RISK-AVERSE SELECTIVE NEWSVENDOR PROBLEMS

by

Arleigh Catherine Waring

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
2012

Doctoral Committee:
Professor H. Edwin Romeijn, Chair
Professor Lawrence Seiford
Associate Professor Amitabh Sinha
Associate Professor Mark Van Oyen

For my grandmothers,
Catherine Carr O'Donnell
and
Mary Gonder Waring

ACKNOWLEDGEMENTS

As this dissertation marks the end of my time at the University of Michigan, I would like to take the time to thank the people who inspired, supported, and encouraged me.

I would like to thank my advisor, Edwin Romeijn, without whom this dissertation would not be possible. I am so grateful that you were willing to work with me in 2009 and I appreciate your guidance and encouragement. Thank you for your time and support.

Many thanks to my dissertation committee. Thanks to Larry Seiford for always taking time to meet with me and brainstorm solutions to the problems I faced in research and in life. Amitabh Sinha always provided valuable insight and feedback for which I am extremely appreciative. Finally, thanks to Mark Van Oyen who found time to review my project in a very hectic year for him.

I owe a debt of gratitude to all the faculty and staff in the Industrial and Operations Engineering department who have supported me throughout my time as a graduate student.

I would also like to thank Joe Hartman and Andrew Ross who advised me throughout my undergraduate career at Lehigh University and encouraged me to go to graduate school.

I would like to acknowledge my mentors, Tershia Pinder-Grover and Toni Benner, who always offered guidance and provided thoughtful advice. I hope they are not offended when I say that I want to be just like them when I grow up.

I am eternally grateful to Ada Barlatt and Stan Dimitrov for their support and encouragement over the past year. Knowing that they believed in me helped me believe in myself.

I am so lucky to have made such wonderful friends over the years and their support has meant the world to me. Special thanks to Katrina Appell, Erika Murdock Balbuena, Megan DeFauw, Rebecca Devlin, Irina Dolinskaya, Michael Lau, Marcial Lapp, Kristin Teufel Miller, Dianne Hosford Morales, Dan Nathan-Roberts, Lauren Spranklin, Tara Terry, Lauren Van Hoesen, Andrew Wenri, Neal Wiggermann, and Allison Williams.

I would like to thank all the McCuen, O'Donnell, and Waring families whose love and support I have relied on throughout my life. I would especially like to thank the Michigan O'Donnell's: Paul, Linda, Erin, Kyle, Dana and Ian, as well as Phil Giroux for making Michigan a home.

Most importantly, I would like to thank my parents, Mary and Craig, and my brothers, Colin and Gavin. Thank you for believing in me, encouraging me, and loving me no matter what. Mom and Dad, I appreciate all that you have done for me and know that there is no way I could ever repay you. Colin, Gavin, and I are so lucky to have you as parents. Please know that we owe all our successes to you.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

CHAPTER I

Introduction

1.1 Problem Overview

Growing interest in risk assessment and operations approaches to managing risk in global supply chains has led to increased research incorporating risk preferences into classic inventory management problems, such as the newsvendor. The goal of any inventory management problem is to successfully choose an inventory quantity to stock prior to the realization of demand so that their objectives are met. In the past, the assumption was that the newsvendor was risk-neutral and therefore its objective was to maximize expected profits. In practice, the decision maker typically incorporates additional factors into their choice of inventory level. This dissertation focuses specifically on risk-averse decision makers. Risk-aversion is commonly defined as a willingness to accept lower returns (expected profit) if there is more certainty in receiving them.

Risk-aversion is a behavior trait ascribed to the decision maker and there are various methods to represent these preferences mathematically in optimization problems. A common approach is to use an expected utility function of the random profit. More recently, inventory problems have incorporated specific risk measures derived from the finance industry. From a finance perspective, a risk measure is a function applied to a model with uncertain payoffs that can be used to determine the amount to be kept in reserve so that the risks taken by financial institutions are accepted by regulators. This dissertation focuses on the application of two such risk measures, namely Value-at-Risk and Conditional Value-at-Risk.

We utilize the measures Value-at-Risk and Conditional Value-at-Risk to study the risk-averse version of the *Selective Newvendor Problem*. This inventory policy problem incorporates a market selection decision into the standard and well-known newsvendor problem. Specifically, we consider a firm that procures and delivers a good within a single selling season in a number of different markets. Prior to the selling season, the firm determines how much to procure and also in which markets to operate. A previously studied risk-neutral version of this problem showed that the optimal order quantity also took the form of a critical fractile solution in which the critical fractile is the proportion of the cost of being understocked to the total cost. The most notable result however was that this critical fractile is independent of the market selection decision yielding the instrumental result that the optimal market selection solution can be found among a small number of candidate solutions satisfying an intuitively appealing ranking structure. We show that this result can be extended to the risk-averse selective problem for a decision maker using either Value-at-Risk or Conditional Value-at-Risk.

Next, we evaluate the effect of changing risk preferences for a selective newsvendor type problem. We analyze how the market selection decision changes as the decision maker's level of risk-aversion fluctuates. Solving the risk-averse problem and the risk-neutral problem separately yields two collections of market selections which suggest satisfying a certain market is optimal. The question that arises is how this collection changes under various risk-averse considerations. A common thought in supply chain is that extending to additional markets increases risk as the exposure to uncertainty increases. However, it is also well known that diversifying an investment portfolio by including additional assets tends to mitigate risk. Our analysis stems from the desire to understand if and how a selective newsvendor's risk changes based on their decision to operate within various markets. We examine this problem by creating a risk-reward Pareto efficient frontier to examine the tradeoff between the specified risk measure (Value-at-Risk and Conditional Value-at-Risk) and expected profit. In the next section we provide some more detailed background regarding risk-aversion and the newsvendor problem as well as a review of pertinent literature. The final section in this chapter details the structure of the dissertation.

1.2 Literature Review

The uncertainty of demand and its effect on inventory management has created a rich field of study in operations research. First referenced by Edgeworth [10] in the late nineