad \quad (3.1), (3.2), \end{aligned}$$

where $f(\mathbf{x}, \tilde{\mathbf{d}})$ is redefined as

$$\begin{aligned} \min \quad & \sum_{j=1}^n (r_j + v_j + g_j) z_j \\ \text{s.t.} \quad & z_j + x_j \geq \tilde{d}_j \quad (j = 1, \dots, n) \\ & z_j \geq 0 \quad (j = 1, \dots, n). \end{aligned}$$

Note that the new term $\sum_{j=1}^n (r_j + v_j) \mu_j$ in the first-stage profit function is also a constant.

Following Beale (1955) and Dantzig (1955), we then use *scenarios* to construct an LP that approximates the simplified SLP. Let t be a positive integer parameter. For $i = 1, \dots, n$ and $s = 1, \dots, t$, let d_j^s be the realisation of \tilde{d}_i in scenario s , and let p_j^s be the probability of the scenario occurring. The LP then consists of maximising

$$\sum_{j=1}^n (r_j + v_j) \mu_j - \sum_{j=1}^n (c_j + v_j) x_j - \sum_{j=1}^n \sum_{s=1}^t p_j^s (r_j + v_j + g_j) z_j^s \quad (3.5)$$

subject to (3.1), (3.2), and

$$z_j^s + x_j \geq d_j^s \quad (j = 1, \dots, n; s = 1, \dots, t) \quad (3.6)$$

$$z_j^s \geq 0 \quad (j = 1, \dots, n; s = 1, \dots, t). \quad (3.7)$$

We remark that the first two terms in the objective function (i.e., the ones involving the μ_j and the x_j) are modelled exactly. Only the third term (i.e., the one involving the z variables) involves an approximation.

We also remark that, when generating the scenarios, one needs only a sample from the *marginal* demand distributions, one for each item, rather than the *joint* demand distribution (see El Agizy, 1967). We consider three scenario-generating methods (see Miller and Rice, 1983):

-
- naive random (a.k.a. Monte Carlo) sampling (RS),
 - the method of “equally likely intervals” (EI),
 - “Gaussian quadrature” (GQ).

In RS, we just sample completely at random and set all probabilities to $1/t$. In EI, the demand realisations are generated using the formula

$$d_j^s = F^{-1} \left(\frac{s}{t+1} \right),$$

for all j and s . All probabilities are then set to $1/t$ as before. In GQ, we match the first $(2t - 1)$ moments of \tilde{d}_j exactly, by finding weights p_j^s and demands d_j^s that satisfy:

$$\sum_{s=1}^t p_j^s (d_j^s)^N = \mathbb{E} \left[\tilde{d}_j^N \right], \text{ for } N = 1, 2, \dots, (2t - 1). \quad (3.8)$$

The weights can be determined by solving a system of simultaneous linear equations (3.8). We omit details, for brevity.

We end this section with some more notations. We let $(\mathbf{x}^*, \mathbf{z}^*)$ denote the optimal primal solution of the LP, and we let $\boldsymbol{\kappa}^*$ and $\boldsymbol{\lambda}^*$ denote the optimal dual vectors for the constraints (3.1) and (3.6), respectively. Finally, we let P and \tilde{P} denote the expected profit of the optimal SLP and LP solutions, respectively. When performing SA and PA concerning a given parameter, we will view P and \tilde{P} as being *functions* of that parameter.

Due to the use of a finite number of scenarios, \tilde{P} is only an approximation of P . As t approaches infinity, however, \tilde{P} will approach P (see, e.g., Birge and Louveaux, 2011; Kall and Wallace, 1994 for results on the accuracy of LP approximations of SLPs).

3.4 Sensitivity and Parametric Analysis

In this section, we show how to use the LP in the previous section to perform SA and PA. The section is structured as follows. Subsection 3.4.1 presents some elementary SA and PA results. Subsection 3.4.2 considers changes in the demand distributions, and Subsection 3.4.3 deals with cross-price elasticities of demand.

3.4.1 Some elementary results

The following three lemmas follow from standard textbook results on LP (e.g., Dantzig, 1963; Vanderbei, 2020):

Lemma 3.4.1. *\tilde{P} is continuous and piecewise-linear in the c_j , r_j , v_j , g_j and b_i parameters. It is convex in the c_j , r_j , v_j and g_j , but concave in the b_i . It is non-increasing in the c_j , v_j and g_j , but non-decreasing in the r_j and b_i .*

Lemma 3.4.2. *At $(\mathbf{x}^*, \mathbf{z}^*)$, we have:*

$$\begin{aligned}\partial\tilde{P}/\partial c_j &= -x_j^* \quad (j = 1, \dots, n) \\ \partial\tilde{P}/\partial b_i &= \kappa_i^* \quad (i = 1, \dots, m).\end{aligned}$$

Lemma 3.4.3. *At $(\mathbf{x}^*, \mathbf{z}^*)$, we also have the following for $j = 1, \dots, n$:*

$$\begin{aligned}\partial\tilde{P}/\partial r_j &= \mu_j - \sum_{s=1}^t p_j^s(z_j^s)^* \\ \partial\tilde{P}/\partial v_j &= \mu_j - x_j^* - \sum_{s=1}^t p_j^s(z_j^s)^* \\ \partial\tilde{P}/\partial g_j &= -\sum_{s=1}^t p_j^s(z_j^s)^*.\end{aligned}$$

We now make two remarks:

- \tilde{P} is also continuous and non-increasing in the a_{ij} parameters, but it is in general neither convex nor concave in those parameters.

-
- In reality, changing one of the selling prices r_j is likely to affect the demand for item j , and it may even affect the demand for other items. Thus, the results that involve the r_j may be inaccurate. We deal with this complication in Subsection 3.4.3.

3.4.2 Changes in the demand distributions

In practice, it may be possible for a marketing team to take actions that affect the demand distributions themselves. Examples of such actions can be found in Darwish et al. (2019), Gerchak and Parlar (1987), Güler (2019), Khouja and Robbins (2003) and Wang (2011). We, therefore, consider SA and PA concerning changes in specific parameters of the marginal demand distributions.

Suppose that the demand of item j can be modelled as:

$$\tilde{d}_j = \mu_j + \tilde{\epsilon}_j \sigma_j, \quad (3.9)$$

where μ_j and σ_j are the mean and standard deviation, and $\tilde{\epsilon}_j$ is itself a random variable with zero mean, unit variance, and known distribution. (This is the case, for example, if \tilde{d}_j follows a normal distribution.) If we increase μ_j by δ , then the objective function (3.5) increases by $(r_j + v_j)\delta$, and the right-hand sides of the corresponding constraints (3.6) all increase by δ . This implies:

Lemma 3.4.4. *If (3.9) holds for item j , then, at $(\mathbf{x}^*, \mathbf{z}^*)$, we have:*

$$\partial \tilde{P} / \partial \mu_j = r_j + v_j - \sum_{s=1}^t (\lambda_j^s)^*.$$

Applying a similar argument to changes in σ_j , we obtain:

Lemma 3.4.5. *If (3.9) holds for item j , then, at $(\mathbf{x}^*, \mathbf{z}^*)$, we have:*

$$\partial \tilde{P} / \partial \sigma_j = - \sum_{s=1}^t \epsilon_j^s (\lambda_j^s)^*.$$

We can also present an analogue of Lemma 3.4.1:

Lemma 3.4.6. *If (3.9) holds for j , then \tilde{P} is continuous, concave and piecewise-linear in both μ_j and σ_j .*

There may of course exist situations when (3.9) doesn't hold. (This happens, for example, if \tilde{d}_j follows a Poisson or negative binomial distribution.) In that case, there does not appear to be a simple closed formula for the partial derivatives, and it is not obvious whether \tilde{P} will be concave in any particular parameter.

3.4.3 Cross-price elasticities

Finally, we consider the case in which there are *cross-price elasticities* in demand (see, e.g., Frank and Cartwright, 2008). That is, a change in the price of one product can affect the mean demand for other products.

We first consider the case in which an *additive* price-demand model has been used (Mills, 1959). In particular, suppose that the demand for product j can be written in the form

$$\tilde{d}_j = \alpha_j + \sum_{k=1}^n \beta_{jk} r_k + \tilde{\epsilon}_j,$$

where the r_k are the product prices, α_j and the β_{jk} are known scalars, and $\tilde{\epsilon}_j$ is a noise term from an assumed distribution.

Now suppose we increase the price of product k by some small $\delta > 0$. This will cause the demand for product j to increase by $\delta\beta_{jk}$. The effect on the LP is as follows:

- The constant term in the objective function (3.5) increases by

$$\delta \left(\mu_k + \sum_{j=1}^n (r_j + v_j) \beta_{jk} \right).$$

- For all s , the coefficient of z_k^s in the same function decreases by δp_k^s .

- For all j and s , the right-hand side of (3.6) increases by $\delta \beta_{jk}$.

From this we obtain

$$\partial \tilde{P} / \partial r_k = \mu_k + \sum_{j=1}^n (r_j + v_j) \beta_{jk} - \sum_{s=1}^t p_k^s (z_k^s)^* - \sum_{j=1}^n \beta_{jk} \sum_{s=1}^t (\lambda_j^s)^*.$$

Now we consider the case of a *multiplicative* price-demand model (Karlin and Carr, 1962). Suppose that the demand for product j takes the form

$$\tilde{d}_j = \alpha_j \tilde{\epsilon}_j \prod_{k=1}^n r_k^{\beta_{jk}}, \quad (3.10)$$

where we now assume that the noise term $\tilde{\epsilon}_j$ comes from a positive distribution (e.g., lognormal).

As before, suppose we increase the price of product k by some small δ . This will cause the demand for product j to increase by approximately $\delta \beta_{jk} \mu_j$, where μ_j is the initial mean demand for item j . The effect on the LP is that:

- The constant term in (3.5) increases by

$$\delta \left(\mu_k + \sum_{j=1}^n (r_j + v_j) \beta_{jk} \mu_j \right).$$

- For all s , the objective coefficient of z_k^s decreases by δp_k^s (as before).
- For all j and s , the right-hand side of (3.6) increases by approximately $\delta \beta_{jk} d_{js}$.

From this we obtain:

$$\partial \tilde{P} / \partial r_k = \mu_k + \sum_{j=1}^n (r_j + v_j) \beta_{jk} \mu_j - \sum_{s=1}^t p_k^s (z_k^s)^* - \sum_{j=1}^n \beta_{jk} \sum_{s=1}^t d_{js} (\lambda_j^s)^*.$$

We remark that, when cross-price elasticities are present, neither P nor \tilde{P} are guaranteed to be non-increasing, non-decreasing, convex, concave or piecewise-linear in the r_j . Indeed, the only thing that we typically know in advance, before actually performing the PA, is that P and \tilde{P} are continuous in the r_j .

3.5 Results with Artificial Data

In this section, we use some artificial MNP instances to explore the way our chosen method behaves, in terms of accuracy and running time, as the number of scenarios and/or problem size grows. All the experiments were performed on an Apple M1 Pro (2021). All the LPs were solved using the simplex-based `lpSolve` package in R.

3.5.1 Instance generation

To begin, we explain how we generate our instances. For simplicity, we assume that the marginal demand distributions are normal. The mean demands μ_j are random integers sampled uniformly between 150 and 250, and the standard deviations σ_j are random integers between 15 and 35. The other parameters are also selected uniformly at random, with $c_j \in [3, 5]$, $r_j \in [8, 10]$, $v_j \in [2, 4]$, $g_j \in [1, 3]$, $a_{ij} \in [4, 8]$ and $b_i \in [4n, 8n]$.


Our instances have n ranging from 2 to 200 and m ranging from 1 to 40. A small instance, with $n = 2$ and $m = 3$, is given in Appendix C. Throughout this section, all reported values are averages over 20 random instances of the given size.

3.5.2 Increasing the number of scenarios

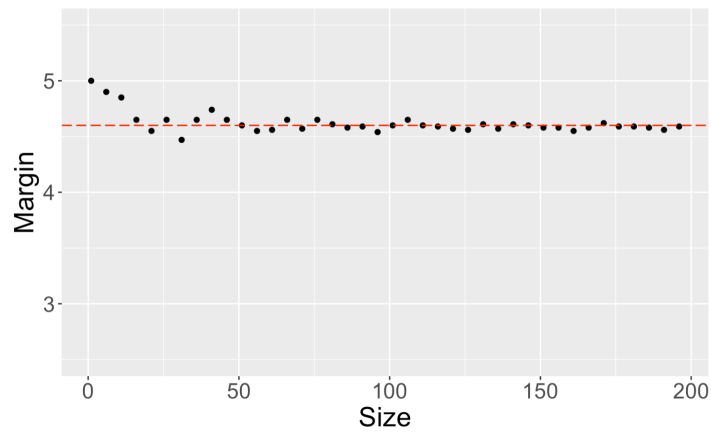
We first consider how the accuracy of the method depends on the number of scenarios. Figure 3.1 shows how the estimate of $\partial\tilde{P}/\partial\mu_a$ varies as the number of scenarios increases, for the small example in Appendix C, for the three different scenario generation schemes mentioned in Section 3.3. The red lines indicate the true value.

As one would expect, the estimate tends to get more accurate with more scenarios, but there is some random variation. Also, the ‘systematic’ scenario generation

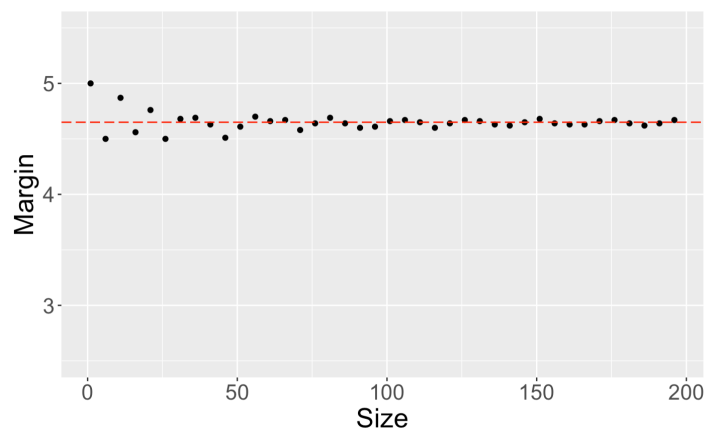
Figure 3.1: Estimated partial derivative vs. number of scenarios



(a) Using RS



(b) Using EI



(c) Using GC

Table 3.1: Number of scenarios required to achieve accuracy within $\pm 5\%$

	Method	Solution output			Marginal gain			
		P	x_a^*	x_b^*	μ_a	μ_b	σ_a	σ_b
Normal distribution	RS	78	69	72	61	55	79	80
	EI	23	22	23	21	18	29	30
	GQ	18	16	16	13	13	22	21
Laplace distribution	RS	75	62	62	60	59	77	78
	EI	16	13	15	13	11	22	21
	GQ	11	10	10	11	10	13	13
Uniform distribution	RS	82	77	79	60	65	76	84
	EI	27	23	25	20	17	31	33
	GQ	23	26	21	14	15	25	27

methods, EI and GC, lead to faster convergence than RS. This behaviour is typical in stochastic programming (Birge and Louveaux, 2011).

Table 3.1 shows, for the same instance, the number of scenarios required to compute various quantities to an accuracy of $\pm 5\%$. The quantities considered are: the expected profit of the optimal solution (P), the optimal order quantities (x_a^* and x_b^*), and the marginal gains with respect to changes in the first and second moments of the marginal demand distributions (μ_a , μ_b , σ_a and σ_b). As before, we include results for RS, IE and GC. Moreover, for completeness, we include results for three different demand distributions: normal, Laplace and uniform.


Table 3.1 shows that EI and GQ need far fewer scenarios than RS, with GQ being the overall winner. Interestingly, estimating P accurately tends to need more scenarios than estimating x_a^* and x_b^* , and estimating the effect of changes in standard deviations needs even more scenarios. On the other hand, estimating the effect of changes in means needs fewer scenarios. We observed the same behaviour with the other instances that we tested. This suggests that, for decision-makers who are only

interested in the effect of increasing mean demands, a small scenario size is sufficient.

We remark that GQ is harder to implement than EI, since the computation of the weights p_j^s and d_j^s in (3.8) involves the solution of a system of simultaneous linear equations. Moreover, in practice, it may be hard to compute reliable estimates of the higher moments of the marginal demand distributions. For these reasons, EI may be preferable to GQ in practice.

Now we turn our attention to running times. Figure 3.2 shows how the time taken to solve the LP changes as the number of scenarios increases (for the same instance, using EI). A regression analysis shows that the empirical growth function is closer to quadratic than exponential. (This is what one would expect from known results on the expected running time of the simplex method; see Shamir, 1987) A best-fit quadratic function is shown in the figure as a red dashed line. The running times when using RS or GQ were similar.

Figure 3.2: Running time vs. Number of scenarios (using EI)



We remark that, in all cases, the time taken to perform the SA was negligible compared to the time taken to solve the LP itself. The time taken to compute the weights in GQ was also negligible. Performing PA, however, took a little longer (up to a few seconds in some cases, depending on the chosen parameter).

Table 3.2: Number of scenarios required to achieve accuracy within $\pm 5\%$ as n increases (*with 25 scenarios under `lp()` function in R)

n	Running time*	Solution output		Marginal gain	
		P	x^*	μ	σ
10	0.029s	30	31	27	31
20	0.156s	24	25	24	28
50	0.889s	19	19	17	22
100	2.472s	12	13	11	16
200	5.293s	9	8	6	11


3.5.3 Increasing the problem size

Next, we examine how the approach behaves as the number of items n increases. For brevity, we give results only for the case in which (a) the marginal demand distributions are normal, and (b) EI is used to generate scenarios. (The results obtained for other distributions and/or scenario-generation methods were similar.)

Table 3.2 shows the following information for various values of n : the time taken to solve the LP (in seconds) when 25 scenarios are used per item, and the number of scenarios required to estimate various quantities within an accuracy of $\pm 5\%$. The quantities considered are: the optimal expected profit (P), the optimal first-stage solution vector (x^*), and the expected marginal gains with respect to changes in the first and second moments of the marginal demand distributions (μ and σ).

As before, estimating the effect of changes in means requires fewer scenarios than estimating the other quantities. More interestingly, the number of scenarios needed to obtain an accurate solution tends to decrease as n increases. A possible explanation of this phenomenon is that the amount of ‘information’ in the LP increases with the number of items, even when the number of scenarios is fixed.

Figure 3.3: Estimate of optimal expected profit vs. number of scenarios for large instance with 200 items (using EI)




Concerning the running time, the figures in the table suggest that the growth is quadratic in n . Fortunately, the running time is very reasonable even for the instance with 200 items. We also remark that, as in the previous subsection, the time taken to perform PA was typically two or three times longer than the time taken to solve the LP.

To gain more insight into the accuracy of the method, we show in Figure 3.3 how the estimated expected profit (\tilde{P}) approaches the true value (P) as the number of scenarios increases, for the instance with 200 items. Interestingly, the estimate approaches the true value from above. This is because EI slightly underestimates the variance of demand, which leads to an over-optimistic profit estimate. Fortunately, the accuracy improves rapidly as we increase the number of scenarios.

Finally, we consider the effect of the number of resources m on the running time. Figure 3.4 shows the average running time in seconds, for $n = 10$ and 25 scenarios, as m ranges from 1 to 40. Each point represents the average over 20 random instances. A regression analysis shows that the empirical growth function is closer to linear than polynomial or exponential.

Figure 3.4: Running time vs. number of resources for a 10-item example (using EI)



3.6 Results with Real Data

We now apply the approach to a real-life case, in which cross-price elasticities of demand play a significant role. Our demand data comes from a medium-sized grocery store which sells a wide range of products, many of which are perishable. The retailer wonders whether better coordination of price-setting across products could increase the total expected profit. Here, we select a small group of products to analyse as an example.

Together with the retailer, we identified the nine most popular products, each from a perishable product category. For reasons of confidentiality, we use category names to refer to them. The retailer was able to estimate the mean and standard deviation of daily demand for each product, along with the revenue and cost parameters. These values are shown in Table 3.3.

The retailer had deliveries every morning, and most of the leftover products at the end of each day were donated to a charity. Thus, the problem can be well approximated by a single-period inventory model and is very close to the MNP described in the previous section.

Table 3.3: Data for a subset of the products

product	demands		price and costs			
	μ	σ	r	c	v	g
bread	87.1	49.8	0.93	0.63	0.21	0.05
egg	57.6	22.8	4.29	3.24	1.03	0.21
fish	44.2	14.1	2.79	1.75	1.20	0.49
fruit	124.1	42.9	4.69	3.35	0.81	0.42
juice	45.3	13.7	3.99	2.56	0.33	0.45
vegetables	1197.5	355.09	2.86	1.96	0.78	0.56
meat	126.8	10.2	20.99	16.67	3.89	2.10
milk	60.2	11.2	1.94	1.28	0.60	0.35
dairy	15.8	9.7	2.28	1.63	0.55	0.13

There also existed five resource constraints. For instance, there was an upper limit on the total amount of liquid products that the retailer could store, and there were lower limits on the sizes of orders from particular brands. The resource constraints are summarised in Appendix D. (They have been slightly modified to fit with the example with nine products.)

Based on our discussions with the retailer, we believe that the *multiplicative* price-demand model (3.10) is suitable for these products, since we assume the elasticities are constant over a range. We also assume for simplicity that the error terms $\tilde{\epsilon}_j$ are i.i.d. lognormal.

The retailer did not have information about elasticities, and it is not easy to acquire accurate elasticity data. Here we use data gathered from the literature to approximate the situation (Henneberry et al., 1999; Liu and Chern, 2003; Zhang and Wang, 2003). The estimated elasticities for our selected products are given in Appendix E.

We remark that bread has a positive elasticity, which suggests that in this situation bread can be considered a ‘Giffen good’ (see, e.g., Masuda and Newman, 1981).

To convert the SLP into an LP, we used EI with 25 scenarios per product. The resulting LP had 9 x variables, 225 z variables, 5 first-stage constraints and 225 second-stage constraints. We were able to solve it in a few seconds using the `lpSolve` package.

The optimal first-stage solution (x^*) and the total expected profit (\tilde{P}) are presented in Table 3.4. We remark that, for all items, the optimal stock level is below the mean demand (compare Tables 3.3 and 3.4). This is because, in this application, the over-stocking costs are higher than the under-stocking costs.

Table 3.4: Optimal first-stage solution and expected profit

x_1^*	x_2^*	x_3^*	x_4^*	x_5^*
62.07	40.88	38.65	102.54	41.26
x_6^*	x_7^*	x_8^*	x_9^*	\tilde{P}
1056.95	119.35	55.76	9.79	1265.19

We are now ready to perform the SA. We explore the sensitivity of \tilde{P} concerning percentage changes in (a) the prices of individual products and (b) the standard deviations of the demands for individual products. The resulting margins, and their ranges of validity (rounded to the nearest percentage point), are shown in Table 3.5.

When interpreting the results in Table 3.5, it should be borne in mind that \tilde{P} is not piecewise-linear in the individual item prices (see the last paragraph in Subsection 3.4.3). This means that the ranges of validity need to be interpreted carefully. In the table, we show the ranges in which the margins remain valid to within an accuracy of $\pm 5\%$. For instance, when the change of bread price is within $[-3\%, 4\%]$, the margin of profit is $0.12 \times (1 \pm 5\%)$.

Table 3.5: Estimated margins and their ranges of validity


	Margins of profit		Ranges of validity		
	% of p_j increase	% of σ_j decrease	% of p_j decrease	% of p_j increase	% of σ_j decrease
bread	0.12	0.19	3	4	100
egg	1.48	0.34	9	9	55
fish	1.09	0.21	2	2	100
fruit	4.29	0.80	3	4	100
juice	0.95	0.23	4	2	100
vegetable	22.23	5.02	3	4	25
meat	22.34	0.77	4	4	100
milk	4.21	0.11	2	2	100
dairy	-1.68	0.08	2	3	100

To explore the behaviour outside the range, one must use PA. In this context, it should be noted that \tilde{P} is not necessarily convex or concave in the individual item prices. This is shown in Figure 3.5, where \tilde{P} varies with respect to changes in the prices of juice and bread. (These curves were computed by solving a large number of similar LPs.) The curves can be either concave or convex, depending on the price elasticity of the product in question. Note also that we assume in this case study that price elasticities are fixed, while in reality, this might not always be true. For example, if the price of bread increased indefinitely, the demand for it would eventually go down, because consumers would switch to other, cheaper products.

Using the output of our procedure, we were able to provide some insights and suggestions for the retailer:

- Increasing the price of vegetables and/or meat is likely to yield a significant increase in profit. (We would however recommend making only small price increases, at least initially, since sudden large price increases could lead to a loss of customer goodwill.)

Figure 3.5: Optimal profit vs. percentage change on price



(a) Change on the price of juice

(b) Change on the price of bread

- Increasing the price of some other products may also increase the expected profit, but only slightly.
- For a specific product (dairy), it may be beneficial to *decrease* the price.
- Decreasing the demand variance could also improve the expected profit. If the marketing budget is limited, the retailer should first consider marketing techniques that could decrease the variance of demand for vegetables. For instance, switching to online retailing may make demand more consistent, thus decreasing the variance (see Darwish et al., 2019; Gerchak and Parlar, 1987; Güler, 2019; Khouja and Robbins, 2003; Wang, 2011 for other examples).

3.7 Concluding Remarks

Although multi-item newsvendor problems have been studied intensively, there are very few works that consider the ordering policy and marketing activities as a whole. This chapter aimed to present a simple sensitivity analysis method that allows one to estimate the value of these activities accurately and quickly, using any decent linear programming solver. The proposed method is robust and doesn't require any additional assumptions on the distribution. Our work provides a useful tool to help

coordinate marketing and inventory decisions in a retail environment. Our extensive computational results, on both artificial and real examples, show that the method yields accurate results quickly.

An interesting (and challenging) topic for future research is the development of methods for performing sensitivity analysis when, in the event of a stock-out, customers may be willing to buy substitute products (Shin et al., 2015). The first steps in this direction have been done by Zhang et al. (2021).

Chapter 4

Naïve Newsvendor Adjustments: Are They Always Detrimental?

Another strand of the literature on newsvendor problems is concerned with the fact that practitioners often make judgemental adjustments to the theoretically “optimal” order quantities. Although the judgemental adjustment is sometimes beneficial, two specific kinds of adjustment are normally considered to be particularly naïve: *demand chasing* and *pull-to-centre*. We discuss how these adjustments work in practice and what they imply in a variety of settings. We argue that even such naïve adjustments can be useful under certain conditions. This is confirmed by experiments on simulated data. Finally, we propose a heuristic algorithm for “tuning” the adjustment parameters in practice.

4.1 Introduction

Single-period stochastic inventory control problems, known as *Newsvendor problems* (NVPs), have received much attention from the Operational Research community (see, e.g., the books Choi, 2012; Hadley and Whitin, 1963; Silver et al., 1998; Zipkin,

2000). In this chapter, we focus on the simplest NVP, in which there is only one product. The demand for the product over the selling period is a random variable \tilde{d} , with known distribution. The product is purchased before the period at a fixed unit price v , and sold during the period at a unit price p . If any excess stock remains at the end of the period, disposal cost c_h is incurred per unit. If there is any unsatisfied demand, a shortage cost c_s is incurred per unit.

For a given x and a given realisation d of \tilde{d} , the realised profit π over the period is:

$$\pi(x, d) = \begin{cases} pd - vx - c_h(x - d), & \text{if } x \geq d \\ px - vx - c_s(d - x), & \text{if } x < d. \end{cases} \quad (4.1)$$

The optimal order quantity that maximises the expected profit is then (Arrow et al., 1951; Choi, 2012):

$$x^* = F^{-1}(\tau), \quad (4.2)$$

where F is the cumulative distribution function for \tilde{d} , $c_o = v + c_h$ is the *overage* cost, $c_u = p - v + c_s$ is the *underage* cost, and the quantile $\tau = c_u/(c_o + c_u)$. We will call a product with $\tau > 0.5$ a “high-margin” product, and a product with $\tau < 0.5$ a “low-margin” product.

In the textbook formulation, it is assumed that one has a correct model of the demand distribution F , with correct parameters. In real life, however, correctness is rarely assured and a model typically misses important information. Moreover, even if the model is correct, the parameters may evolve over time (for example, due to market shocks or product innovations by competitors). For these reasons, decision-makers often make so-called “judgemental adjustments” to the theoretically “optimal” order quantities. The literature on this topic is extensive (e.g., Lau and Bearden, 2013; Lau et al., 2014; Schweitzer and Cachon, 2000).

Although judgemental adjustments can exist in many forms (see Goodwin and Wright, 2014; Kahn, 2014; Ord et al., 2017 for example), two specific kinds of

adjustment have received the most attention in the NVP literature (Benzion et al., 2008; Bolton and Katok, 2008; Bostian et al., 2008; Lau and Bearden, 2013; Lau et al., 2014). The first, called *pull-to-centre*, means adjusting the order quantity towards the mean demand. The second, called *demand chasing*, means adjusting it towards the demand in the immediately preceding period. (Some literature suggests they are not independent.) Although these two kinds of adjustment are regarded as especially naïve, considerable evidence for their existence has been found by scholars, and several theories have been developed to explain them (Benzion et al., 2008; Bolton and Katok, 2008; Bostian et al., 2008; Cui et al., 2013; Feng et al., 2011; Moritz et al., 2013; Wu and Niederhoff, 2014). Yet, to our knowledge, there has not been any numerical study of the effect that these two naïve adjustment mechanisms are likely to have on the long-term expected profit in the NVP.

In this chapter, we look at naïve adjustment from a different point of view. We begin by arguing that such adjustments may be useful under certain conditions, due to the fact that any given statistical model of the demand is unlikely to be completely accurate. In particular, we suggest that a modest amount of naïve adjustment may be beneficial in two specific situations that are highly important in practice; namely, (i) when a relatively small amount of demand data is available, and (ii) when the true demand model is unknown. This idea is tested via extensive computational experiments. We then consider the possibility of applying naïve adjustments in an automated fashion. For this purpose, we propose a simple heuristic for tuning the adjustment parameters. Finally, we test the heuristic on a real-life example, with encouraging results.

The chapter is organised as follows. In Section 4.2, we summarise the existing literature. In Section 4.3, we argue in favour of naïve adjustment, and propose a mathematical model which incorporates both types of adjustment. In Section 4.4, we conduct experiments on simulated data under different conditions and discuss

the results. In Section 4.5, we present the tuning heuristic and apply it to a real-life example. Finally, Section 4.6 contains some concluding remarks.

4.2 Literature Review

In this section, we review and discuss the existing literature on naïve adjustment. Subsections 4.2.1 and 4.2.2 are concerned with demand-chasing and pull-to-centre, respectively.

4.2.1 Demand chasing effect

The *demand chasing* effect (DC), in the NVP context, describes a phenomenon whereby decision-makers tend to adjust their order quantity towards the realised demand in the previous operating period.

To our knowledge, the first paper to give empirical evidence for DC was Schweitzer and Cachon (2000). They observed decision makers over fifteen consecutive ordering periods, for a selection of products, each with known distribution. They showed that the participants systematically deviated from the optimal order quantity. They also gave a tentative explanation for the phenomenon, based on the avoidance of *regret*. The idea is that, if decision-makers fail to choose the *ex-post* optimal order quantity in a given period, they regret their decision, which leads them to adjust the quantity in the next period.

After the publication of Schweitzer and Cachon (2000), evidence for DC appeared in many papers. Bolton and Katok (2008) repeated the experiment, but with 100 decision periods. They found that the participants tended to improve over time, but only very slowly. Benzion et al. (2008) studied the convergence of the participant's behaviour, and argued that the order quantities from decision-makers converge to a

level different from the one which optimises expected profit. Other relevant works include Bostian et al. (2008), Cui et al. (2013), Feng et al. (2011), Moritz et al. (2013) and Wu and Niederhoff (2014).

Lau and Bearden (2013) argued that some of the statistical techniques used in the behavioural experiments were flawed. See the recent paper by Kirshner and Moritz (2020) for a discussion of this issue.

Several models of DC were proposed in the above-mentioned papers. For brevity, we present only the simplest model, which appeared in Bostian et al. (2008). It takes the form:

$$x_t = x_{t-1} + \beta(d_{t-1} - x_{t-1}), \quad (4.3)$$

where x_t is the actual order quantity in time period t , d_{t-1} is the realised demand in the previous period and $\beta > 0$ is the *DC parameter*. A higher β indicates a stronger demand chasing effect. (In rare cases, one might observe $\beta < 0$, as a “pull forward in demand”.)

We remark that the model (4.3) is equivalent to “simple exponential smoothing” or SES (Brown, 1956), in the so-called “error-correction” form. We observe that it ignores the NVP solution and asymptotically converges to the mean demand. Moreover, it assumes a fixed β for all periods, implying that practitioners adjust the orders every period by the same proportion. Fortunately, the latter assumption does not cause serious problems in practice. Indeed, even if practitioners adjust the order with different quantities over time, their behaviour can be modelled on average using (4.3). Furthermore, there is empirical evidence that, for human decision makers, the value of their β will eventually converge to a single value over time (Schweitzer and Cachon, 2000; Zhang and Siemsen, 2019).

4.2.2 Pull-to-centre effect

The *pull-to-centre* effect (PtC) describes a phenomenon when a decision maker adjusts the order quantity towards the mean demand (Zhang and Siems, 2019). This phenomenon has also been referred as mean anchor (Schweitzer and Cachon, 2000) and central tendency bias (Bostian et al., 2008) in literature. In the rest of this chapter, we use the term PtC for consistency.

To our knowledge, the first paper to give empirical evidence for PtC was again Schweitzer and Cachon (2000). Their interpretation of PtC is that decision-makers order less than the optimal amount for high-margin products, but more for low-margin products. They also discussed several possible causes for the phenomenon, including risk and loss aversion, underestimation of opportunity cost, and waste aversion. They also discussed a possible explanation in terms of “prospect theory” (Kahneman and Tversky, 1979).

Alternative explanations of PtC include adaptive learning (Benzion et al., 2008), decision noise and optimisation error (Su, 2008), overconfidence bias (Ren and Croson, 2013), and psychological costs associated with leftovers and stockouts (Ho et al., 2010). We remark that Lau et al. (2014) argued that some of the statistical techniques used to detect PtC were flawed, just as Lau and Bearden (2013) argued for DC.

There is some evidence that individual differences can affect the behaviour of the decision-maker in NVPs. De Vericourt et al. (2013) showed that males tend to take more risks than females, which leads them to order more, on average. Cui et al. (2013) and Feng et al. (2011) showed that differences in nationality correlate with different biases while making newsvendor decisions.

Benzion et al. (2008) found that decision-makers tend to be more biased towards the mean demand in earlier periods than in later periods. This suggests that training

could be of some benefit to decision-making. Additional discussions of training effects can be found in Bolton and Katok (2008), Bostian et al. (2008), Ren and Croson (2013) and Zhang and Siemsen (2019).

Following the works of Benzion et al. (2008) and Bostian et al. (2008), the PtC effect can be expressed mathematically as:

$$x_t = (1 - \gamma)x_t^* + \gamma\hat{\mu}_t, \quad (4.4)$$

where $\hat{\mu}_t$ is the estimated mean demand for the period t (which can be obtained with a forecasting technique), and $0 < \gamma < 1$ is the *PtC parameter*. In this case, the order quantity can be viewed as a weighted average of the “textbook” order quantity x_t^* , and the estimated mean $\hat{\mu}_t$. A higher γ indicates a stronger PtC effect.

Note that the model (4.4), like the DC model (4.3), assumes that the adjustment parameter is constant over time and that adjustments happen on each observation. We argue that these assumptions are reasonable because they express a behaviour on average, similar to how the DC behaviour is modelled via (4.3).

4.3 An Alternative Perspective and Model

In this section, we argue that there are some positive aspects to naïve adjustment. We also present a new adjustment model, which allows one to perform demand-chasing and pull-to-centre in combination if desired.

4.3.1 In favour of naïve adjustment

As one can see from Section 4.2, the previous literature on judgemental adjustment has assumed, either implicitly or explicitly, that DC and PtC are harmful. In this chapter, we take a different point of view: we argue that DC and PtC may

sometimes be *beneficial* in practice. To see why, note that the “textbook” NVP formula (4.2) applies only when one has an accurate statistical model of the demand distribution. In reality, of course, such a model is rarely available. As a result, the textbook formula may give the wrong answer in practice, either underestimating or overestimating the optimal order level. In some circumstances, therefore, DC and/or PtC might help rather than hinder.

To be more specific, we suggest that a modest amount of “naïve” adjustment may be beneficial in two practically important situations:

1. When the demand model is correct, but there is insufficient data to estimate its parameters accurately.
2. When the demand model is mis-specified.

We will test these hypotheses using simulation experiments in the next section, modelling the two situations.

There is another key difference between our work and the existing literature. As mentioned above, the latter relies almost exclusively on data collected from behavioural experiments with human subjects. Here, by contrast, we will use simulated data, since it allows us to conduct extensive experiments very easily.

4.3.2 An integrated adjustment model

To proceed, we make some additional remarks about the DC model (4.3). In our view, it is unlikely to be a good model of human behaviour. Indeed, we have already observed that it effectively estimates the mean demand, and does not take cost information into account.

In an attempt to remedy the above weakness, we now propose a “two-stage” model

of naïve adjustment, in which DC takes place after PtC:

$$\begin{aligned}x'_t &= (1 - \gamma)x_t^* + \gamma\hat{\mu}_t \\x_t &= x'_t + \beta(d_{t-1} - x_{t-1}).\end{aligned}\tag{4.5}$$

The idea here is that we first take the “textbook” order quantity x_t^* , and apply PtC with parameter γ . This yields an adjusted order quantity, here denoted by x'_t . After that, we adjust x'_t itself, by applying DC with parameter β .

Unlike the classical DC model (4.3), the two-stage model (4.5) yields non-trivial estimates of the optimal order quantity, rather than merely estimating the mean demand. We remark that we are not claiming that human practitioners actually use such a model consciously. But in practice, decision-makers cannot distinguish the two effects and adjust based on their domain knowledge. We argue that the resulting adjustments will contain both PtC and DC parts. However, our goal is not to explore the behaviour of decision-makers in practice, but to investigate whether the model (4.5) might be useful in the two situations with the model specification mentioned in the previous subsection.

Inserting the first equation in the second one in (4.5), we obtain a unified formulation for DC and PtC, which summarises the order adjustment in one formula:

$$x_t = (1 - \gamma)x_t^* + \gamma\hat{\mu}_t + \beta(d_{t-1} - x_{t-1}).\tag{4.6}$$

This makes it clearer that the adjusted order quantity is a linear combination of three terms: system order quantity, mean and actual demand. We will focus our investigation on the parameters β and γ taking values between 0 and 0.5 (This is a common assumption in the literature, e.g. Benzion et al., 2008; Bolton and Katok, 2008; Bostian et al., 2008), although the theoretical parameters range might be wider.

It is important to note that, for a given t , the estimate $\hat{\mu}_t$ is itself based on d_1, \dots, d_{t-1} , and so are the quantities x_t^* and x_{t-1} . Thus, all three quantities are

subject to estimation errors.

4.4 Experiments on Simulated Data

In this section, we perform extensive computational experiments, to test the two hypotheses mentioned in the previous section. In Subsection 4.4.1, we describe our methodology. The hypotheses themselves are tested in Subsections 4.4.2 and 4.4.3, respectively.

We assume initially that $\tau = 0.7$, a value commonly used in the NVP literature to approximate real-life problems (e.g., fashion retail, nurse staffing) (Alfares and Elmorra, 2005; Lariviere and Porteus, 2001). We show later in this section that the results of our experiment hold with other NVP parameters as well.

4.4.1 Methodology

The first step is to construct 500 time series, each consisting of 200 consecutive demand realisations, which is sufficiently long as shown in behavioural experiments (Benzion et al., 2008; Bostian et al., 2008; Schweitzer and Cachon, 2000). To do this, we use the `arma.sim()` function from the `stats` package in R. We assume that the “true” DGP for the demands is an ARIMA(1,0,1) process, with an initial mean of 10,000. We also assume that the noise term is normally distributed with a standard deviation of 100.

We use an ARIMA model because it is popular in the NVP literature, and we choose a model with two parameters so that we can explore the effects of both over- and under-parametrisation. For a given time series and a given $t = 1, \dots, 200$, we let d_t denote the demand realisation in the t^{th} time period.

Now suppose that we have selected a forecasting model. This can be the correct model, i.e., ARIMA(1,0,1), or an incorrectly specified model, such as AR(1). Suppose also that we have selected the adjustment parameters β and γ . We do the following for each time series:

1. For $t = 21, \dots, 200$, we use the `arima()` function in the `stats` package to produce maximum-likelihood estimates of the mean and standard deviation of demand in time period t . We let $\hat{\mu}_t$ and $\hat{\sigma}_t$ denote these estimates. We note that for $t \leq 20$, the results may be biased due to the shortage of observations.
2. For $t = 21, \dots, 200$, we use τ , $\hat{\mu}_t$ and $\hat{\sigma}_t$ to compute the “textbook” optimal order quantity for time period t using the formula (4.2). We let x_t^* denote this quantity.
3. Finally, we simulate the adjustment process. To avoid systematic bias, we assume that $x_{20} = d_{20}$. For $t = 21, \dots, 200$, we assume that the amount ordered at the start of period t follows formula (4.6).

To quantify the effect of adjustment, we proceed as follows. For a given series and for $t = 21, \dots, 200$, we compute

$$\text{PPL}(x_t) = 100 \left[\frac{\pi(d_t, d_t) - \pi(x_t, d_t)}{\pi(d_t, d_t)} \right]. \quad (4.7)$$

(Here, “PPL” stands for ‘percentage profit loss’; see Liu et al., 2022.) It shows the percentage of profit that would be lost due to using each method instead of knowing the true demand. We also compute the “relative profit improvement” (also known as “forecast value added” in some contexts; see Gilliland, 2010):

$$\text{RPI}(x_t) = 1 - \frac{\text{PPL}(x_t)}{\text{PPL}(x_t^*)}. \quad (4.8)$$

It is a relative measurement comparing the adjusted performance with the unadjusted. Intuitively, the mean of the $\text{RPI}(x_t)$, overall 500 time series, represents the

improvement in profit (if any) gained by using the chosen adjustment in period t . The higher the value is, the better the performance of the approach is. If the value is negative then this means that the approach is worse than the benchmark.

4.4.2 When the DGP is known

We first report results for the case in which the DGP is known, but the model parameters need to be estimated. In particular, we assume that we are using an ARIMA(1,0,1) model, but with unknown parameters.

In Table 4.1, we show the mean RPI for different values of the adjustment parameters. The heading “short dataset” indicates that the mean RPI is computed over the interval $t \in [21, 110]$, and the heading “long dataset” indicates that the mean is computed over the interval $t \in [111, 200]$.

Table 4.1 indicates that a modest amount of adjustment can be beneficial, especially when the number of historical demand observations is short. This is due to the statistical model being unable to estimate its parameters accurately on insufficient data. Therefore, a modest amount of adjustment can provide additional information. On the other hand, too much adjustment leads to a loss. We also find that the RPI is more sensitive to the choices of β , meaning that demand chasing has a bigger influence in this case. We mark that this is also true for other data inputs.

To explore this effect in more detail, we show in Figure 4.1 a plot of average RPI against the length of the dataset, for three different values of (β, γ) , namely, $(0, 0.4)$, $(0.1, 0)$ and $(0.2, 0.1)$. It can be seen that the average RPI is well above zero initially, but decreases, and eventually becomes negative.


A tentative explanation is that the maximum-likelihood estimates of $\hat{\mu}_t$ and $\hat{\sigma}_t$ are prone to errors when the number of observations is small. It may even be that the

Table 4.1: Average RPI with varying adjustment parameters and with short/long datasets ($\tau = 0.7$)

Short dataset		γ				
β	0	0.1	0.2	0.3	0.4	0.5
0	–	3.5%	3.4%	3.3%	3.2%	2.8%
0.1	3.8%	5.0%	4.8%	4.7%	4.6%	3.8%
0.2	4.2%	5.1%	5.0%	4.6%	3.0%	2.4%
0.3	2.5%	2.8%	2.6%	1.5%	0.9%	0.2%
0.4	-2.4%	-1.3%	-1.9%	-3.1%	-3.9%	-5.0%
0.5	-5.1%	-3.9%	-4.4%	-5.9%	-6.8%	-7.7%


Long dataset		γ				
β	0	0.1	0.2	0.3	0.4	0.5
0	–	-0.5%	-0.9%	-1.3%	-1.7%	-2.1%
0.1	-0.7%	-0.2%	-0.5%	-0.7%	-0.9%	-1.3%
0.2	0.1%	0.3%	0.2%	0.1%	-0.2%	-0.7%
0.3	-2.1%	-1.6%	-2.6%	-3.3%	-4.1%	-5.1%
0.4	-6.7%	-6.4%	-7.5%	-8.4%	-9.2%	-10.2%
0.5	-11.1%	-10.5%	-11.8%	-12.8%	-13.8%	-14.7%

Figure 4.1: Average RPI of judgemental adjustments vs. dataset length ($\tau = 0.7$)



estimates suffer from some kind of systematic bias, which decreases over time. By performing a small amount of adjustment, we shift the order quantity toward the true optimal value. This effect vanishes as more data becomes available.


Figure 4.2: Heat map of RPIs for different combinations of adjustment parameters (data length = 20)



In Figure 4.2, we use a heat map to compare the performance of adjustment with different parameter values, for the case in which the number of past observations is exactly 20 (i.e., $t = 21$). As we observed in Figure 4.1, the dominance relationship between parameter pairs (one parameter pair generating higher RPI than the other) is not influenced by data length. Therefore, the choice of $t = 21$ can amplify the results, making it easier for us to observe, without creating additional biases. One can clearly see from Figure 4.2 that large adjustments are harmful, while modest amounts of adjustment, on the other hand, can generate positive RPI. From our data, the most profitable option is to set β to around 0.2 and γ to around 0.05. A possible explanation for this phenomenon is that the “textbook” order quantity will be close to the theoretical one, but it would need to be adjusted by the order quantity and demand on the previous observation. The PtC effect needs to be reduced in

comparison with the DC effect.

Figure 4.3: Boxplots of the RPI for different values of τ and different combinations of adjustment parameters



Next, we examine the effect of τ on the performance of adjustment. Figure 4.3 shows boxplots of the RPI for four different values of τ (0.3, 0.5, 0.7 and 0.9), and the same three values for (β, γ) as before. Here, $t = 21$ as before. Each boxplot shows the range, median and quartiles over the 500 time series.

Interestingly, adjustment yields a benefit in every case, *except* when $\tau = 0.5$. Moreover, we can see in Figure 4.3 that the performance of adjustments when $\tau = 0.3$ is very similar to that when $\tau = 0.7$. This suggests that the effect of adjustment may be symmetric around 0.5. Note also that, even when $\tau = 0.5$, the adjustment does not cause any noticeable loss of profit.

All things considered, it appears that, when the model is correct but the data length is short, a modest amount of naïve adjustment can be beneficial instead of harmful. This goes against the prevailing view in the literature that DC and PtC are invariably damaging. We believe that the discrepancy is mainly due to the model correctness assumption made in the literature. When model correctness is assured,

there is no doubt that any kind of adjustment will be harmful to the profit. Indeed, the behavioural results are in line with the results that we obtained with the longest datasets in our experiment, where there was sufficient data to properly estimate all parameters.

4.4.3 When the model is misspecified

In this subsection, we examine the effect of model misspecification on the relative performance of the judgemental adjustments. We consider three scenarios of model misspecification:


1. The model is under-parametrised (i.e., omits one or more important variables), which typically leads to biased estimates of parameters;
2. The model is over-parametrised (i.e., has one or more redundant variables), which usually leads to inefficient estimates of parameters;
3. The model has the correct parameters, but the assumed distribution of the error term is wrong, which can lead to biased quantile estimates.

For scenario #1, we use an MA(1) model to fit the underlying demand data. For scenario #2, we use an ARIMA(2,0,1) model. For scenario #3, the data is generated using a modified version of ARIMA model, in which the error term follows the Laplace instead of the normal distribution. When estimating the optimal order quantity, however, we use the incorrect assumption that the error term follows the normal distribution.

Since we wish to focus on the effect of model misspecification, rather than the effect of a lack of data (as in the previous subsection), we report the mean RPIs when $t = 200$, when plenty of demand data is available. As before, however, all means are taken over 500 time series.

Figure 4.4 shows the boxplots for scenarios #1 and scenario #2, together with the boxplots for the correctly specified case, for comparison. Here, τ is equal to 0.7. As before, results are reported for three different settings of (β, γ) .

Figure 4.4: Boxplots of the RPI for three different models and three different combinations of adjustment parameters ($\tau = 0.7, t = 200$)




It is apparent that, when the forecasting model omits important variables, naïve adjustment can yield a significant increase in profit. The improvement seems to be present also when the model has redundant variables, but the effect is less pronounced. This is probably due to the nature of the judgemental adjustment. In the case where important information is omitted from the model, there is a good chance that adjustment can provide supplementary information. On the other hand, when the model contains redundant variables, no additional information is needed.

Heatmaps for scenario #1 and scenario #2 are presented in Appendix F. They confirm again our findings in Figure 4.4. It appears that the optimal PtC parameter (γ) is larger than the optimal DC parameter (β) in the under-parametrised case. This is probably because omitting variables in the model leads to an over-estimate of the variance, leading to order quantities that are further from the mean than

required. This bias is remedied by using PtC, since the order quantity is brought closer to the mean of the data. On the other hand, the improvement from adjustment when the model has redundant variables is limited, because the model overfits the data, thus having a lower variance than needed.

Finally, we consider scenario #3, in which the underlying assumption on the error term distribution is wrong. Figure 4.5 shows the RPI boxplots for this case. It can be seen that adjustment significantly improves the expected profit when $\tau = 0.9$, but the effect disappears when $\tau = 0.5$. This can be explained by the relative shapes of the normal and Laplace distributions. Since the largest difference between the distributions is in the “tails”, more adjustment will be needed when τ is at an extreme value (i.e., close to either 0 or 1).

Figure 4.5: Boxplots of the RPI when the assumed distribution of error term is wrong ($t = 200$)



Interestingly, PtC seems to be of more benefit than DC in scenario #3. This is probably because the wrong assumption about the error term distribution leads to a systematic overestimation of the variance. Just as in the under-parametrised case, this over-estimation was alleviated by PtC.

Once again, we find that naïve adjustments can be beneficial instead of harmful. The results are similar to what we found in Subsection 4.4.2 and can be overall summarised as:

1. DC is more useful in the short data length case. This is probably because of the small sample bias, which vanishes on larger samples.
2. PtC is more beneficial than DC in the case of model misspecification. This is because of the systematic underestimation of variance in case of omitted variables or wrong model form.
3. In the case when the DGP is known, the benefits from adjustments only depend on data length and adjustment parameters.
4. The improvement from adjustments appears when the model omits important variables, or the distributional assumption is wrong.
5. When the model has redundant variables, the improvement from both DC and PtC is less pronounced.

In general, we can conclude that DC is likely to be beneficial when there is a shortage of demand data, whereas PtC is likely to bring value when the variance of demand is overestimated.

4.5 Tuning Algorithm

In the previous section, we showed that naïve adjustments can improve the expected profit when the data is insufficient and/or the demand model is misspecified. It is however not clear how one might choose suitable parameter values when faced with a specific NVP instance. In this section, we propose and test a simple heuristic algorithm for parameter “tuning”. We believe that this tuning algorithm may be of

interest to both academics and practitioners. The method is explained in Subsection 4.5.1. In Subsection 4.5.2, we apply our method to a real-life NVP instance, for which the true model is not known. Finally, in Subsection 4.5.3, we present and compare the results with and without the application of adjustment with “tuned” parameters.

4.5.1 Procedure

Let us suppose that we have a forecasting model that, for each period in the training set, $t \in [1, s]$, is able to provide an estimate of the mean demand $\hat{\mu}_t$. Moreover, let us assume that we are also able to estimate the “textbook” optimal order quantity x_t^* . We remind the reader that our adjusted order quantity takes the form:

$$x_t = (1 - \gamma)x_t^* + \gamma\hat{\mu}_t + \beta(d_{t-1} - x_{t-1}).$$

The x_t can be calculated for all $t \in [1, s]$ based on the available mean and actual demand, the “textbook” order quantity and some values of β and γ . To determine the values of parameters, we solve an optimisation problem over the training set. Following suggestions in Ban and Rudin (2019) and Liu et al. (2022), we use a non-standard “loss function” for this purpose. The function is chosen to maximise the profit over the training set, instead of minimising the MSE or MAE in the usual way. That is, we estimate β and γ by maximising the in-sample empirical profit:

$$\max_{\beta, \gamma} \sum_{t=2}^s \pi(x_t^*, \hat{\mu}_t, x_t, d_t). \quad (4.9)$$

We remark that the function to be maximised in (4.9) is continuous and concave. On the other hand, it is not differentiable in general. This means that in order to maximise the profit we need to use derivative-free optimisation algorithms, such as Nelder-Mead (Nelder and Mead, 1965; Rios and Sahinidis, 2013).

4.5.2 Real life example

Here we present an example of the application of our approach to real data. The data we use comes from a medium-sized grocery store which sells a wide range of products, many of which are perishable. It includes daily demands for each product, for a period of around 9 weeks, which ran from mid-October to December. For this study, we selected four typical products with very different data structures and NVP parameters. In particular, we made sure that the selected products have a range of prices, demands and critical quantiles, in order to make the experiment less biased. For reasons of confidentiality, we refer to these products as simply A, B, C and D.

To give the reader some sense of the data, we provide time-series plots in Figure 4.6 and summarise the cost parameters in Table 4.2.

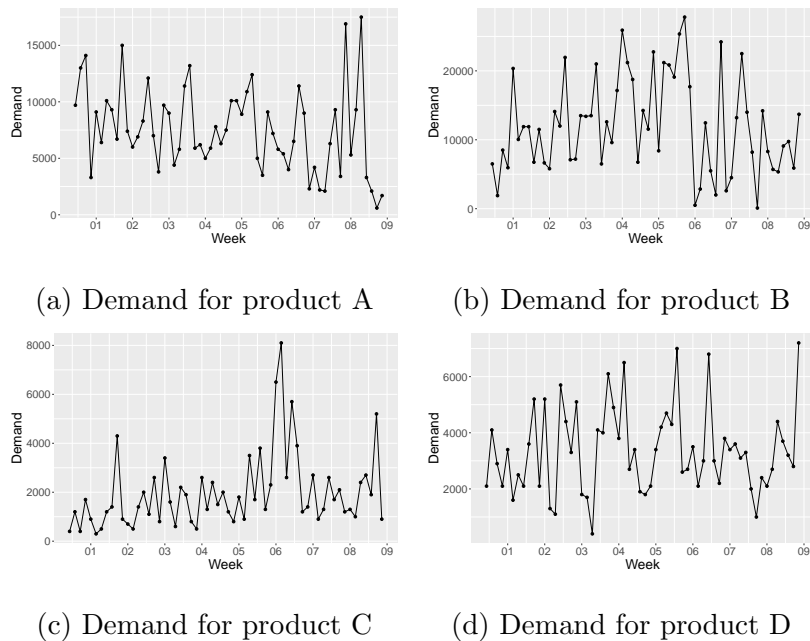


Figure 4.6: Demand time-series for real-life case

Following standard practice in forecasting, we use a rolling-origin method (Tashman, 2000), with constant in-sample size. For each product, on each iteration, we use three-fifths of the data as the training set and perform a one-step-ahead forecast.

Table 4.2: Data for a subset of the products

Products	Price and Costs				Critical Quantile
	p	v	c_h	c_s	
<i>A</i>	2.96	1.28	0.49	0.51	0.55
<i>B</i>	11.98	4.13	2.49	1.33	0.58
<i>C</i>	2.86	1.96	0.78	0.56	0.35
<i>D</i>	4.29	3.24	1.03	0.21	0.23


To perform a fair comparison, and reduce the possibility of bias in our choice of model, we simply applied one of the most popular automatic techniques for forecasting: the `ets()` function from the R `forecast` package (Hyndman et al., 2020). This function attempts to select the most appropriate ETS model, using the Akaike Information Criterion. The pre-tuning decision (x_t^*) is computed using the output from the traditional forecasting procedure, while the tuned decision (x_t) is computed using the output from the forecasting and tuning procedures in combination.

4.5.3 Results

We now present the results obtained with our tuning algorithm. Figure 4.7 displays box plots of the RPI, taken over the iterations, for each of the four products. We remind the reader that a positive RPI indicates that adjustment has been beneficial (see Subsection 4.4.1). The plots indicate that the RPI is positive for all four products in all situations. Thus, the tuned order decisions outperform the pre-tuned ones for all four products.

Table 4.3 summarises performance in terms of service levels. The row labelled ‘target’ shows the critical quantile that maximises the expected profit for the given cost and price parameters. The next two rows show the achieved service level without and with tuning, respectively. In all four cases, the achieved service level with tuning

Figure 4.7: Boxplot of the out-of-sample RPI. The black lines in the boxes represent mean values



is much closer to the target one than the one without tuning.

Table 4.3: Achieved Service level of each method

	Product A	Product B	Product C	Product D
Target	0.55	0.58	0.35	0.23
Pre-tuning	0.75	0.80	0.13	0.17
Tuned	0.65	0.62	0.25	0.29

To gain additional insight, we repeated the entire experiment using four other popular forecasting methods. The results are presented in Table 4.4. From the table, one can see that the tuning algorithm yields a positive out-of-sample RPI in every single case.

The experiment in this section supports our findings in the simulation study. We show that because the true model is not known, the proposed tuning algorithm leads to improvements, bringing the order closer to the correct level.

Table 4.4: RPI results obtained when using adjustment with other forecasting methods

Products	Methods	In-sample	Out-of-sample
Product A	Mean	4.2%	2.8%
	S-Mean	0.4%	2.4%
	S-Naïve	3.0%	2.4%
	ARIMA	1.0%	0.4%
Product B	Mean	3.9%	4.1%
	S-Mean	1.1%	2.9%
	S-Naïve	2.5%	2.8%
	ARIMA	1.9%	1.2%
Product C	Mean	4.1%	4.2%
	S-Mean	2.8%	2.2%
	S-Naïve	2.6%	2.5%
	ARIMA	0.2%	1.1%
Product D	Mean	4.8%	3.1%
	S-Mean	1.2%	3.5%
	S-Naïve	3.3%	4.5%
	ARIMA	2.6%	1.1%

4.6 Concluding Remarks

Although there is considerable literature on judgemental adjustment for newsvendor problems, it has been assumed up to now that ‘demand chasing’ and ‘pull-to-centre’ are especially naïve, and likely to lead to losses in profit. In this chapter, we have shown that, surprisingly, these adjustment procedures can lead to increased profits in some situations. In particular, they can be useful when (a) there is not enough data available to estimate parameters accurately, and (b) the demand model is misspecified. Interestingly, DC appears to be more useful under condition (a), while PtC seems to be of more benefit under condition (b). In general, this is because in

the case of (a) the order estimates suffer from some kind of systematic bias due to short data length; while in situation (b) the estimated variance is often higher than needed.

We also proposed a simple heuristic for tuning the adjustment parameters. Using a real-life example, we show that the tuned orders outperform the pre-tuned ones in terms of the achieved profit, and also led to a service level closer to the target one.

There are several interesting topics for further research. First, one could attempt to characterise other scenarios under which naïve adjustments tend to be beneficial. Second, one could examine the effects of other forms of adjustment. Third, it might be beneficial to conduct behavioural experiments, in the lab and/or field, to confirm the simulation results. Finally, it would be interesting to extend the research to multi-item newsvendor problems, either with or without substitution effects between products.

Chapter 5

News vendor Conditional Value-at-Risk Minimisation on Non-parametric Approach

Next, we turn our attention to alternative objective functions. In most of the literature on NVPs, the objective is to maximise the total expected profit. Some recent works, however, are concerned with the minimisation of the *conditional value-at-risk* (CVaR), a most preferable risk measure in financial risk management. Unfortunately, CVaR estimation involves considering observations with extreme values, which poses problems for both parametric and non-parametric methods. In this chapter, we propose an alternative non-parametric approach to CVaR minimisation with feature-based demand data. We note that our proposed method uses only a small proportion of data and that the empirical risk generated by our method converges to the true risk under suitable assumptions. Using both simulation and real-life case studies, we show that the proposed method can be very useful in practice, allowing the decision-makers to suffer less downside loss in extreme cases while requiring reasonable computing effort.

5.1 Introduction

In this chapter, we focus on *Newsvendor Problems* (NVPs), by which we mean single-period inventory control problems with stochastic demand. In early works on NVPs (Arrow et al., 1951; Morse and Kimball, 1951), it is assumed that the demand in each time period comes from a known probability distribution, and the objective is to determine the order quantity that maximises the expected profit.

Recently, several works have considered a variant of the NVP in which the objective is to minimise the *conditional value-at-risk* (CVaR). The motivation for this is that CVaR is currently a very popular risk measure in financial risk management, as pointed out in Rockafellar and Uryasev (2002). In the work of Gotoh and Takano (2007), a closed-form solution was given for the CVaR-minimisation NVP. Moreover, a mean-CVaR criterion was considered. Then, Jammerneegg and Kischka (2007) proposed an extended model where the inventory manager can control internal and customer-oriented performance measures. Chen et al. (2009), later on, investigated the optimal pricing and ordering decisions in a single framework. The idea of risk aversion in the context of pricing competition was further studied in Wu et al. (2014). Other relevant literature can be found in Abdel-Aal and Selim (2017), Cheng et al. (2009), Wu et al. (2013) and Xinsheng et al. (2015).

In all of the above-mentioned works, it is assumed that the demand comes from a known family of probability distributions with known parameters. In real life, unfortunately, model correctness is rarely assured. Assuming that historical data is available, one can attempt to address this issue by decomposing the problem into a forecasting phase and an optimisation phase, commonly called the disjoint approach (Liu et al., 2022) or SEO approach (Ban and Rudin, 2019). However, if the forecasting model is misspecified, and/or there is substantial noise in the data, then this might impact the optimisation phase in an unexpected way, possibly lead-

ing to sub-optimal solutions, even resulting in nonsensical negative order decisions. Moreover, given that the CVaR concerns observations with extreme values, which are often treated as outliers in traditional statistical approaches, the computed order quantities could underestimate the downside risk and lead to significant losses in extreme cases (Gençay et al., 2003; Yao et al., 2013). The forecasting accuracy can be slightly improved by considering feature-based demand data (Vapnik, 1998), or by using an alternative statistical approach, such as bootstrapping (Efron and Tibshirani, 1994) or extreme value theory (De Haan and Ferreira, 2006). However, the performance depends heavily on the form of the profit function.

To get around these difficulties, one could use a single, non-parametric approach, in which the order quantities are determined directly from the data based on an assumed model or filter. Common non-parametric approaches for classic NVP include sample average approximation (SAA) (Levi et al., 2015), smart “predict then optimise” (SPO) loss (Elmachtoub and Grigas, 2017), NV-features (Ban and Rudin, 2019) and IMEO (Liu et al., 2022). However, none of them is directly applicable to the CVaR minimisation NVP, due to the nonlinearity and non-differentiability of the loss function. Works have been done in the field of CVaR estimation. In the work of Chun et al. (2012), the authors proposed a mixed quantile regression method to estimate the CVaR using a formulation similar to quantile regression. Then, a superquantile regression method was proposed by Rockafellar et al. (2014), derived based on the risk quadrangle. The method of superquantile regression (SQR) was then extended by Harsha et al. (2015) and Miranda (2014), in which the authors considered novel decomposition methods that enable the formulation to be empirically more tractable. Unfortunately, all those methods were originally designed to fit the CVaR of an observable variable (demand itself). With slight modifications, SQR can be applied to the CVaR minimisation NVP, but it is still very sensitive to the choice of profit function and can be extremely computationally expensive under large instances. Moreover, we mark that the SQR may be biased under certain

circumstances.

In this chapter, we propose a *non-parametric feature-based approach of CVaR minimisation*, which we call “NPC” for short. We consider both an empirical model and an adaptive model. We give rigorous proof that under suitable assumptions, the true risk is well estimated by the risk functions of our models. We also perform extensive experiments, on both artificial and real data, to examine the performance of NPC under different settings.

The highlights of our approach include:

1. The approach works directly with historical data and considers features related to the demand, while only requiring a small proportion of data points.
2. The NPC is very robust with regards to the data structures and is adaptive to different forms of profit function, both linear and nonlinear.
3. Experiment results show that the computed order quantities from NPC lead to equal or less downside loss in extreme cases than competing methods.
4. From a statistical viewpoint, the estimated parameters of NPC are statistically explainable and can be easily applied to prescriptive analytics to provide additional operational insights.

The chapter is organised as follows. We review some well-known results on the classic single-period NVP in Section 5.2 and define the profit function of a nonlinear NVP as well. In Section 5.3, the CVaR is introduced in a general form following Rockafellar and Uryasev (2002), and the closed-form solution of the NVP under CVaR minimisation is developed following Gotoh and Takano, 2007. In Section 5.4, we present the method of NPC in detail, including both an empirical model and an adaptive model. We prove that the risk generated by our adaptive model converges to the true risk under suitable distribution assumptions. Sections 5.5 and 5.6 give

computational results on artificial data. Section 5.7, on the other hand, applies NPC to a real-life example. Finally, Section 5.8 contains some concluding remarks.

5.2 Single-Item Newsvendor Problems

In the simplest NVP, as defined, for example, by Choi (2012), a company purchases goods at the beginning of a time period, and aims to sell them by the end of the period. The demand during the period is a random variable \tilde{d} with known probability density function f and cumulative distribution function F . We are also given parameters $c, r, v, g \in \mathbb{Q}$, where

- c is the cost of purchasing one unit of item;
- r is the revenue gained by selling one unit of item;
- v is the disposal cost of each unsold unit of item;
- g is the shortage cost of each unit of unsatisfied demand for item.

We assume without loss of generality that $r > c \geq 0$, $c > -v$ and $g \geq 0$. We permit v to be positive or negative. (A negative value could indicate that excess items can be sold at a discounted price.) We remark that g may be used to represent the “loss of customer goodwill” incurred by stockouts.

The retailer must decide how many units of the item to order before the start of the sales period. We let x denote the number of units ordered. We assume for simplicity that x is continuous. For a given value of x , and a given realisation d of \tilde{d} , the profit over the period is:

$$\pi(x, d) := r \min\{x, d\} - cx - v[x - d]^+ - g[d - x]^+. \quad (5.1)$$

In the classical NVP, the goal is to find a value for x that maximises the total expected profit, which can be given in a closed form (Arrow et al., 1951; Choi, 2012):

$$x^* = F^{-1} \left(\frac{U}{E + U} \right), \quad (5.2)$$

where F^{-1} is the inverse of the distribution function F , $E := c + v$ denotes the overstock cost, and $U := r - c + g$ denotes the understock cost. Moreover, it is easy to prove that the expected profit is a concave function of x (Arrow et al., 1951).

In the general nonlinear NVP, the profit function takes the form:

$$\pi(x, d) := \begin{cases} R(x, d) - C(x, d) - V(x, d), & \text{for } x \geq d \\ R(x, d) - C(x, d) - G(x, d), & \text{for } x < d, \end{cases} \quad (5.3)$$

where R , C , V and G are now *functions* rather than constants.

The nonlinear NVP can be seen as an extension of the classical NVP, as it enables one to model more real-life problems, e.g. with nonlinear shortage cost due to the damage of reputation. The detailed motivation can be found in Khouja (1995), Liu et al. (2022) and Pantumsinchai and Knowles (1991). In general, however, a closed-form solution in terms of a quantile is unlikely to exist for nonlinear NVPs. In such cases, one could resort to numerical integration or simulation techniques to solve the problem.

5.3 NVPs Under CVaR Minimisation

Let $L(x, d) := -\pi(x, d)$ denote the magnitude of the loss for a given realisation d of \tilde{d} and a fixed x , and let

$$\Phi(\eta|x) := \mathbb{P}\{L(x, \tilde{d}) \leq \eta\} \quad (5.4)$$

denote the distribution function of L . We can deduce that $\Phi(\eta|x)$ is a positive, non-decreasing function with $\lim_{\eta \rightarrow -\infty} \Phi(\eta|x) = 0$ and $\lim_{\eta \rightarrow \infty} \Phi(\eta|x) = 1$. For simplicity,

we assume that both x and \tilde{d} are continuous. For $\beta \in [0, 1)$, we define the β -VaR of the distribution by

$$\alpha(x, \beta) := \inf_{\eta \in \mathbb{R}} \{\eta | \Phi(\eta|x) \geq \beta\} = \inf_{\eta \in \mathbb{R}} \{\eta | \mathbb{P}\{L(x, \tilde{d}) \leq \eta\} \geq \beta\}. \quad (5.5)$$

Note that α is a function dependent on β and x . For any $\alpha \in \mathbb{R}$, we can then write β as

$$\beta = \mathbb{P}\{L(x, \tilde{d}) \leq \alpha\} = \Phi(\alpha|x). \quad (5.6)$$

A β -tail distribution function that focuses on the upper tail part of the loss distribution can be formed as (Rockafellar and Uryasev, 2002):

$$\Phi_\beta(\eta|x) := \frac{\Phi(\eta|x) - \beta}{1 - \beta}, \quad \text{for } \eta \geq \alpha(x, \beta). \quad (5.7)$$

Therefore, the β -conditional value-at-risk (β -CVaR) of the loss L can be defined as

$$\psi_\beta(x) := \mathbb{E}_\beta \left[L(x, \tilde{d}) \right], \quad (5.8)$$

where $\mathbb{E}_\beta[\cdot]$ is the expectation operator under the β -tail distribution. We plot an illustrative distribution function of $\Phi(\eta|x)$ and $\Phi_\beta(\eta|x)$ in Figure 5.1. It is easy to see that the β -tail distribution is formed by picking the top $(1 - \beta)$ proportion of $\Phi(\eta|x)$ values, scaling those values by an affine transformation, and setting the rest of the $\Phi(\eta|x)$ values to 0. We remark that the construction of the β -tail distribution is the theoretical foundation for our non-parametric CVaR minimisation method.

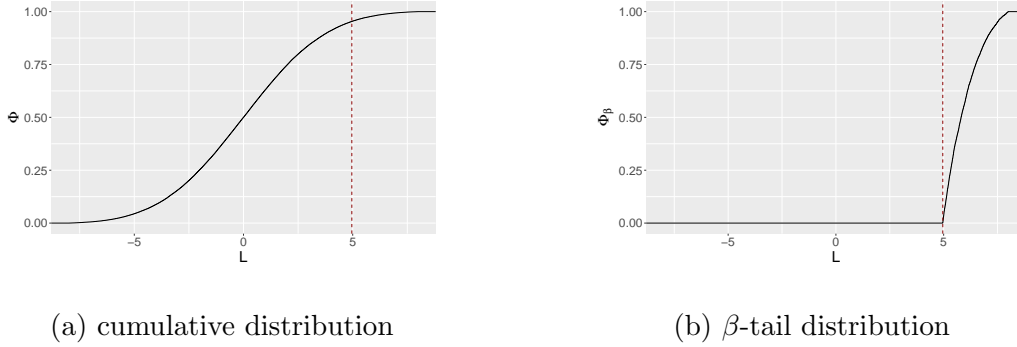
To simplify the procedure for locating α , Rockafellar and Uryasev (2002) defined an auxiliary function:

$$F_\beta(x, \alpha) := \alpha + \frac{1}{1 - \beta} \mathbb{E} \left[[L(x, \tilde{d}) - \alpha]^+ \right]. \quad (5.9)$$

It has been shown in their work that:

$$\min_{x \in X} \psi_\beta(x) = \min_{(x, \alpha) \in X \times \mathbb{R}} F_\beta(x, \alpha), \quad (5.10)$$

Figure 5.1: The cumulative distribution function of $L(x, \tilde{d})$ and the β -tail distribution.



where $X \subseteq \mathbb{R}$ is a feasible region. This relation shows that the minimal value $\psi_\beta(x^*)$ can be achieved by minimising the function $F_\beta(x, \alpha)$ with respect to $x \in X$ and $\alpha \in \mathbb{R}$, simultaneously. With an optimal solution (x^*, α^*) to the right-hand side optimisation problem in Equation (5.10), x^* is an optimal solution of the left-hand side one.

From Equation (5.1), (5.9) and (5.10), the solution to the CVaR version of a (linear) NVP can be given in a closed form (Gotoh and Takano, 2007):

$$\begin{cases} x^* = \frac{E+W}{E+U} F^{-1} \left(\frac{U(1-\beta)}{E+U} \right) + \frac{U-W}{E+U} F^{-1} \left(\frac{E\beta+U}{E+U} \right), \\ \alpha^* = \frac{E(U-W)}{E+U} F^{-1} \left(\frac{E\beta+U}{E+U} \right) - \frac{U(E+W)}{E+U} F^{-1} \left(\frac{U(1-\beta)}{E+U} \right), \end{cases} \quad (5.11)$$

where we recall that $E := c + v$ and $U := r - c + g$, and we set $W := r - c = U - g$.

In particular, when $g = 0$, we have a simpler result:

$$x^* = F^{-1} \left(\frac{U(1-\beta)}{E+U} \right), \quad \alpha^* = -Ux^*. \quad (5.12)$$

We see that the difference between the solutions x^* given by Equation (5.1), (5.11) or (5.12) depends only on two parameters g and β . In particular, when $g = 0$, the difference is only the coefficient in the argument of the inverse F^{-1} . Moreover, when $\beta = 0$, the solutions in Equation (5.1) and (5.11) under CVaR minimisation reduce to the classical expected profit maximisation solution in Equation (5.12).

This consequence is consistent with the definition of the β -CVaR. The difference can be easily visualised, as seen in Figure 3 in Appendix G with artificial data.

5.4 Non-Parametric CVaR Minimisation

As mentioned in Section 5.1, there exist two main issues with the conventional methods for NVPs under CVaR minimisation:

- Given the observations, with extreme values often treated as outliers, the traditional parametric methods could underestimate the downside risk from the tail, resulting in biased order quantities and leading to a significant loss (Gençay et al., 2003).
- The existing non-parametric methods depend heavily on the data structures and the profit function. Moreover, given that they need to consider all historical observations, they can be very computationally expensive for large data sets.

To get around these difficulties, we propose an alternative non-parametric approach. We assume that the historical data are $[(\mathbf{z}_1, d_1), \dots, (\mathbf{z}_s, d_s)]$. For $t = 1, \dots, s$, each $\mathbf{z}_t := [z_t^1, \dots, z_t^p]$ represents p features related to the demand, such as trend, seasonality, prices, promotions and so on. We consider $x = x(\mathbf{z})$. The problem now becomes that of finding a function $x : \mathbb{R}^p \rightarrow \mathbb{R}$ and a value α that optimise the risk function $F_\beta(x, \alpha)$ with respect to the distribution of (\mathbf{z}_t, d_t) . Then, we can evaluate x at a data point in the proceeding period, i.e. $x_{s+1} = x(\mathbf{z}_{s+1})$.

5.4.1 Empirical CVaR minimisation via NPC

We carry out the CVaR minimisation by minimising the auxiliary function (5.9),

$$\min_{x,\alpha} F_\beta(x, \alpha) = \min_{x,\alpha} \left(\alpha + \frac{1}{1-\beta} \mathbb{E} [[L(x, d) - \alpha]^+] \right). \quad (5.13)$$

In practice, the underlying distribution L is often unknown *a priori*. Even if we could find such a distribution for L , the multidimensional integral for the expectation in (5.9) cannot be accurately computed for high dimensional data (Nemirovski et al., 2009). Instead of computing the exact integral, we compute (5.9) in an empirical fashion. The empirical risk minimisation problem can be written as $\min_{x,\alpha} \tilde{F}_\beta(x, \alpha)$ where

$$\tilde{F}_\beta(x, \alpha) := \alpha + \sum_{t=1}^s \frac{[L(x, d_t) - \alpha]^+}{(1-\beta)s}. \quad (5.14)$$

We call $\tilde{F}_\beta(x, \alpha)$ the *empirical risk* and $F_\beta(x, \alpha)$ the *true risk*. Section 4.1 in Von Luxburg and Schölkopf, 2011 proves the following result.

Theorem 5.4.1. *For a fixed x , by the law of large numbers, the empirical risk converges to the true risk as the sample size s goes to infinity, i.e., $\tilde{F}_\beta(x, \alpha) \rightarrow F_\beta(x, \alpha)$ for $s \rightarrow \infty$. The Chernoff inequality also gives a bound which states how likely it is that, the empirical risk is close to the actual risk (Chernoff, 1952): $\mathbb{P} \left(|\tilde{F}_\beta(x, \alpha) - F_\beta(x, \alpha)| \geq c \right) \leq 2e^{-2sc^2}$, where c is any small positive constant.*

Remark 5.4.1. *Theorem 5.4.1 shows that the probability of large deviations of the empirical risk from the true risk decays exponentially as s increases. The data set we used for experiments in Section 5.5–5.6 usually has $s \approx 300$. In this case, the probability that the empirical risk deviates from the true risk by 0.1 is less than 0.4%.*

Corollary 5.4.1. *For a fixed function x , the empirical risk $\tilde{F}_\beta(x, \alpha)$ is an unbiased and consistent estimate of the true risk $F_\beta(x, \alpha)$.*

Proof. Proof of Corollary 5.4.1 To prove that the estimate is unbiased, we use the following equality.

$$\begin{aligned}
\mathbb{E}(\tilde{F}_\beta(x, \alpha)) &= \mathbb{E}\left(\alpha + \sum_{t=1}^s \frac{[L(x, d_t) - \alpha]^+}{(1 - \beta)s}\right) \\
&= \alpha + \frac{s}{(1 - \beta)s} \mathbb{E}[L(x, d_t) - \alpha]^+ \\
&= \mathbb{E}(F_\beta(x, \alpha)).
\end{aligned}$$

Note that the first and last equality is by definition of \tilde{F}_β and F_β respectively. The second equality follows from the linearity of expectation and from the identical, independent nature of d_t . Also, by Theorem 5.4.1, $\tilde{F}_\beta(x, \alpha) \rightarrow F_\beta(x, \alpha)$. This proves that the estimate is consistent. \square

5.4.2 Adaptive CVaR minimisation via NPC

The benefit of using this empirical formulation is that it does not rely on the distribution of demand or the linearity of the profit function. This approach is also less susceptible to modelling bias. However, it still requires the entire data set of historical observation. Thus, the method will still be computationally expensive on large data sets.

To deal with this drawback, we propose an adaptive way of selecting the data for NPC. Instead of minimising empirical risk using the whole data set $\{(z_t, d_t)\}_{1 \leq t \leq s}$, we carefully select a $2 \times (1 - \beta)$ portion of the data and use the reduced data set to minimise an adaptive risk function. (The value of β is normally selected to be 90% or 95% in practice.) In this subsection, we first give a step-by-step explanation of our selection criterion and the adaptive NPC algorithm. We then prove that, under suitable assumptions, the empirical risk function with the reduced data set converges to the true risk function as $s \rightarrow \infty$.

Selection Criteria

Suppose that from the observation $[d_1, \dots, d_s]$, one could decompose the time series into a systematic component, T , and an irregular (noise) component, ϵ . After the decomposition, the set $\{d_t\}_{1 \leq t \leq m}$ corresponds to a set of noise values $\{\epsilon_t\}_{1 \leq t \leq m}$. We can also re-write the loss function as

$$L(x, d_t) = L(\tilde{x}(x, T_t), \epsilon_t) = L(x, T_t, \epsilon_t) \quad \text{for all } 1 \leq t \leq s. \quad (5.15)$$

Now, for a fixed \tilde{x} , we observe that in the NVP, the loss function takes a large value if and only if the noise term ϵ_t takes extreme values. This observation motivates us to design an adaptive selection criterion.

We define the ‘worst’ scenarios as the ‘smallest’ and the ‘largest’ $(1 - \beta)$ proportion of the data in regard to their noise ϵ . We denote the selected noises in ascending order as

$$\mathcal{E} := \{\epsilon_{i_1}, \dots, \epsilon_{i_m}, \epsilon_{i_{m+1}}, \dots, \epsilon_{i_{2m}}\} \quad (5.16)$$

where $m = \lceil (1 - \beta)s \rceil$. The first m items of \mathcal{E} are the m smallest ϵ_t values, and the last m items are the m largest. We define the index set of the chosen data as

$$M := \{i_1, \dots, i_{2m}\}. \quad (5.17)$$

Now, we minimise the adaptive risk function with respect to the reduced data set, i.e.:

$$\min_{x, \alpha} \hat{F}_\beta(x, \alpha), \quad \text{where } \hat{F}_\beta(x, \alpha) := \alpha + \sum_{t \in M} \frac{[L(x, T_t, \epsilon_t) - \alpha]^+}{m}. \quad (5.18)$$

It is worth noting that this adaptive model only requires a $2 \times (1 - \beta)$ proportion of the data, significantly reducing the computational effort. The experiments in the following sections show that this approximate format can outperform benchmark methods easily.

Proof of convergence

In this subsection, we prove that under suitable assumptions, the adaptive risk function with respect to the reduced data set $\hat{F}_\beta(x, \alpha) \rightarrow F_\beta(x, \alpha)$ as $s \rightarrow \infty$. In other words, the adaptive risk function obtained from our carefully selected $2 \times (1 - \beta)$ proportion of data approximates the true risk function.

To complete the proof, we need the following assumptions.

- **Noise Distribution Assumption:** Assume the noises $\{\epsilon_t\}_{1 \leq t \leq s}$ are independent and identically distributed random variables from a distribution with zero mean. Let $\Phi_\epsilon(\eta) = \mathbb{P}(\epsilon < \eta)$ be the cumulative distribution function of the distribution, such that $\lim_{\eta \rightarrow -\infty} \Phi_\epsilon(\eta) = 0$ and $\lim_{\eta \rightarrow \infty} \Phi_\epsilon(\eta) = 1$.
- **Continuity Assumption:** Let $(\mathcal{X}, \mathcal{T}, (-\infty, \infty))$ be the feasible region for the distribution of (x, T_t, ϵ) . The loss function $L(\cdot, \epsilon)$ is continuous with respect to ϵ for all $(x, T_t) \in (\mathcal{X}, \mathcal{T})$.
- **Tail Assumption:** For all $(x, T_t) \in (\mathcal{X}, \mathcal{T})$, we assume that at least one tail of the loss function is unbounded as $|\epsilon| \rightarrow \infty$. Namely, one of the below cases is true,

$$\lim_{\epsilon \rightarrow -\infty} L(x, T_t, \epsilon) \rightarrow \infty, \lim_{\epsilon \rightarrow \infty} L(x, T_t, \epsilon) \text{ is bounded}, \quad (5.19)$$

$$\text{or } \lim_{\epsilon \rightarrow \infty} L(x, T_t, \epsilon) \rightarrow \infty, \lim_{\epsilon \rightarrow -\infty} L(x, T_t, \epsilon) \text{ is bounded}, \quad (5.20)$$

$$\text{or } \lim_{\epsilon \rightarrow \pm\infty} L(x, T_t, \epsilon) \rightarrow \infty. \quad (5.21)$$

The key idea of the proof is that, on the one hand, when we calculate the true risk $F_\beta(x, \alpha) = \alpha + \frac{1}{(1-\beta)} \mathbb{E} \left[[L(x, \tilde{d}) - \alpha]^+ \right]$, only the data points (T_t, ϵ_t) that correspond to $L \geq \alpha$ have an impact on the expectation. On the other hand, using our selection criterion, the $2 \lceil (1 - \beta)s \rceil$ indices selected in M are sufficient to cover the data points that generate non-zero expectation in $F_\beta(x, \alpha)$. The first statement is proven in Lemma 5.4.1 and the second statement is proven in Theorem 5.4.2.

Lemma 5.4.1. For a fixed x , we consider the data set $\{T_t, \epsilon_t\}_{1 \leq t \leq m}$. Let α be the **risk threshold** such that exactly $\lceil (1 - \beta)s \rceil$ values of the loss function $\{L(x, T_t, \epsilon_t)\}_{1 \leq t \leq m}$ have a larger value than α . We denote the index set as S , such that

$$S := \{t | L(x, T_t, \epsilon_t) \geq \alpha\}. \quad (5.22)$$

Then,

$$\tilde{F}_\beta(x, \alpha) = \alpha + \frac{1}{|S|} \sum_{t \in S} (L(x, T_t, \epsilon_t) - \alpha) \rightarrow F_\beta(x, \alpha). \quad (5.23)$$

where $|S|$ denotes the size of a set S .

Proof. Proof of Lemma 5.4.1 We simplify the expected value in $F_\beta(x, \alpha)$,

$$\begin{aligned} \mathbb{E} \left[[L(x, \tilde{d}) - \alpha]^+ \right] &= \mathbb{E} \left[[L(x, \tilde{d}) - \alpha]^+ | L(x, \tilde{d}) \geq \alpha \right] (1 - \Phi(\alpha|x)) \\ &\quad + \underbrace{\mathbb{E} \left[[L(x, \tilde{d}) - \alpha]^+ | L(x, \tilde{d}) \leq \alpha \right]}_{=0} \Phi(\alpha|x) \\ &= \mathbb{E} \left[(L(x, \tilde{d}) - \alpha) | L(x, \tilde{d}) \geq \alpha \right] (1 - \Phi(\alpha|x)) \end{aligned}$$

where Φ was defined in (5.4). From (5.6), we deduce that the expected size for S is $(1 - \beta)s$ or, in the integer case, m . In view of Theorem 5.4.1, for any fixed x , we can use an empirical estimation to approximate $F_\beta(x, \alpha)$

$$\frac{1}{|S|} \sum_{t \in S} (L(x, T_t, \epsilon_t) - \alpha) \rightarrow \mathbb{E} \left[(L(x, \tilde{d}) - \alpha) | L(x, \tilde{d}) \geq \alpha \right] \quad \text{as } s \rightarrow \infty \quad (5.24)$$

Note that

$$\begin{aligned} \mathbb{E} \left[(L(x, \tilde{d}) - \alpha) | L(x, \tilde{d}) \geq \alpha \right] &= \frac{1}{(1 - \Phi(\alpha|x))} \mathbb{E} \left[[L(x, \tilde{d}) - \alpha]^+ \right] \\ &= \frac{1}{(1 - \beta)} \mathbb{E} \left[[L(x, \tilde{d}) - \alpha]^+ \right]. \end{aligned} \quad (5.25)$$

Adding α on (5.24)–(5.25) and substituting $|S| = m$ gives the result. \square

Theorem 5.4.2. Under the noise distribution, continuity and tail assumptions, for any fixed x , there exists a $\beta \in (0, 1)$ such that

$$\hat{F}_\beta(x, \alpha) \rightarrow F_\beta(x, \alpha), \quad \text{as } s \rightarrow \infty. \quad (5.26)$$

Proof. Proof of Theorem 5.4.2 **Step 1: The monotonically increasing tail of loss function.** In this proof, we assume without loss of generality that the tail assumption holds in the case of $\lim_{\epsilon \rightarrow \infty} L(x, T_t, \epsilon) \rightarrow \infty$ only. The cases of $\lim_{\epsilon \rightarrow -\infty} L(x, T_t, \epsilon) \rightarrow \infty$ only and $\lim_{\epsilon \rightarrow \pm\infty} L(x, T_t, \epsilon) \rightarrow \infty$ follow similarity.

The loss function is unbounded as $\epsilon \rightarrow \infty$ and continuous with respect to $\epsilon \in \mathbb{R}$. There exists a positive constant $C_1 \in \mathbb{R}$ large enough, such that for any $C > C_1$, there exists an $\epsilon_c > 0$, such that $L(x, T_t, \epsilon) \geq C$ for all $\epsilon \geq \epsilon_c$ and $L(x, T_t, \epsilon) < C$ for all $\epsilon \leq \epsilon_c$. Specifically, $L(\cdot, \cdot, \epsilon)$ is a bijective map for $\epsilon \in (\epsilon_c, \infty)$ and $L \in (C, \infty)$, such that

$$L(x, T_t, \epsilon) \geq C \quad \Leftrightarrow \quad \epsilon \geq \epsilon_c. \quad (5.27)$$

Step 2: The risk threshold. Let $\alpha(\beta)$ be the risk threshold defined as in Lemma 5.4.1. We write the CDF of ϵ as

$$1 - \Phi_\epsilon(\alpha) = \mathbb{P}(\epsilon \geq \alpha(\beta)) = 1 - \beta. \quad (5.28)$$

By the distribution assumption, there is a $\beta \in (0, 1)$, such that we can generate a risk threshold $\alpha > C_1$.

Step 3: Relationship between S and M . Now consider the set $\{\epsilon_t\}_{t \in S}$. For all $t \in S$, we have $L(x, T_t, \epsilon_t) \geq \alpha > C_1$. By (5.27), we deduce that $\epsilon_t \geq \epsilon_j$ where $t \in S$ and $j \notin S$. Specifically, S contains the indices of the largest $\lceil (1 - \beta)s \rceil$ values among $\{\epsilon_t\}_{1 \leq t \leq s}$. This corresponds to indices $\{i_{m+1}, \dots, i_{2m}\}$ in M (where M is as defined in (5.17)). So we deduce that $S \subset M$.

Step 4: Convergence.

$$\begin{aligned}\hat{F}_\beta(x, \alpha) &= \alpha + \frac{1}{m} \sum_{t \in M} [L(x, d_t) - \alpha]^+ \\ &= \alpha + \frac{1}{m} \left[\sum_{t \in S} [L(x, d_t) - \alpha]^+ + \underbrace{\sum_{t \in M/S} [L(x, d_t) - \alpha]^+}_{=0 \text{ by definition of } S} \right] \\ &= \alpha + \frac{1}{m} \sum_{t \in S} (L(x, d_t) - \alpha) \rightarrow F_\beta(x, \alpha)\end{aligned}$$

as $s \rightarrow \infty$. The last line was proven in Lemma 5.4.1. \square

Numerical examples and implementation


We illustrate Theorem 5.4.2 with some numerical examples. We generate the demands d_t with the noise values ϵ_t following a mean zero normal distribution. It is straightforward to verify that such a distribution satisfies the noise distribution assumption. We consider a linear loss function and a nonlinear loss function in Figure 5.2. Both loss functions satisfy the continuity assumption and the tail assumption. As a result, we see that for all plots, the set S (shaded area) is contained in the set M (the region bounded by dashed lines and the vertical edges of the graphs). These examples confirm our proof in Theorem 5.4.2.

One last remark is that in the adaptive NPC method, we will also have reduced computational complexity for the function x . In the simplest case, we write x in the following form,


$$x(\mathbf{z}_t) := \mathbf{z}_t^\top \boldsymbol{\gamma} = \sum_{j=1}^m z_t^j \gamma^j, \quad (5.29)$$

where $\boldsymbol{\gamma} \in \mathbb{R}^p$, together with α , are the parameters to be optimised in the CVaR minimisation. In this case, the estimated parameters can be viewed as the ‘effective ratio’ of given features (e.g. the order should be increased by γ^j unit in order to achieve minimum CVaR if the feature z^j is increased by one unit). Other representations of x can be polynomials (e.g. with quadratic regularisation terms) or


Figure 5.2: Illustrative example for Theorem 5.4.2. The linear loss function as in Equation (5.35) and a nonlinear loss function as in Equation (5.37).




(a) Linear loss function, $x = 5$



(b) Linear loss function, $x = 8$



(c) Nonlinear loss function, $x = 5$



(d) Nonlinear loss function, $x = 8$

lags (e.g. with previous observations). Since L can be a function of any level of complexity, minimising (5.18) is, in general, a continuous nonlinear optimisation problem. Under reasonable assumptions on the functions R , C , V and G in (5.3), and the function $x(\cdot)$ itself, the function (5.18) will be convex, but not necessarily everywhere differentiable. Unfortunately, general-purpose algorithms for nonlinear optimisation are not guaranteed to converge to a global minimum, due to the lack of everywhere-differentiability. Fortunately, the experiments in Section 5.5 and Section 5.6 indicate that this does not cause serious problems.

5.5 Baseline Experiment

In order to assess the performance of the proposed method (adaptive NPC), and to understand its strengths and weakness, we conduct simulation experiments in R

4.2.1 with an Apple M1 Pro (2021) machine. In Subsection 5.5.1, we discuss the setup for our baseline experiment. In Subsection 5.5.2 the simplest case is studied, in which the profit function is linear. The case in which the profit function is nonlinear is discussed in Subsection 5.5.2.

5.5.1 Experimental setup

For our baseline experiments, we consider NVPs with artificial demand data, and we suppose initially that there are 4 features related to the demand, each containing 500 observations (the cases with other numbers of features will be discussed later). We generate each feature from a seasonal ARIMA process and we generate the demand as:

$$d_t := b_0 + b_1 z_t^1 + b_2 z_t^2 + b_3 z_t^3 + b_4 z_t^4 + r_t, \quad (5.30)$$

where z_t^p is the realisation of feature p at time t , and r_t is a realised error generated by an additive (weighted) mixture of `rnorm()`, `rlaplace()` and `rt()` functions. The choice of b_p for the features, ϕ and θ for the ARIMA process, and other parameters for the generation of error terms are all selected randomly. We choose a seasonal ARIMA model since it is one of the most popular statistical models in the literature (for example, see Syntetos et al., 2016). The detailed parameter values for the baseline setup and the demand series can be seen in Table 5 and Figure 4 in Appendix H. The results of the experiment with other parameter values will be discussed in Section 5.6. The scripts of all experiments have been made available on Github (Liu, 2022).

All experiments are performed on a rolling-origin basis with 1-step ahead order forecasts (Tashman, 2000), in which we fix the *origin size* (holdout sample size) to be 50, 100, 150, 200, 250 or 300, and the *iteration number* (number of shifts) to be 50, 100, 150 and 200. For each pair of *origin size* and *iteration number*, we use the following two quantities as measurements:

-
- β -Downside Loss (β - DL) = $\frac{1}{n} \sum_{t=1}^n L_t$: This measures the average value of the largest $(1 - \beta)$ cases of losses, where $n = \lceil (1 - \beta) * \textit{iteration number} \rceil$, and L_t is ranked in descending order. It is desirable for this value to be as small as possible.
 - Service Level (SL) = $\frac{1}{N} \sum_{t=1}^N \mathbb{I}(x_t \geq d_t)$: This measures the proportion of cases in which the demand is successfully fulfilled, where $N = \textit{iteration number}$, and $\mathbb{I}(\cdot)$ is the indicator function. In the ideal situation, SL should be as close to the target service level as possible.

In the proposed method (“NPC”), we use the `optim()` function from the `stats` package for R and the Limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm (L-BFGS) for the estimation of the parameters of the model. The L-BFGS algorithm has been shown to perform well in similar nonlinear programming tasks in an NVP context (Liu et al., 2022; Liu and Nocedal, 1989).

To get around the scale issues, we consider two benchmark methods, with which we can compute the relative β - DL and relative SL (note that the “instability” issue and the “negativity” issue of relative measurement do not incur in our experiment):

- The benchmark method - Sample weighted average (“SA”): With historical demand $[d_1, \dots, d_s]$, the order quantity x_{s+1} is set to be a weighted average of empirical quantiles under Equation (5.11).
- The benchmark method - Under correctly specified model (“UM”): A method uses `lm()` function from `stats` package for R to forecast the next period demand considering all features and all observations, and determines the order quantity with Equation (5.11). We make sure the distribution of the error terms is correctly assumed.

Besides the benchmarks, three competing methods are also considered:

-
- A non-featured method (“NF”) that applies the `auto.arima()` function from `forecast` package for R to the demand series itself in the forecasting phase, and determines the order quantity with Equation (5.11).
 - A superquantile regression method (“SQR”) that uses `rq()` to determine the order quantity (Rockafellar et al., 2014).
 - A method (“PLM”) that uses `lm()` to forecast but with only observations from the ‘worst’ $2 \times (1 - \beta)$ proportion of scenarios (the term ‘worst’ is defined in Subsection 5.4.2).

In the baseline experiment, we consider a linear profit function

$$\pi(x, d) := 20 \min\{x, d\} - 8x + 3[x - d]^+ + 7[d - x]^+, \quad (5.31)$$

and a nonlinear profit function

$$\pi(x, d) := 20 \min\{x, d\} - 8x - 4[x - d]^+ + 5 \mathbb{E}[\min\{[x - d]^+, u\}] - 0.01 ([d - x]^+)^2, \quad (5.32)$$

where $u \sim \mathcal{N}(30, 5^2)$. These settings are consistent with the work of Liu et al. (2022). The optimal service levels to maximise the expected profit for these two functions are 0.35 and 0.62.

5.5.2 Results of the baseline experiment

Here, we present the results from our experiment, where the parameters described in Subsection 5.5.1 are used.

To get some sense of the experimental procedure and the calculation of the relative measurement, we first present an example in Appendix I, where only two methods are considered, NPC and SA. The relative β -DL and relative SL in this example

can be calculated as:

$$\text{relative } \beta\text{-DL} = \frac{DL_{SA} - DL_{NPC}}{DL_{SA} - DL_{UM}} = 93\% \quad (5.33)$$

$$\text{relative } SL = 1 - \left| \frac{SL_{NPC} - SL_{UM}}{SL_{SA} - SL_{UM}} \right| = 20\% \quad (5.34)$$

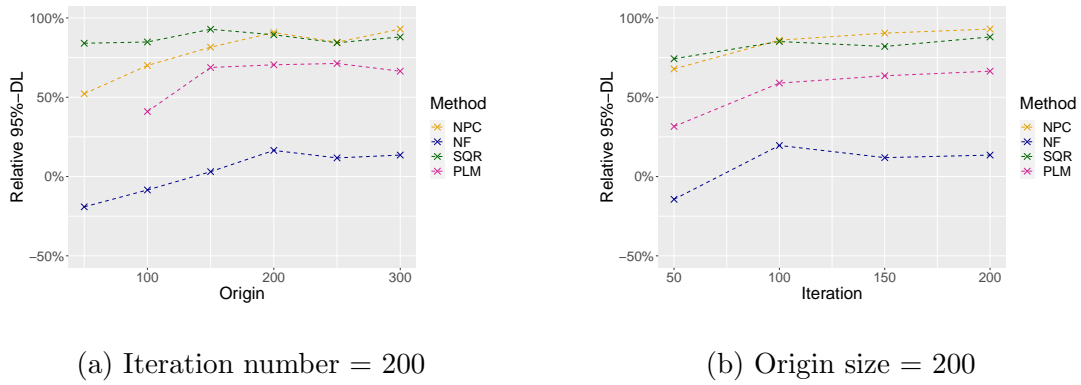
These values can be interpreted as: by using NPC, the decision makers will suffer 93% less loss than in the case of SA in the worst 5% of scenarios, and the service level they achieve will be 20% closer to the target service level. Here we note that the service level achieved by CVaR minimisation may lie far away from the service level achieved by expectation maximisation, due to their different natures. In practice, a decision maker will need to balance the benefit of reducing downside loss and the harm of decreasing service level (as they are inseparable in most cases).

Linear profit function

In Figure 5.3a, we present the relative $\beta\text{-DL}$ from our baseline experiment, where multiple choices of *origin size* are considered under the given linear profit function when *iteration number* = 200. We remark that the result from the PLM method under *origin size* = 50 is excluded from the plot, as it is far below zero (This is probably due to the drawback of using Least-squares estimation under a small sample size). In general, we see that both the proposed NPC method and the SQR method achieve a high relative $\beta\text{-DL}$. In fact, their performance is quite close, even though SQR uses all historical observations and NPC only uses a small proportion of them. The results from the other two methods are less appealing, as PLM generates very frustrating performance when *origin size* = 50, and NF barely improves the loss compared to the benchmark SA method.

In Figure 5.3b, we focus on the case of *origin size* = 300, and present the relative $\beta\text{-DL}$ under multiple choices of *iteration number*. We can see that the results are very similar to what we found in Figure 5.3a, where the NPC and SQR methods

Figure 5.3: Relative 95%-DL under linear profit function



outperform the other two. In Appendix J, we present the results from all other choices of *origin size*, *iteration number* and β in detail, where we include the relative *SL* as well.

We note that, as we are using relative measurements, the results seem to be “stable” among all choices of parameters. This is to be expected, given that the absolute performance of all methods is influenced by parameters at the same time. From the results, we can say that the NFC method shows very strong robustness, as its performance is very close to the SQR method (and the UM method) under all cases in regard to the relative β -DL, even though it only uses a proportion of data. Using the same amount of data, the PLM method, however, performs poorly in most cases.


The results are not unexpected. As there is no significant upward and/or downward trend in the demand series, as seen in Figure 4 in Appendix H, it is totally understandable why the two non-featured methods, SA and NF, perform similarly. (Though the performance of NF improves slightly as the *origin size* and *iteration number* increase.) For the method of PLM, the nature of its loss function is to minimise the MSE, leading it to be under-fitted when the data is limited. With the same amount of data, the NPC method, on the other hand, adopts a different loss function and focuses on extreme scenarios, making good use of all the selected data. The SQR method also performs well in this experiment. However, as we can see

from Figures 5.3a, 5.3b and Table 6 in Appendix J, it gets slightly outperformed by NPC when *origin size* and/or *iteration number* is large. This is presumably due to the presence of bias mentioned in Section 5.1. This bias is amplified when the profit function is nonlinear and/or the error term distribution is changed, as we will see in Subsection 5.5.2 and 5.6.3.

Nonlinear profit function

Here, we present the results from our baseline experiment with a nonlinear profit function. As one can see from Equation (5.31) and Equation (5.32), the major differences between these two forms of profit function concern the penalties incurred by disposal and shortage. Instead of a fixed disposal cost, we now allow the excess products to be sold on a salvage market. Instead of a fixed shortage cost, we consider a quadratic cost function. As CVaR minimisation focuses on extreme cases, these differences may be amplified in our experiments and lead to results very different from those of Subsection 5.5.2. Given that a closed-form solution does not exist for the given nonlinear function, one can use the technique proposed by Kyparisis and Koulamas, 2018, or other numerical approaches, to verify that the quantiles to minimise 95%-CVaR and 90%-CVaR are approximately 0.13 and 0.16.

Figure 5.4: Relative 95%-DL when iteration number = 200 under nonlinear profit function



In Figure 5.4, we present the relative β -DL with multiple choices of *origin size* when *iteration number* = 200. It can be seen from the figure that the NPC method outperforms all other methods in regard to relative β -DL under all values of *origin size*. Moreover, we find that the relative performance of the SQR method decreases as *origin size* increases. This can be further investigated by looking at the absolute performance in Appendix J. We see that the β -DL from SQR method does not improve as *origin size* increases, while the β -DL from SA does, leading to an overall decrease in relative β -DL. One possible explanation for this phenomenon is that, in the SQR method, the loss function targets the extreme demand realisation instead of the extreme profit realisation directly. Therefore, under the nonlinear relationship between demand and profit, this loss function could be heavily biased. Thus, it is no surprise that the performance of SQR does not improve when increasing *origin size*. On the other hand, the NPC targets extreme profit realisation.

We would like to stress that, unlike the parametric methods, NPC does not need any complicated numerical optimisation or simulation methods to estimate the optimal order quantity – it does that directly. In addition, NPC requires only a proportion of data under the selection criterion, yielding results in a more efficient way. Overall, we see that NPC performs at least as well as SQR under linear profit functions while outperforming all other methods under nonlinear profit functions. We examine the robustness of the NPC method in the next section.

5.6 Experiments With Other Parameters

In this section, we extend our experiment to other parameters. In particular, we vary the number of features to be considered in Subsection 5.6.1. Then, we present results with other profit functions in Subsection 5.6.2. Finally, we consider other forms of the error term in 5.6.3. We remark that we have also experimented with

other data-generating models, e.g. ETS, TBATS. We do not present the results here, as they are very similar to the ones presented below.

5.6.1 Varying the number of features

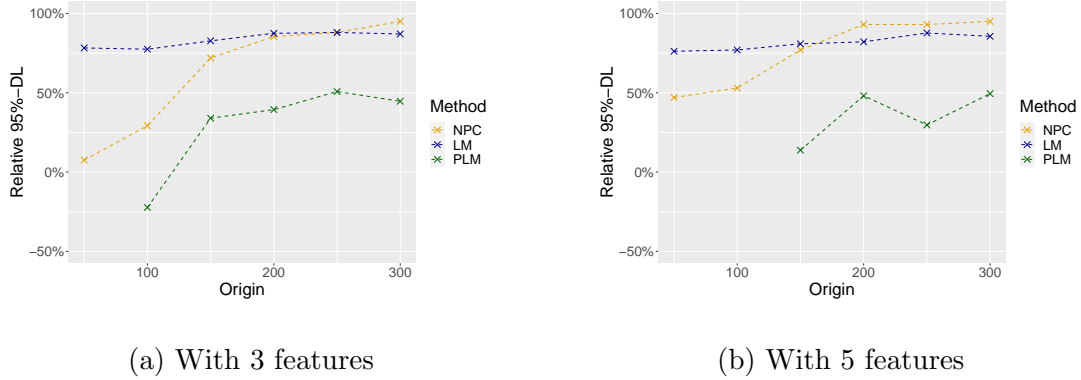
Now, we focus on the number of features. In particular, we consider the number of features to be adopted by the method, instead of the overall feature numbers (as this has negligible impact). This is motivated by the fact that, in reality, decision-makers are rarely able to guarantee the quality of feature choices (Heinze et al., 2018). Therefore, it makes sense for us to consider the performance of our proposed method in the case of model misspecification. To do that, we consider the relative β -DL of the NPC method, the PLM method and one other method:

- A regression method (“LM”) that uses `lm()` to forecast with the same number of features as used in NPC.

Besides, we also make sure that the PLM method uses the same number of features as used in NPC and LM. We remark that the NF method and the SQR method are excluded from this comparison, for the obvious reason that they do not require any features in the computation. Without changing other settings, we now consider the cases where the method uses 3 features or 5 features instead, while using the same data set as before. These represent the cases of model under-fitting and model over-fitting, respectively. To avoid redundancy, here we only present the results with a linear profit function, as the results with a nonlinear profit function were very similar.

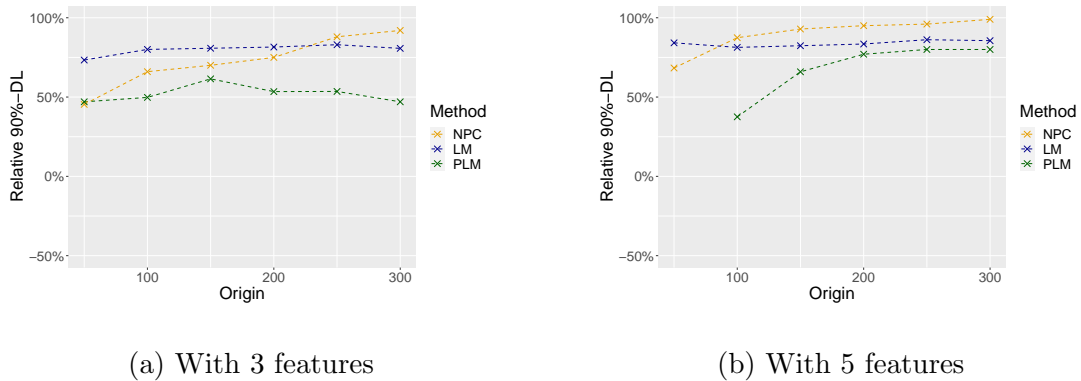
We see from Figure 5.5 that NPC performs better than PLM for all origin sizes, in both the under-fitting and over-fitting cases. However, its performance is worse than LM when *origin size* is small, especially in the under-fitting case. The performance

Figure 5.5: Relative 95%-DL when iteration number = 200 under linear profit function with other numbers of features



improves as *origin size* increases. This is not completely unexpected. As the NPC method uses only a small proportion of the data, it could be more vulnerable than other methods when *origin size* is small, especially when some information is missing due to under-fitting. Fortunately, we can see that using the same amount of data, the performance of NPC is significantly better than the performance of PLM.

Figure 5.6: Relative 90%-DL when iteration number = 200 under linear profit function with other numbers of features



In Figure 5.6, we present the results where $\beta = 90\%$. In this setting, more data is used in NPC and PLM methods. We see that the performance of NPC is still slightly worse than LM when *origin size* is small, but the gap is much smaller than in the case when $\beta = 95\%$. Besides, we find that the NPC method outperforms LM

as long as the *origin size* is larger than 100 in the over-fitting case, and 250 in the under-fitting case. We remark that the results with other values of *iteration number* were very similar to the case when *iteration number* = 200. Therefore, we do not present them here.

To sum up, we find that the proposed NPC method is more vulnerable than other methods when *origin size* is small, especially when the model is under-fitting. Luckily, this drawback is tolerable, as our motivation in proposing an alternative method was to reduce the computational effort with large instances. Even in the case when *origin size* = 300 (where NPC outperforms LM), the NPC method requires only 30 observations with $\beta = 95\%$, fewer than that required by LM when the *origin size* = 50.

5.6.2 With other profit functions

In the previous subsections, we tested the performance of our approach with one linear profit function and one nonlinear profit function, under different conditions. In this subsection, we consider four additional profit functions, two linear and two nonlinear, to examine the sensitivity of our method to the parameters of the profit function. All other settings are consistent with our baseline experiment. We call the functions in the baseline experiment “Linear 0” and “Nonlinear 0”, and we define “Linear 1”, “Linear 2” and “Nonlinear 1” and “Nonlinear 2” as follows:

- Linear 1:

$$\pi(x, d) = 20 \min\{x, d\} - 8x - 3[x - d]^+ - 7[d - x]^+. \quad (5.35)$$

- Linear 2:

$$\pi(x, d) = 20 \min\{x, d\} - 8x + 7[x - d]^+ + 3[d - x]^+. \quad (5.36)$$

-
- Nonlinear 1:

$$\pi(x, d) = 20 \min\{x, d\} - 8x - 4[x - d]^+ - 0.01 ([d - x]^+)^2. \quad (5.37)$$

- Nonlinear 2:

$$\pi(x, d) = 20 \min\{x, d\} - 8x + 5 \mathbb{E}[\min\{[x - d]^+, u\}], \quad (5.38)$$

where $u \sim \mathcal{U}(0, 15)$.

“Linear 1” and “Linear 2” are consistent with the work of Liu et al. (2022), while “Nonlinear 1” and “Nonlinear 2” are derived from it. The optimal service levels to maximise the expected profit for these four functions are 0.63, 0.9, 0.56 and 0.71, respectively. Although the order quantity that minimises the CVaR is usually very different from the one that achieves maximum expected profit, we find the optimal service levels of functions influence the performance of NPC.

As before, we present the Relative 95%-DL of all methods when *origin size* = 50 and *iteration number* = 50 with a linear profit function, as well as the case when *origin size* = 300 and *iteration number* = 200. We can see from Table 5.1 that the NPC method perform well under all linear and nonlinear profit functions, and its performance converges to the UM method when *origin size* and *iteration number* increase. Specifically, the NPC method achieves better results when the optimal service level of the profit function is away from 0.5. This normally means that either the overstock cost or the understock cost is much higher than the other. It appears that, when the profit function is heavily skewed, the NPC method is more efficient than competing methods to prevent downside loss in extreme cases. The phenomenon can be explained by the difference between the nature of the NPC method and competing methods. All competing methods compute results indirectly, as they work with extreme demand observations first and apply the output to the risk function second. However, the method of NPC works with extreme risks directly.

Table 5.1: Relative 95%-*DL* under other profit functions (negative values are excluded)

		Methods			
Origin size = 50/Iteration number = 50		NPC	NF	SQR	PLM
Linear 0 (0.35)		60%	/	68%	/
Linear 1 (0.63)		50%	3%	60%	13%
Linear 2 (0.9)		66%	28%	57%	/
Nonlinear 0 (0.62)		85%	/	1%	/
Nonlinear 1 (0.56)		56%	4%	5%	16%
Nonlinear 2 (0.71)		77%	/	20%	/

		Methods			
Origin size = 300/Iteration number = 200		NPC	NF	SQR	PLM
Linear 0 (0.35)		90%	/	72%	47%
Linear 1 (0.63)		84%	13%	67%	30%
Linear 2 (0.9)		94%	1%	74%	27%
Nonlinear 0 (0.62)		93%	/	4%	/
Nonlinear 1 (0.56)		92%	10%	49%	75%
Nonlinear 2 (0.71)		97%	54%	39%	43%

5.6.3 With other forms of the error term

Finally, we consider the influence of the error term. In our baseline experiment, the error term was generated by a mixture of `rnorm()`, `rlaplace()` and `rt()` functions with random parameters. Therefore, we have not yet examined how the proposed NPC method performs in the presence of heavy tails or light tails. Given the drawbacks of traditional parametric methods on treating outliers, we could expect the gap in performance between NPC and PLM to be larger with light-tailed error terms than with heavy-tailed ones. We don't focus on the comparison of NPC and SQR in this case as they are influenced by the distribution of error term in a similar way. The experiment is conducted with some additional instances. We call the error

term in the baseline experiment “Error 0”, and we define “Error 1” and “Error 2” as follows:

- Error 1: We use `rnorm()` with $\mu = 0$ and $\sigma = 100$ as a light-tail case.
- Error 2: We use `rt()` with $\mu = 0$, $\sigma = 100$, $\nu = 5$ as a heavy-tail case.

We remark that it is not possible for the decision-maker to know the exact distribution of the error term *a priori* in reality. Therefore, in our experiment, we let our parametric methods assume that the distribution is normal in all cases, and we make sure that in our setting, the variance is the same in each instance.

Table 5.2: Relative 95%-*DL* under other forms of the error term (negative values are excluded)

		Methods			
Origin size = 50/Iteration number = 50		NPC	NF	SQR	PLM
	Error 0	60%	/	68%	/
	Error 1	71%	/	87%	/
	Error 2	55%	/	82%	/
		Methods			
Origin size = 300/Iteration number = 200		NPC	NF	SQR	PLM
	Error 0	86%	/	72%	49%
	Error 1	90%	6%	85%	40%
	Error 2	84%	9%	84%	60%

The results in Table 5.2 meet our expectation, as the gap of performance between NPC and PLM is indeed very large in “Error 1”, and it is smaller in “Error 2”. A possible explanation is that the NPC method works directly with the data, and does not rely on the assumption of normality, while the PLM, using the same amount of data, relies on normality. The gap is largest under light tails, since the downside loss in extreme cases in light tail distribution is more likely to be treated as outliers than

in heavy tail by traditional methods like PLM. Moreover, considering NPC itself, its performance improves as data size becomes larger and it performs best against other methods when the distribution of the error term is light-tailed. Although SQR performs slightly better than NPC when the data size is small, the NPC outperforms again once the data size becomes large.

5.6.4 Sum up

From the results of the experiments, we find that:

1. Compared to competing methods, NPC performs better when the data size is large. Due to the fact that NPC only requires a small proportion of data, it can significantly reduce computational effort under large data sets.
2. The NPC performs well with nonlinear NVP, since it does not make any assumption on the linearity of the functions and works directly on historical data.
3. NPC is more vulnerable than other methods when the data size is small, especially when some information is missing due to under-fitting.
4. The proposed method performs best, compared to the competing methods when the profit function is heavily skewed and/or when the distribution of the error term is light tailed. In both cases, the traditional methods are likely to treat extreme data as outliers and underestimate the downside risk.

5.7 Real-life Example

In this section, we examine the performance of NPC with a real-life example within a food bank. A food bank is an emergency feeding organisation providing hunger relief

to families living in poverty. Each food bank covers a given region, and the decision maker has to prepare food on a weekly basis for its distribution day (normally a Sunday). The food preparation problem within the food bank can be approximately fitted by the nonlinear NVP model. The goal of the problem is to determine the amount of food to prepare that fulfils the demand. In the simplest case, we assume the consumption of each individual is the same, and we can just use the 'number of visits' as our demand. Moreover, we assume both x and \tilde{d} (under the same scale) to be continuous. Yet, we should note in particular that:

1. The demand in food banks normally has a smaller variance than the demand considered in other classic inventory management problems. Thus, instead of the expected profit, the CVaR is more of our interest.
2. The opportunity cost of overage is linear, as the food bank can easily get rid of the leftovers. However, the cost of underage is believed to be quadratic.

Derived from Davis et al. (2014) and Riches (2018), this problem can be approximated as:

$$\pi(x, d) = \eta[x - d]^+ + \zeta([d - x]^+)^2. \quad (5.39)$$

where η denotes the overage opportunity cost, including but not limited to transportation fee, management cost and disposal fee; ζ denotes the underage opportunity cost, including but not limited to loss of goodwill and additional management cost. The objective is to minimise π , which we call the 'unit of risk' for generality. The parameter values in our application were estimated to be:


- $\eta = 15, \zeta = 1$.

The data we use comes from a local food bank in Durham. It includes the total number of visits on each distribution day for 104 weeks from July 2020 to June 2022.

We also consider 10 relevant features within the same time scope, as seen in Table 8 in Appendix K.

To get some sense of the data, we provide a time-series plot for the number of visits in Figure 5.7. It can be seen that the number of visits to the food bank shows multiple levels of seasonality, monthly and seasonally, and that, rather surprisingly, the number of visits in winter (weeks 10-30 and 60-80) is lower than in the rest of the year. We think that this could be due to substitution effects from other forms of winter-exclusive aid, such as winter appeals or Christmas grants.

Figure 5.7: Time-series plot for the number of visits



Again, we use SA and UM as benchmarks. This time, we consider 10 methods that include a different number of features:

- Non-feature: NF
- Seasonal feature (9-10): PLM-0, LM-0, NPC-0
- Local feature (5-10): PLM-1, LM-1, NPC-1
- National feature (1-10): PLM-2, LM-2, NPC-2.

To compare the performance of the methods, we obtain their 1-step ahead forecasts with rolling horizon, where *origin size* is 60 and the origin is shifted 44 times. For each forecasted value, we compute the overage/underage amount and the cost.

Table 5.3: Relative performance when $\beta = 0.95/0.90$ with 10 methods

$\beta = 0.95$		Measurements			
Methods	rMAE	rMPS	rRMSE	Relative 95%-DL	Relative SL
NF	91%	/	91%	7%	0%
PLM-0	/	/	/	/	/
PLM-1	/	/	91%	/	/
PLM-2	/	/	77%	/	10%
LM-0	74%	/	78%	2%	10%
LM-1	39%	92%	65%	45%	15%
LM-2	29%	85%	50%	92%	20%
NPC-0	83%	88%	72%	68%	15%
NPC-1	49%	63%	54%	85%	30%
NPC-2	29%	33%	50%	96%	30%

$\beta = 0.90$		Measurements			
Methods	rMAE	rMPS	rRMSE	Relative 90%-DL	Relative SL
NF	89%	99%	97%	13%	1%
PLM-0	/	/	/	1%	/
PLM-1	/	/	90%	1%	/
PLM-2	92%	/	78%	5%	10%
LM-0	70%	/	78%	4%	10%
LM-1	38%	91%	66%	46%	15%
LM-2	38%	85%	55%	93%	25%
NPC-0	83%	89%	72%	68%	15%
NPC-1	48%	63%	54%	88%	30%
NPC-2	27%	33%	50%	98%	30%

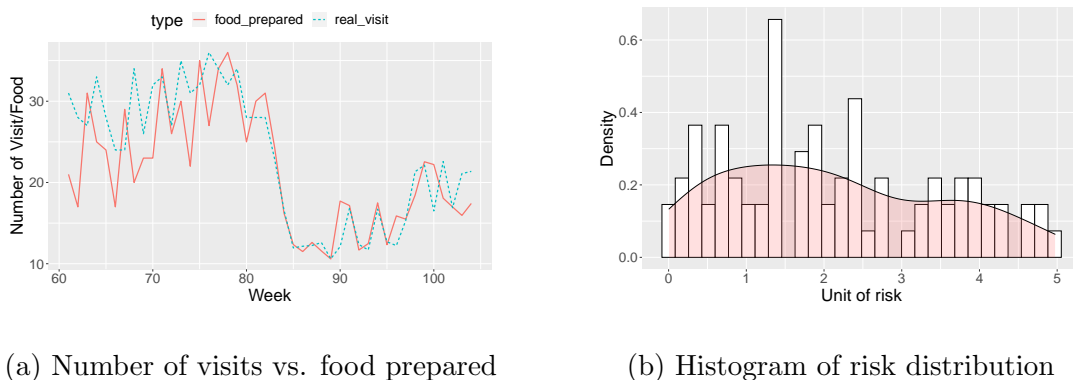
We summarise the results in Table 5.3, where rMAE denotes the Relative Mean Absolute Error for the visit estimation, rMPS denotes the Relative Mean Pinball Score and rRMSE denotes the Relative Root Mean Square Error (Davydenko and Fildes, 2013). We recall that low rMAE, rMPS and rRMSE are favourable, while

high Relative DL and Relative SL are favourable.

From Table 5.3, we can see that the NPC method with the national feature performs best, having the lowest error and highest relative DL. Moreover, under the same number of features and the same choice of β , the NPC outperforms other competing methods. This is expected. Since the true distributions of the time series and the error term are both unknown, and the cost function is nonlinear, the tail of the downside loss is hard to capture. All competing methods suffer from the underestimation of the downside loss from the tail. Moreover, NPC requires only a proportion of input data, significantly improving the computing speed. For instance, when $\beta = 0.90$ and national features are considered, the NPC method computes 5 times faster than the LM-2 method. Thus, NPC can not only help the decision maker to achieve lower downside risk but also works more efficiently overall.

Last but not least, let us now suppose the food bank indeed implements our NPC-2 approach with $\beta = 0.90$ for its food preparation decisions. We wish to gain some insights into the predictions made by the approach. Figure 5.8 provides the results acquired by NPC-2.

Figure 5.8: Results of NPC-2 for week 61-104



We derive the following insights from Figure 5.7 and 5.8, which could be useful for both the visitor and the food bank:

-
1. Due to the policy of minimising CVaR, the food bank may understock their food in winter.
 2. For food bank visitors, if possible, we would suggest visiting the food bank in summer, while try acquiring support from other sources in winter.
 3. The risk of using NPC-2 is close to ‘evenly’ distributed according to Figure 5.8b. The food bank will not face a significant loss, but it also cannot achieve minimum overall risk.
 4. We suggest the food bank systematically obtain and record data that may be associated with the demand since the proposed method performs better as more features are included.
 5. Keep track of the performance of the model over time, as the parameters of the model may evolve.

5.8 Concluding Remarks

In this chapter, we proposed an alternative non-parametric method (NPC) for CVaR minimisation. Unlike the existing methods, the NPC method requires only a small proportion of the data, significantly reducing the computational effort. Besides, it works directly with the data, not relying on any assumption on the demand distribution. Our experiments with both artificial and real-life data indicate that our proposed method is very robust with regard to different data structures, and it can handle easily both linear and nonlinear profits. On the other hand, one should be careful using NPC when the sample size is small, especially when the model is under-fitting, as it can be more vulnerable than other competing methods in this case. Luckily, this drawback is tolerable, as our motivation in proposing NPC was to reduce the computational effort with large instances.

There are several interesting topics for further research. First, as observed in our experiments, the performance of NPC suffers from model under-fitting. Therefore, it would be interesting to extend the current NPC model to deal with this drawback. For instance, one can try introducing an additional parameter that controls the data usage manually (to a value other than $2 \times (1 - \beta)$). Second, it would be desirable to develop a variable selection mechanism in NPC, so as to prevent the model from over-fitting automatically, e.g. by cross-validation or a step-wise technique based on an information criterion. Finally, although we focused our research on the NVP, the proposed method could be valuable in fields other than inventory control, such as Finance, Logistics or Manufacturing.

Chapter 6

Conclusion

In this last chapter, we summarise the contributions of the thesis and make some suggestions for future work.

6.1 Summary

Newsvendor problems (NVPs), also known as single-period stochastic inventory control problems, form an important topic in Operational Research and Management Science (OR/MS). Since the introduction of the classical NVP in the 1950s, researchers have considered several extensions of the problem, including variants with multiple product types, resource constraints, price settings, alternative objective functions, and so on. Moreover, several fields of study besides OR/MS have been proven to be relevant to the problem, such as Operations Management, Economics, Risk Analysis and Psychology. In this thesis, we extended the literature in four directions:

1. We have developed an integrated algorithm that determines the order quantities directly from past data and work for both NVP and NNVP;

-
2. We have formed a way to use sensitivity analysis to process information gained in the optimisation phase to inform marketing decisions;
 3. We have investigated the outcomes of judgemental adjustment under different conditions and situations, and provided a heuristic algorithm for “tuning” the adjustment parameters in practice;
 4. We have developed an alternative non-parametric approach to determine the order quantities that minimise the CVaR in the Newsvendor context.

In Chapter 2, we introduced an “Integrated Method for Estimation and Optimisation” or IMEO, which is extended from the method of Ban and Rudin (2019). Instead of minimising the expected opportunity cost, IMEO attempts to maximise the expected profit, and this turns out to be an important distinction in the nonlinear case. We showed that IMEO reduces to the method of Ban and Rudin (2019) in the linear case, when both methods turned out to be equivalent to quantile regression.

In our experiments on artificial data, we showed that IMEO performs at least as well as benchmark methods, in terms of both mean percentage profit loss and service level. We also found that IMEO is more robust to model misspecification than the existing approaches.

In Chapter 3, we presented a method for performing Sensitivity Analysis (SA) and Parametric Analysis (PA) for multi-item NVPs, using a combination of discrete approximation and linear programming. The approach is very general, being able to handle changes in costs, prices and resource availabilities. Under certain conditions, it can also handle changes in the demand distribution itself. In our view, this offers a useful tool to help coordinate marketing and inventory decisions in a retail environment, providing insights that are not immediately obvious otherwise. Extensive computational results, on both artificial and real examples, showed that the method

yields on average accurate results in reasonable computational times.

In Chapter 4, we turned our attention to some judgemental adjustment procedures that have traditionally been regarded as rather naïve. We showed that, surprisingly, the “naïve” adjustment procedures can lead to increased profits in some situations. In particular, they can be useful when (a) there is not enough data available to estimate parameters accurately, and (b) the demand model is misspecified. Interestingly, “demand chasing” appears to be more useful under condition (a), while “pull-to-centre” seems to be of more benefit under condition (b). In general, this is because in case (a) the order estimates suffer from some systematic bias due to short data length; whereas in case (b) the estimated variance is often higher than needed.

We then considered the possibility of applying naïve adjustments in an automated fashion. For this purpose, we proposed a simple heuristic for tuning the adjustment parameters. Using a real-life example, we showed that the tuned orders outperformed the pre-tuned ones in terms of the achieved profit, and also led to a service level closer to the target one.

Finally, in Chapter 5, we proposed an alternative non-parametric method (NPC) for CVaR minimisation. We considered both empirical and adaptive approaches. We gave rigorous proof that under suitable assumptions, the true risk is well estimated by the risk functions of NPC. Unlike the existing methods, the NPC requires only a small portion of data, significantly reducing the computational effort. Besides, it works directly with the data, not relying on any assumptions about demand distribution. Our experiments with both artificial and real-life data indicate that the proposed method is robust with regard to different data structures, and it can handle easily both linear and nonlinear profits.

This thesis combined the four directions of extensions into an integrated inventory

control framework. Chapter 2 and Chapter 5 considered the decision-making on the ordering phase. For decision-makers interested in expected profit maximisation solutions, Chapter 2 provided a novel data-driven algorithm. If decision-makers focused on the CVaR minimisation instead, the solution method in Chapter 5 becomes a handful. Chapter 4, on the other hand, investigated the usefulness of naïve adjustment, and proposed an alternative way to improve statistical decision using automated parameter selection. Finally, Chapter 3 used the information gained in the NVP optimisation model to inform marketing decisions. It allows decision-makers to balance marketing activities and inventory decisions to achieve maximum profit. Overall, this thesis built a framework for inventory control decision-making than the classic NVP, and created a handful of tools for real-life decision-makers.

6.2 Further Research

We believe that the work presented in this thesis has the potential to be extended in a number of ways. Here, we outline some suggestions for further work that stem from this thesis.

Considering the article presented in Chapter 2, first, it would be interesting to study the performance of IMEO with other demand models, such as ETS (Hyndman et al., 2020) and ARIMA (Box et al., 2015). Second, it would be desirable to develop a variable selection mechanism in IMEO. The conventional disjoint method allows one to do this in the first phase, for example by using cross-validation or a stepwise technique based on information criteria (Konishi and Kitagawa, 2008), while Ban and Rudin (2019) use regularisation for the selection and estimation. While these are good approaches, they require large samples and are computationally expensive. Developing a more efficient feature selection method can be considered one of the fruitful future research directions. Third, we focused our research on

NVPs, but IMEO could be potentially extended to multi-period and multi-item inventory problems. An especially challenging case for IMEO would be to deal with stock-outs, when customers may switch to different products (Shin et al., 2015).

Considering the article presented in Chapter 3, one potential weakness is that the cross-price elasticity data used in Subsection 3.6 was acquired from the literature instead of a real-life case. It could perhaps be beneficial to examine the proposed method again, with real-life elasticity data, also measuring the effect of the uncertainty coming from the estimation of the parameters on the profit. In terms of the methodology itself, it would be useful (though challenging) to develop methods for performing sensitivity analysis in the presence of product substitution.

Considering the article presented in Chapter 4, first, it might be beneficial to conduct behavioural experiments, in the lab and/or in the field, to confirm the simulation results. Second, one could attempt to characterise other scenarios under which naïve adjustments tend to be beneficial. Given that the demand models frequently miss important information in real life, we would expect the number of those scenarios to be large. Third, one could examine the effects of other forms of adjustment, such as scenario forecasting.

Finally, we consider the article presented in Chapter 5. First, as observed in our experiments, the performance of NPC sometimes suffers from model under-fitting. Therefore, it would be interesting to extend the current NPC model to deal with this. For instance, one could try introducing an additional parameter that controls the data usage manually (to a value other than $2 \times (1 - \beta)$). Second, it would be desirable to develop a variable selection mechanism in NPC, to prevent the model from over-fitting automatically, e.g. by cross-validation or a step-wise technique based on information criteria. Finally, though we focused our research on the NVP, the proposed method could be valuable in fields other than inventory control, e.g. Finance, Logistics.

On a broader picture, we think it would be also interesting to test other variants of NVP model with the ideas mentioned in the thesis. For instance, one could extend the data-driven methods mentioned in Chapter 2 and Chapter 5 to fit in a MNVP model with substitution. Moreover, one could attempt to develop an integrated price/order decision strategy that takes pricing and inventory decisions into account at the same time, by adding additional decision variables, for example. Last but not least, it would be beneficial to look into NVPs with multi-objective strategy and Pareto analysis.

Appendices

A Proof of Equivalence to Quantile Regression

Proof. In quantile regression, the objective function is defined as (Koenker and Hallock, 2001):

$$\min \sum_{t=1}^s \rho_{\tau}(y_t - q(\mathbf{x}_t)),$$

where $\rho_{\tau}(u) = u(\tau - \mathbb{I}_{(u < 0)})$, and \mathbb{I} is an indicator function. Thus, simply by setting $\tau = c_u/(c_o + c_u)$, we have:

$$\begin{aligned} & \min \sum_{t=1}^s (c_u[y_t - q(\mathbf{x}_t)]^+ + c_o[q(\mathbf{x}_t) - y_t]^+) \\ &= \min(c_o + c_u) \sum_{t=1}^s \left(\frac{c_u}{c_o + c_u} [y_t - q(\mathbf{x}_t)]^+ + \frac{c_o}{c_o + c_u} [q(\mathbf{x}_t) - y_t]^+ \right) \\ &= \min \sum_{t=1}^s (\tau[y_t - q(\mathbf{x}_t)]^+ + (1 - \tau)[q(\mathbf{x}_t) - y_t]^+) \\ &= \min \sum_{t=1}^s \rho_{\tau}(y_t - q(\mathbf{x}_t)). \end{aligned}$$

□

B Proof of Maximising-Minimising Transformation

Proof. We have:

$$\min[a, b] = a - [a - b]^+,$$

and

$$a - b = [a - b]^+ - [b - a]^+.$$

We can transform:

$$\begin{aligned} \pi(q(\mathbf{x}_t), y_t) &= p \min[q(\mathbf{x}_t), y_t] - vq(\mathbf{x}_t) - c_h[q(\mathbf{x}_t) - y_t]^+ - c_s[y_t - q(\mathbf{x}_t)]^+ \\ &= p\{q(\mathbf{x}_t) - [q(\mathbf{x}_t) - y_t]^+\} - vq(\mathbf{x}_t) - c_h[q(\mathbf{x}_t) - y_t]^+ - c_s[y_t - q(\mathbf{x}_t)]^+ \\ &= (p - v)q(\mathbf{x}_t) - (c_h + p)[q(\mathbf{x}_t) - y_t]^+ - c_s[y_t - q(\mathbf{x}_t)]^+. \end{aligned}$$

Therefore, we have (since y_t is fixed):

$$\begin{aligned} \max \sum_{t=1}^s \pi(q(\mathbf{x}_t), y_t) &= \max \sum_{t=1}^s \{(p - v)q(\mathbf{x}_t) - (c_h + p)[q(\mathbf{x}_t) - y_t]^+ - c_s[y_t - q(\mathbf{x}_t)]^+\} \\ &= \max \sum_{t=1}^s \{(p - v)[q(\mathbf{x}_t) - y_t] - (c_h + p)[q(\mathbf{x}_t) - y_t]^+ - c_s[y_t - q(\mathbf{x}_t)]^+\} \\ &= \max \sum_{t=1}^s \{(p - v)[q(\mathbf{x}_t) - y_t]^+ - (p - v)[y_t - q(\mathbf{x}_t)]^+ \\ &\quad - (c_h + p)[q(\mathbf{x}_t) - y_t]^+ - c_s[y_t - q(\mathbf{x}_t)]^+\} \\ &= \min \sum_{t=1}^s \{(v + c_h)[q(\mathbf{x}_t) - y_t]^+ + (p - v + c_s)[y_t - q(\mathbf{x}_t)]^+\} \\ &= \min \sum_{t=1}^s \{c_o[q(\mathbf{x}_t) - y_t]^+ + c_u[y_t - q(\mathbf{x}_t)]^+\}. \end{aligned}$$

□

C Small Artificial Instance

Suppose that $n = 2$ and $m = 3$. The products are called a and b , and the resources are called A , B and C . The selling prices and the order, holding and shortage costs are $(8, 3, 2, 1)$ and $(6, 3, 4, 3)$ for products a and b , respectively. The mean demands and standard deviations are $(210, 5)$ and $(210, 6)$, respectively. One unit of product a requires 4, 7 and 8 units of resource A , B and C , respectively. One unit of product b requires 6, 5 and 8 units of resources A , B and C . The resource availabilities are 2200, 2500 and 3500, respectively. We suppose that we have generated twelve scenarios for each product, by RS, as shown in Table 1.

Table 1: Twelve possible demand realisations for two products

Products	Demand Realisations											
	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}	s_{11}	s_{12}
product a	200	220	180	190	190	210	240	250	200	190	210	240
product b	250	230	200	180	210	210	170	150	180	220	260	260

The LP for this example has 26 variables and 27 constraints. It can be solved in less than one second. Table 2 shows some of the resulting estimates of various partial derivatives. It can be shown that these estimates are all accurate to within $\pm 5\%$.

Table 2: Influence of changes in costs, prices, resources and demand

Products	Margins					
	$\tilde{P}/\partial r_j$	$\tilde{P}/\partial c_j$	$\tilde{P}/\partial v_j$	$\tilde{P}/\partial g_j$	$\tilde{P}/\partial \mu_j$	$\tilde{P}/\partial \sigma_j$
product a	199.40	-207.14	-7.74	-10.60	4.50	-4.03
product b	195.83	-210.00	-14.17	-14.17	2.64	-5.26

D Constraints for Real-Life Example

Table 3: Constraints for real-life case

bread	egg	fish	fruit	juice	vegetable	meat	milk	dairy	direction	RHS
0	0	0	1	0	1	0	0	0	\leq	1200
1	0	1	1	0	0.1	1	1	0	\leq	550
0	0	0	0	0	0	0	1	1	\geq	30
0	0	0	0	1	0	0	1	0	\leq	300
0	1	0	0	0	0	0	0	1	\leq	60

E Elasticities for Real-Life Example

Table 4: Cross-price elasticities for 9 key products

	bread	egg	fish	fruit	juice	vegetable	meat	milk	dairy
bread	0.06	-0.03	-0.05	-0.01	0.00	-0.05	0.00	0.00	-0.01
egg	-0.23	-0.22	-0.04	-0.06	0.07	-0.03	0.04	-0.09	-0.08
fish	-0.15	-0.05	-0.53	0.00	-0.06	0.01	0.02	-0.16	-0.04
fruit	-0.09	-0.04	0.00	-0.60	-0.04	-0.03	0.00	0.30	-0.08
juice	-0.06	0.00	-0.09	-0.06	-0.85	0.00	0.00	0.14	-0.05
vegetable	0.00	0.00	0.00	0.10	0.00	-0.54	0.01	0.20	-0.03
meat	0.00	0.02	0.08	-0.04	0.00	-0.02	-0.48	0.17	-0.23
milk	0.00	-0.12	-0.01	0.07	0.08	0.08	0.02	-0.96	0.00
dairy	-0.04	-0.01	-0.01	-0.01	0.00	-0.01	-0.07	0.00	-0.55

F Heatmaps for model misspecification

Figure 1: RPI heat map for the under-parametrised case ($\tau = 0.7, t = 200$)






Figure 2: RPI heatmap for the over-parametrised case ($\tau = 0.7, t = 200$)



G Expectation maximisation vs. CVaR minimisation

Figure 3: Difference between the expectation maximisation solution and the CVaR minimisation solution



In Figure 3, we mark three order quantities. They fulfil the objectives of expectation maximisation, CVaR minimisation and risk averse profit maximisation ($0.7 \times \text{Mean} - 0.3 \times \text{CVaR}$), respectively. We can see the order quantity that minimise CVaR is lower than the order quantity that maximise expectation. However, this is parameter-dependant, as the CVaR minimisation order quantity is a weight average of critical quantiles. When the overage cost is significantly larger than the underage cost, the CVaR minimisation quantity will be, with no doubt, larger than the expectation maximisation quantity.


H Baseline experiment parameters

Table 5: Baseline experiment parameters

b_0	b_1	b_2	b_3	b_4	$\theta_{1,1}^1$	$\theta_{12,1}^1$	$\theta_{1,1}^2$	$\phi_{1,1}^2$	$\theta_{12,1}^2$
500	0.642	0.354	0.407	0.521	0.3	0.5	0.2	0.5	0.1
$\theta_{1,1}^3$	$\phi_{1,1}^3$	$\phi_{12,1}^3$	$\phi_{1,1}^4$	$\phi_{1,2}^4$	$\theta_{12,1}^4$	$\theta_{12,2}^4$	rnorm	rlaplace	rt
0.3	0.2	0.1	0.1	0.2	0.1	0.1	$\mu = 0$	$\mu = 0$	$\mu = 0$
							$\sigma = 100$	$b = 71$	$\sigma = 100$
									$\nu = 5$

We mark that the t-distribution is believed to have heavy-tail, and the normal distribution is believed to have light-tail. We use a mix of t-distribution, normal distribution and Laplace distribution with random weights to simulate real circumstance where we have no information about the shape of error distribution in prior.

Figure 4: Baseline experiment features and demand series



I Histogram example

Figure 5: Profit histogram on the non-parametric method and benchmark method



In this example, the origin size = 200, iteration number = 100, and $\beta = 0.95$. Therefore, in the histogram, 100 profit realisations are considered for each methods, and the average profit/loss for the worst 5 cases are marked by dashed lines. We also marked the corresponding performance of the UM method, where all features and all observations are considered, by a red line dashed line. Here, we have $DL_{SA} = -319.76$, $DL_{NPC} = -2355.24$, $DL_{UM} = -2508.41$, $SL_{SA} = 86\%$, $SL_{NPC} = 90.5\%$ and $SL_{UM} = 88.5\%$.

J Baseline experiment full results

Table 6: Relative β -DL/Relative SL for all choices of parameters under linear profit function when $\beta = 0.95$ (or 0.9)

Relative β -DL/ SL		Origin size						
Iteration	Method	50	100	150	200	250	300	
50	NPC	60%/60%	60%/50%	76%/80%	99%/33%	74%/20%	68%/-	
		(55%/0%)	(78%/100%)	(95%/25%)	(89%/33%)	(58%/33%)	(83%/0%)	
	NF	-39%/0%	-1%/0%	11%/40%	-13%/67%	8%/-100%	-14%/-	
		(-55%/0%)	(-5%/25%)	(0%/0%)	(3%/-67%)	(-48%/-100%)	(-3%/-33%)	
	SQR	68%/20%	88%/25%	98%/60%	99%/99%	89%/0%	74%/-	
		(84%/0%)	(68%/75%)	(99%/75%)	(99%/67%)	(90%/67%)	(95%/67%)	
	PLM	-8%/40%	7%/25%	7%/60%	99%/33%	57%/-300%	32%/-	
		(-88%/-60%)	(72%/75%)	(89%/25%)	(77%/100%)	(37%/67%)	(81%/100%)	
	100	NPC	26%/17%	54%/0%	87%/6%	99%/15%	76%/5%	86%/67%
			(60%/75%)	(97%/63%)	(93%/14%)	(81%/20%)	(75%/0%)	(91%/33%)
		NF	-73%/0%	0%/200%	5%/67%	-6%/0%	-5%/0%	20%/6%
			(-19%/50%)	(-2%/13%)	(-1%/-29%)	(-10%/-100%)	(-21%/-67%)	(5%/0%)
SQR		79%/-17%	96%/-50%	99%/89%	98%/100%	69%/-200%	85%/67%	
		(79%/50%)	(77%/75%)	(94%/71%)	(99%/60%)	(91%/100%)	(99%/50%)	
PLM		-71%/-33%	49%/-60%	80%/56%	97%/-100%	44%/-20%	59%/89%	
		(4%/-75%)	(81%/75%)	(79%/29%)	(67%/60%)	(66%/50%)	(88%/67%)	
150		NPC	65%/43%	68%/50%	90%/14%	99%/25%	80%/56%	90%/56%
			(77%/20%)	(93%/30%)	(90%/8%)	(83%/14%)	(83%/0%)	(93%/44%)
		NF	-25%/29%	-11%/0%	6%/86%	-1%/0%	19%/22%	12%/67%
			(-30%/30%)	(-2%/-20%)	(-7%/-25%)	(-9%/-85%)	(-12%/-10%)	(19%/0%)
	SQR	87%/-71%	92%/-50%	98%/71%	81%/50%	88%/44%	82%/44%	
		(74%/60%)	(83%/70%)	(92%/83%)	(97%/71%)	(97%/80%)	(99%/78%)	
	PLM	-45%/-44%	57%/-50%	79%/43%	86%/-100%	68%/56%	64%/67%	
		(-21%/-20%)	(74%/100%)	(75%/75%)	(69%/14%)	(75%/60%)	(83%/100%)	
	200	NPC	76%/72%	68%/0%	77%/0%	73%/44%	71%/73%	90%/20%
			(71%/40%)	(91%/73%)	(90%/13%)	(87%/36%)	(84%/15%)	(87%/5%)
		NF	-24%/29%	-8%/-80%	0%/80%	-8%/11%	-5%/27%	-14%/100%
			(-21%/20%)	(-2%/-13%)	(-8%/-7%)	(-7%/-18%)	(4%/-8%)	(14%/16%)
SQR		70%/-43%	77%/-50%	75%/40%	78%/78%	76%/55%	72%/0%	
		(71%/73%)	(82%/67%)	(92%/93%)	(99%/64%)	(97%/92%)	(99%/95%)	
PLM		42%/-43%	38%/-40%	44%/40%	39%/44%	51%/73%	47%/80%	
		(-14%/-7%)	(60%/93%)	(76%/87%)	(71%/45%)	(68%/77%)	(88%/95%)	

Table 7: Relative β -DL/Relative SL for all choices of parameters under nonlinear profit function when $\beta = 0.95$ (or 0.9)

Relative β -DL/ SL		Origin size					
Iteration	Method	50	100	150	200	250	300
50	NPC	85%/20%	99%/50%	56%/-	66%/88%	99%/25%	99%/100%
		(87%/100%)	(78%/50%)	(63%/0%)	(75%/50%)	(95%/ -)	(85%/90%)
	NF	-17%/-100%	-84%/0%	8%/-	6%/13%	13%/0%	-14%/11%
		(-39%/-100%)	(-10%/-50%)	(11%/50%)	(4%/13%)	(4%/ -)	(-9%/20%)
	SQR	1%/-40%	5%/0%	6%/-	23%/0%	43%/-50%	37%/-11
(-8%/-40%)		(-16%/-50%)	(-3%/-15%)	(12%/13%)	(18%/ -)	(18%/-20%)	
PLM	-10%/-100%	-2%/-50%	39%/-	36%/88%	63%/100%	57%/88	
	(-32%/100%)	(-19%/0%)	(52%/100%)	(20%/88%)	(52%/ -)	(72%/100%)	
100	NPC	99%/60%	99%/100%	87%/40%	80%/58%	84%/27%	99%/57%
		(99%/67%)	(97%/100%)	(83%/80%)	(92%/25%)	(88%/60%)	(91%/14%)
	NF	-43%/-100%	-19%/0%	23%/0%	9%/8%	-3%/0%	-12%/-14%
		(-51%/33%)	(-31%/0%)	(6%/0%)	(-7%/-25%)	(0%/-10%)	(-3%/42%)
	SQR	-32%/-50%	-15%/-25%	0%/14%	43%/-8%	46%/-36%	19%/-85%
(-15%/-13%)		(-19%/-20%)	(-3%/-14%)	(22%/-37%)	(33%/-50%)	(32%/-14%)	
PLM	-21%/-19%	-23%/-45%	39%/80%	53%/91%	74%/72%	62%/14%	
	(4%/33%)	(81%/67%)	(79%/100%)	(67%/63%)	(66%/70%)	(88%/85%)	
150	NPC	85%/30%	99%/14%	86%/22%	99%/17%	99%/100%	88%/13%
		(74%/90%)	(93%/50%)	(88%/100%)	(91%/23%)	(93%/83%)	(67%/50%)
	NF	-16%/-100%	-13%/40%	14%/11%	-6%/11%	-3%/12%	-11%/-33%
		(-37%/100%)	(-21%/100%)	(2%/-100%)	(-5%/-5%)	(0%/-33%)	(-6%/50%)
	SQR	-13%/-90%	-10%/-18%	27%/-88%	33%/-23%	24%/12%	10%/18%
(22%/-11%)		(6%/-46%)	(2%/-30%)	(6%/-29%)	(3%/-21%)	(5%/-16%)	
PLM	-13%/-29%	-10%/-16%	45%/77%	60%/88%	59%/25%	67%/-67%	
	(3%/-11%)	(52%/66%)	(59%/75%)	(49%/76%)	(43%/50%)	(51%/88%)	
200	NPC	76%/10%	56%/18%	78%/69%	96%/67%	75%/21%	93%/85%
		(60%/87%)	(95%/75%)	(77%/80%)	(88%/83%)	(97%/83%)	(97%/85%)
	NF	-24%/50%	-6%/-33%	1%/7%	-9%/-16%	-4%/14%	-12%/-15%
		(-34%/25%)	(-20%/-50%)	(-2%/-8%)	(-6%/15%)	(-4%/-16%)	(-6%/33%)
	SQR	-8%/-55%	12%/-23%	25%/-76%	16%/-91%	17%/-21%	4%/-11%
(-16%/-35%)		(-38%/11%)	(9%/-93%)	(18%/-92%)	(30%/-30%)	(-19%/-88%)	
PLM	-12%/-16%	-5%/-100%	46%/61%	42%/91%	52%/-42%	-2%/-15%	
	(-9%/-20%)	(55%/-10%)	(57%/75%)	(30%/76%)	(37%/50%)	(9%/66%)	

K Features for food preparation problem

Table 8: Relevant features to food preparation problem within food bank

No.	Feature
1	UK inflation data (monthly)
2	UK unemployment rate (monthly)
3	UK economics index (weekly)
4	FTSE 100 close price (weekly)
5	Durham birth registered (weekly)
6	Durham death registered (weekly)
7	Durham Covid-19 cases (weekly)
8	Durham crime index (weekly)
9	UK Bank holidays dummies
10	Seasonality dummies

Bibliography

- Abdel-Aal, M.A. and S. Selim (2017). ‘Risk-averse multi-product selective news-vendor problem with different market entry scenarios under CVaR criterion’. In: *Computers & Industrial Engineering* 103, pp. 250–261.
- Abdel-Malek, L.L. and N. Areeratchakul (2007). ‘A quadratic programming approach to the multi-product newsvendor problem with side constraints’. In: *European Journal of Operational Research* 176, pp. 1607–1619.
- Abdel-Malek, L.L. and R. Montanari (2005a). ‘An analysis of the multi-product newsboy problem with a budget constraint’. In: *International Journal of Production Economics* 97, pp. 296–307.
- (2005b). ‘On the multi-product newsboy problem with two constraints’. In: *Computers & Operations Research* 32, pp. 2095–2116.
- Abdel-Malek, L.L., R. Montanari and L.C. Morales (2004). ‘Exact, approximate, and generic iterative models for the multi-product newsboy problem with budget constraint’. In: *International Journal of Production Economics* 91, pp. 189–198.
- Al Thobaity, A. and F. Alshammari (2020). ‘Nurses on the frontline against the COVID-19 pandemic: an integrative review’. In: *Dubai Medical Journal* 3, pp. 87–92.
- Alfares, H.K. and H.H. Elmorra (2005). ‘The distribution-free newsboy problem: Extensions to the shortage penalty case’. In: *International Journal of Production Economics* 93, pp. 465–477.

-
- Anvari, M (1987). ‘Optimality criteria and risk in inventory models: The case of the newsboy problem’. In: *Journal of the Operational Research Society* 38, pp. 625–632.
- Arrow, K.J., T. Harris and J. Marschak (1951). ‘Optimal inventory policy’. In: *Econometrica* 19, pp. 250–272.
- Ban, G.-Y. and C. Rudin (2019). ‘The big data newsvendor: practical insights from machine learning’. In: *Operations Research* 67, pp. 90–108.
- Bassok, Y., R. Anupindi and R. Akella (1999). ‘Single-period multiproduct inventory models with substitution’. In: *Operations Research* 47, pp. 632–642.
- Beale, E.M.L. (1955). ‘On minimizing a convex function subject to linear inequalities’. In: *Journal of the Royal Statistical Society* 17, pp. 173–184.
- Ben-Daya, M. and A. Raouf (1993). ‘On the constrained multi-item single-period inventory problem’. In: *International Journal of Operations & Production Management* 13, pp. 104–112.
- Benzion, U. et al. (2008). ‘Decision-making and the newsvendor problem: an experimental study’. In: *Journal of the Operational Research Society* 59, pp. 1281–1287.
- Bertsimas, D. and N. Kallus (2020). ‘From predictive to prescriptive analytics’. In: *Management Science* 66, pp. 1025–1044.
- Bertsimas, D. and A. Thiele (2005). *A data-driven approach to newsvendor problems*. Tech. rep. Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology.
- Beutel, A.-L. and S. Minner (2012). ‘Safety stock planning under causal demand forecasting’. In: *International Journal of Production Economics* 140, pp. 637–645.
- Birge, J.R. and F. Louveaux (2011). *Introduction to Stochastic Programming*. 2nd ed. New York: Springer.

-
- Bolton, G.E. and E. Katok (2008). ‘Learning by doing in the newsvendor problem: a laboratory investigation of the role of experience and feedback’. In: *Manufacturing & Service Operations Management* 10, pp. 519–538.
- Bostian, A.A., C.A. Holt and A.M. Smith (2008). ‘Newsvendor “pull-to-center” effect: adaptive learning in a laboratory experiment’. In: *Manufacturing & Service Operations Management* 10, pp. 590–608.
- Box, G.E.P. et al. (2015). *Time Series Analysis: Forecasting and Control*. Hoboken, NJ: Wiley.
- Brown, R.G. (1956). *Exponential Smoothing for Predicting Demand*. Cambridge, MA: Little.
- Bruzda, J. (2016). ‘Quantile forecasting in operational planning and inventory management—an initial empirical verification’. In: *Dynamic Econometric Models* 16, pp. 5–20.
- Chen, S.-C., Y.-H. Lai and S.-L. Tsay (2020). ‘Nursing perspectives on the impacts of COVID-19’. In: *Journal of Nursing Research* 28, pp. 1682–3141.
- Chen, Y., M. Xu and Z.G. Zhang (2009). ‘A risk-averse newsvendor model under the CVaR criterion’. In: *Operations Research* 57, pp. 1040–1044.
- Cheng, L., Z. Wan and G. Wang (2009). ‘Bilevel newsvendor models considering retailer with CVaR objective’. In: *Computers & Industrial Engineering* 57, pp. 310–318.
- Chernoff, H. (1952). ‘A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations’. In: *Annals of Mathematical Statistics* 23, pp. 493–507.
- Choi, T.-M. (2012). *Handbook of Newsvendor Problems*. New York: Springer.
- Chun, S.Y., A. Shapiro and S. Uryasev (2012). ‘Conditional value-at-risk and average value-at-risk: Estimation and asymptotics’. In: *Operations Research* 60, pp. 739–756.

-
- Coronini-Cronberg, S., E.J. Maile and A. Majeed (2020). ‘Health inequalities: the hidden cost of COVID-19 in NHS hospital trusts?’ In: *Journal of the Royal Society of Medicine* 113, pp. 179–184.
- Cui, Y. et al. (2013). ‘Chinese perspective on newsvendor bias: an exploratory note’. In: *Journal of Operations Management* 31, pp. 93–97.
- Dantzig, G.B. (1955). ‘Linear programming under uncertainty’. In: *Management Science* 1, pp. 197–206.
- (1963). *Linear Programming and Extensions*. Princeton, NJ: Princeton University Press.
- Darwish, M.A., M. Alkhedher and A. Alenezi (2019). ‘Reducing the effects of demand uncertainty in single-newsvendor multi-retailer supply chains’. In: *International Journal of Production Research* 57, pp. 1082–1102.
- Davis, L.B. et al. (2014). ‘Scheduling food bank collections and deliveries to ensure food safety and improve access’. In: *Socio-Economic Planning Sciences* 48, pp. 175–188.
- Davydenko, A. and R. Fildes (2013). ‘Measuring forecasting accuracy: The case of judgmental adjustments to SKU-level demand forecasts’. In: *International Journal of Forecasting* 29, pp. 510–522.
- De Haan, L. and A. Ferreira (2006). *Extreme Value Theory: An Introduction*. New York: Springer.
- De Vericourt, F. et al. (2013). ‘Sex, risk and the newsvendor’. In: *Journal of Operations Management* 31, pp. 86–92.
- Edgeworth, F.Y. (1888). ‘The Mathematical Theory of Banking’. In: *Journal of the Royal Statistical Society* 51, pp. 113–127.
- Eeckhoudt, L., C. Gollier and H. Schlesinger (1995). ‘The risk-averse (and prudent) newsboy’. In: *Management Science* 41, pp. 786–794.
- Efron, B. and R. Tibshirani (1994). *An Introduction to the Bootstrap*. Boca Raton, FL: Chapman & Hall.

-
- El Agizy, M. (1967). ‘Two-stage programming under uncertainty with discrete distribution function’. In: *Operations Research* 15, pp. 55–70.
- Elmachtoub, A.N. and P. Grigas (2017). ‘Smart “predict, then optimize”’. In: *Management Science* 68, pp. 9–26.
- Erlebacher, S.J. (2000). ‘Optimal and heuristic solutions for the multi-item newsvendor problem with a single capacity constraint’. In: *Production and Operations Management* 9, pp. 303–318.
- Farrar, D.E. and R.R. Glauber (1967). ‘Multicollinearity in regression analysis: the problem revisited’. In: *Review of Economic and Statistics* 49, pp. 92–107.
- Feng, T., L.R. Keller and X. Zheng (2011). ‘Decision making in the newsvendor problem: A cross-national laboratory study’. In: *Omega* 39, pp. 41–50.
- Fernandez, R. et al. (2020). ‘Implications for COVID-19: a systematic review of nurses’ experiences of working in acute care hospital settings during a respiratory pandemic’. In: *International Journal of Nursing Studies* 111.
- Fourer, R. (2021). ‘Optimisation software survey’. In: *OR/MS Today* 48.
- Frank, R.H. and E. Cartwright (2008). *Microeconomics and Behavior*. Boston, MA: McGraw-Hill.
- Gal, T. (1995). *Postoptimal Analyses, Parametric Programming, and Related Topics*. 2nd ed. Berlin: De Gruyter.
- Gal, T. and H.J. Greenberg, eds. (1997). *Advances in Sensitivity Analysis and Parametric Programming*. Dordrecht: Kluwer.
- Gale, D., H.W. Kuhn and A.W. Tucker (1951). ‘Linear programming and the theory of games’. In: *Activity Analysis of Production and Allocation*. Ed. by T.C. Koopmans. New York: Wiley, pp. 317–335.
- Gençay, R., F. Selçuk and A. Ulugülyağci (2003). ‘High volatility, thick tails and extreme value theory in value-at-risk estimation’. In: *Insurance: Mathematics and Economics* 33, pp. 337–356.

-
- Gerchak, Y. and M. Parlar (1987). ‘A single period inventory problem with partially controllable demand’. In: *Computers & Operations Research* 14, pp. 1–9.
- Gilliland, M. (2010). *The Business Forecasting Deal*. Hoboken, NJ: Wiley.
- Gondzio, J. (2012). ‘Interior point methods 25 years later’. In: *European Journal of Operational Research* 218, pp. 587–601.
- Goodwin, P. and G. Wright (2014). *Decision Analysis for Management Judgement*. 5th ed. Chichester, UK: Wiley.
- Gotoh, J.-Y. and Y. Takano (2007). ‘Newsvendor solutions via conditional value-at-risk minimization’. In: *European Journal of Operational Research* 179, pp. 80–96.
- Güler, M.G. (2019). ‘Advertising and forecasting investments of a newsvendor’. In: *4OR* 24, pp. 45–73.
- Hadley, G. and T.M. Whitin (1963). *Analysis of Inventory Systems*. Englewood, NJ: Prentice-Hall.
- Halman, N., J.B. Orlin and D. Simchi-Levi (2012). ‘Approximating the nonlinear newsvendor and single-item stochastic lot-sizing problems when data is given by an oracle’. In: *Operations Research* 60, pp. 429–446.
- Harsha, P., R. Natarajan and D. Subramanian (2015). *A data-driven, distribution-free, multivariate approach to the price-setting newsvendor problem*. Tech. rep. Working paper, IBM TJ Watson Research Center, Yorktown Heights, NY.
- He, B. et al. (2012). ‘The timing of staffing decisions in hospital operating rooms: Incorporating workload heterogeneity into the newsvendor problem’. In: *Manufacturing & Service Operations Management* 14, pp. 99–114.
- Heinze, G., C. Wallisch and D. Dunkler (2018). ‘Variable selection - A review and recommendations for the practicing statistician’. In: *Biometrical Journal* 60, pp. 431–449.

-
- Henneberry, S.R., K. Piewthongngam and H. Qiang (1999). ‘Consumer food safety concerns and fresh produce consumption’. In: *Journal of Agricultural and Resource Economics* 24, pp. 1–16.
- Ho, T.-H., N. Lim and T.H. Cui (2010). ‘Reference dependence in multilocation newsvendor models: a structural analysis’. In: *Management Science* 56, pp. 1891–1910.
- Hochbaum, D.S. (1995). ‘A nonlinear knapsack problem’. In: *Operations Research Letters* 17, pp. 103–110.
- Hodges, S.D. and P.G. Moore (1970). ‘The product-mix problem under stochastic seasonal demand’. In: *Management Science* 17, B107–B144.
- Huber, J. et al. (2019). ‘A data-driven newsvendor problem: From data to decision’. In: *European Journal of Operational Research* 278, pp. 904–915.
- Hyndman, R. et al. (2020). *forecast: Forecasting functions for time series and linear models*. R package version 8.13.
- Jammerneegg, W. and P. Kischka (2007). ‘Risk-averse and risk-taking newsvendors: a conditional expected value approach’. In: *Review of Managerial Science* 1, pp. 93–110.
- Kahn, K.B. (2014). *New Product Forecasting: An Applied Approach*. New York: Routledge.
- Kahneman, D. and A. Tversky (1979). ‘Prospect theory: an analysis of decision under risk’. In: *Econometrica* 47, pp. 263–291.
- Kall, P. and S.W. Wallace (1994). *Stochastic Programming*. Chichester, UK: Wiley.
- Karlin, S. and C.R. Carr (1962). ‘Prices and optimal inventory policy’. In: *Studies in Applied Probability and Management Science*. Ed. by K.J. Arrow, S. Karlin and H. Scarf. Stanford, CA: Stanford University Press, pp. 159–172.
- Karmarkar, U.S. (1994). ‘A robust forecasting technique for inventory and leadtime management’. In: *Journal of Operations Management* 12, pp. 45–54.

-
- Kashefi, M.A. (2016). ‘Effect of salvage market on strategic technology choice and capacity investment decision of firm under demand uncertainty’. In: *Facility Design eJournal* 17, pp. 140–155.
- Khouja, M. (1995). ‘The newsboy problem under progressive multiple discounts’. In: *European Journal of Operational Research* 84, pp. 458–466.
- (1999). ‘The single-period (news-vendor) problem: literature review and suggestions for future research’. In: *Omega* 27, pp. 537–553.
- Khouja, M. and S.S. Robbins (2003). ‘Linking advertising and quantity decisions in the single-period inventory model’. In: *International Journal of Production Economics* 86, pp. 93–105.
- Kim, G., K. Wu and E. Huang (2015). ‘Optimal inventory control in a multi-period newsvendor problem with non-stationary demand’. In: *Advanced Engineering Informatics* 29, pp. 139–145.
- Kirshner, S.N. and B.B. Moritz (2020). ‘Comment on “Newsvendor demand chasing revisited”’. In: *Management Science* 66, pp. 6062–6064.
- Kocherginsky, M., X. He and Y. Mu (2005). ‘Practical confidence intervals for regression quantiles’. In: *Journal of Computational and Graphical Statistics* 14, pp. 41–55.
- Koenker, R. (2005). *Quantile Regression*. Cambridge, UK: Cambridge University Press.
- Koenker, R. and K.F. Hallock (2001). ‘Quantile regression’. In: *Journal of Economic Perspectives* 15, pp. 143–156.
- Koenker, R. and J.A.F. Machado (1999). ‘Goodness of fit and related inference processes for quantile regression’. In: *Journal of the American Statistical Association* 94, pp. 1296–1310.
- Konishi, S. and G. Kitagawa (2008). *Information Criteria and Statistical Modeling*. New York: Springer.

-
- Korpela, J. and M. Tuominen (1996). ‘Inventory forecasting with a multiple criteria decision tool’. In: *International Journal of Production Economics* 45, pp. 159–168.
- Kyparisis, G.J. and C. Koulamas (2018). ‘The price-setting newsvendor with non-linear salvage revenue and shortage cost’. In: *Operations Research Letters* 46, pp. 64–68.
- Lariviere, M.A. and E.L. Porteus (2001). ‘Selling to the newsvendor: An analysis of price-only contracts’. In: *Manufacturing & Service Operations Management* 3, pp. 293–305.
- Lau, H.-S. and A. Lau (1995). ‘The multi-product multi-constraint newsboy problem: Applications, formulation and solution’. In: *Journal of Operations Management* 13, pp. 153–162.
- (1996). ‘The newsstand problem: A capacitated multi-product single period inventory problem’. In: *Operations Research* 94, pp. 29–42.
- (1997). ‘Some results on implementing a multi-item multi-constraint single-period inventory model’. In: *International Journal of Production Economics* 48, pp. 121–128.
- Lau, N. and J.N. Bearden (2013). ‘Newsvendor demand chasing revisited’. In: *Management Science* 59, pp. 1245–1249.
- Lau, N., S. Hasija and J.N. Bearden (2014). ‘Newsvendor pull-to-center reconsidered’. In: *Decision Support Systems* 58, pp. 68–73.
- Levi, R., G. Perakis and J. Uichanco (2015). ‘The data-driven newsvendor problem: new bounds and insights’. In: *Operations Research* 63, pp. 1294–1306.
- Liberopoulos, G., I. Tsikis and S. Delikouras (2010). ‘Backorder penalty cost coefficient “b”: what could it be?’ In: *International Journal of Production Economics* 123, pp. 166–178.

-
- Lim, J.S. and M. O'Connor (1995). 'Judgemental adjustment of initial forecasts: Its effectiveness and biases'. In: *Journal of Behavioral Decision Making* 8, pp. 149–168.
- Liu, C. (2020). *Forecasting-and-Newsvendor*. URL: <https://github.com/Joshua960411/Forecasting-and-Newsvendor>.
- (2022). *Github repository – NPC*. URL: <https://github.com/JoshuaLiu111/NPC>.
- Liu, C., A.N. Letchford and I. Svetunkov (2022). 'Newsvendor problems: an integrated method for estimation and optimisation'. In: *European Journal of Operational Research* 300, pp. 590–601.
- Liu, D.C. and J. Nocedal (1989). 'On the limited memory BFGS method for large scale optimization'. In: *Mathematical Programming* 45, pp. 503–528.
- Liu, K.E. and W.S. Chern (2003). *Food demand in urban China: an application of a multi-stage censored demand system*. Tech. rep. Department of Agricultural, Environmental and Development Economics, Ohio State University.
- Liu, Q. et al. (2020). 'The experiences of health-care providers during the COVID-19 crisis in China: a qualitative study'. In: *The Lancet Global Health* 8, e790–e798.
- Liyanage, L.H. and J.G. Shanthikumar (2005). 'A practical inventory control policy using operational statistics'. In: *Operations Research Letters* 33, pp. 341–348.
- Marsten, R. et al. (1990). 'Interior point methods for linear programming: Just call Newton, Lagrange, and Fiacco and McCormick!' In: *Interfaces* 20, pp. 105–116.
- Masuda, E. and P. Newman (1981). 'Gray and Giffen goods'. In: *Economic Journal* 91, pp. 1011–1014.
- Matsuyama, K. (2006). 'The multi-period newsboy problem'. In: *European Journal of Operational Research* 171, pp. 170–188.
- Miller, A.C. and T.R. Rice (1983). 'Discrete approximations of probability distributions'. In: *Management Science* 29, pp. 352–362.
- Mills, E.S. (1959). 'Uncertainty and price theory'. In: *Quarterly Journal of Economics* 73, pp. 116–130.

-
- Miranda, S.I. (2014). *Superquantile regression: theory, algorithms, and applications*.
Tech. rep. Naval Postgraduate School, Monterey, CA.
- Moon, I. and E.A. Silver (2000). ‘The multi-item newsvendor problem with a budget constraint and fixed ordering costs’. In: *Journal of the Operational Research Society* 51, pp. 602–608.
- Moritz, B.B., A.V. Hill and K.L. Donohue (2013). ‘Individual differences in the newsvendor problem: behavior and cognitive reflection’. In: *Journal of Operations Management* 31, pp. 72–85.
- Morse, P.M. and G.E. Kimball (1951). *Methods of Operations Research*. Revised.
Cambridge, MA: MIT Press.
- Nahmias, S. and C.P. Schmidt (1984). ‘An efficient heuristic for the multi-item newsboy problem with a single constraint’. In: *Naval Research Logistics Quarterly* 31, pp. 463–474.
- Nelder, J.A. and R. Mead (1965). ‘A simplex method for function minimization’. In: *The Computer Journal* 7, pp. 308–313.
- Nemirovski, A. et al. (2009). ‘Robust stochastic approximation approach to stochastic programming’. In: *SIAM Journal on Optimization* 19, pp. 1574–1609.
- NHS Statistics on COVID-19 hospital activity* (2020). URL: www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity.
- Niederhoff, J.A. (2007). ‘Using separable programming to solve the multi-product multiple ex-ante constraint newsvendor problem and extensions’. In: *European Journal of Operational Research* 176, pp. 941–955.
- Ord, K., R.A. Fildes and N. Kourentzes (2017). *Principles of Business Forecasting*.
New York: Wessex Press.
- Pal, B., S.S. Sana and K. Chaudhuri (2015). ‘A distribution-free newsvendor problem with nonlinear holding cost’. In: *International Journal of Systems Science* 46, pp. 1269–1277.

-
- Pantumsinchai, P. and T.W. Knowles (1991). ‘Standard container size discounts and the single-period inventory problem’. In: *Decision Sciences* 22, pp. 612–619.
- Petruzzi, N.C. and M. Dada (1999). ‘Pricing and the newsvendor problem: A review with extensions’. In: *Operations Research* 47, pp. 183–194.
- Porteus, E.L. (2002). *Foundations of Stochastic Inventory Theory*. Stanford, CA: Stanford University Press.
- Prékopa, A. (2013). *Stochastic Programming*. Dordrecht: Springer.
- R Core Team (2020). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing.
- Ren, Y. and R. Croson (2013). ‘Overconfidence in newsvendor orders: an experimental study’. In: *Management Science* 59, pp. 2502–2517.
- Riches, G. (2018). *Food Bank Nations: Poverty, Corporate Charity and the Right to Food*. London: Routledge.
- Rios, L.M. and N.V Sahinidis (2013). ‘Derivative-free optimisation: a review of algorithms and comparison of software implementations’. In: *Journal of Global Optimisation* 56, pp. 1247–1293.
- Rockafellar, R.T., J.O. Royset and S.I. Miranda (2014). ‘Superquantile regression with applications to buffered reliability, uncertainty quantification, and conditional value-at-risk’. In: *European Journal of Operational Research* 234, pp. 140–154.
- Rockafellar, R.T. and S. Uryasev (2002). ‘Conditional value-at-risk for general loss distributions’. In: *Journal of Banking & Finance* 26, pp. 1443–1471.
- Scarf, H. (1958). *Studies in the Mathematical Theory of Inventory and Production*. Stanford, CA: Stanford University Press.
- Schweitzer, M.E. and G.P. Cachon (2000). ‘Decision bias in the newsvendor problem with a known demand distribution: experimental evidence’. In: *Management Science* 46, pp. 404–420.

-
- Shamir, R. (1987). ‘The efficiency of the simplex method: a survey’. In: *Management Science* 33, pp. 301–334.
- Shin, H. et al. (2015). ‘A classification of the literature on the planning of substitutable products’. In: *European Journal of Operational Research* 246, pp. 686–699.
- Silver, E.A., D.F. Pyke and R. Peterson (1998). *Inventory Management and Production Planning and Scheduling*. 3rd ed. New York: Wiley.
- Solis, F.J. and R.J.-B. Wets (1981). ‘Minimization by random search techniques’. In: *Mathematics of Operations Research* 6, pp. 19–30.
- Su, X. (2008). ‘Bounded rationality in newsvendor models’. In: *Manufacturing & Service Operations Management* 10, pp. 566–589.
- Syntetos, A.A. et al. (2016). ‘Supply chain forecasting: Theory, practice, their gap and the future’. In: *European Journal of Operational Research* 252, pp. 1–26.
- Tashman, L.J. (2000). ‘Out-of-sample tests of forecasting accuracy: An analysis and review’. In: *International Journal of Forecasting* 16, pp. 437–450.
- Turken, N. et al. (2012). ‘The multi-product newsvendor problem: review, extensions, and directions for future research’. In: *Handbook of Newsvendor Problems*. Ed. by T.-M. Choi. New York: Springer, pp. 3–39.
- Vanderbei, R.J. (2020). *Linear Programming: Foundations and Extensions*. New York: Springer.
- Vapnik, V. (1998). ‘The support vector method of function estimation’. In: *Nonlinear Modeling*. Ed. by J.A.K. Suykens and J. Vandewalle. Boston, MA: Springer, pp. 55–85.
- Von Luxburg, U. and B. Schölkopf (2011). ‘Statistical learning theory: Models, concepts, and results’. In: *Handbook of the History of Logic*. Ed. by D.M. Gabbay, S. Hartmann and J. Woods. Vol. 10. Amsterdam: Elsevier, pp. 651–706.
- Wang, T. (2011). ‘Sensitivity analysis for newsvendor model with general advertisement sensitive demand’. In: *Procedia Engineering* 15, pp. 1209–1213.

-
- Webby, R. and M. O'Connor (1996). 'Judgemental and statistical time series forecasting: a review of the literature'. In: *International Journal of Forecasting* 12, pp. 91–118.
- Wendell, R.E. (1985). 'The tolerance approach to sensitivity analysis in linear programming'. In: *Management Science* 31, pp. 564–578.
- Wu, M., S.X. Zhu and R.H. Teunter (2013). 'The risk-averse newsvendor problem with random capacity'. In: *European Journal of Operational Research* 231, pp. 328–336.
- (2014). 'A risk-averse competitive newsvendor problem under the CVaR criterion'. In: *International Journal of Production Economics* 156, pp. 13–23.
- Wu, X. and J.A. Niederhoff (2014). 'Fairness in selling to the newsvendor'. In: *Production and Operations Management* 23, pp. 2002–2022.
- Xinsheng, X. et al. (2015). 'Optimal decisions for the loss-averse newsvendor problem under CVaR'. In: *International Journal of Production Economics* 164, pp. 146–159.
- Yao, F., H. Wen and J. Luan (2013). 'CVaR measurement and operational risk management in commercial banks according to the peak value method of extreme value theory'. In: *Mathematical and Computer Modelling* 58, pp. 15–27.
- Zhang, B. (2012). 'Multi-tier binary solution method for multi-product newsvendor problem with multiple constraints'. In: *European Journal of Operational Research* 218, pp. 426–434.
- Zhang, B., X. Xu and Z. Hua (2009). 'A binary solution method for the multi-product newsboy problem with budget constraint'. In: *International Journal of Production Economics* 117, pp. 136–141.
- Zhang, J., W. Xie and S.C. Sarin (2021). 'Multiproduct newsvendor problem with customer-driven demand substitution: a stochastic integer program perspective'. In: *INFORMS Journal on Computing* 33, pp. 1229–1244.

Zhang, W. and Q. Wang (2003). *Changes in China's urban food consumption and implications for trade*. Tech. rep. Department of Agricultural Economics, Michigan State University.

Zhang, Y. and E. Siensen (2019). 'A meta-analysis of newsvendor experiments: revisiting the pull-to-center asymmetry'. In: *Production and Operations Management* 28, pp. 140–156.

Zipkin, P.H. (2000). *Foundations of Inventory Management*. Burr Ridge, IL: McGraw-Hill.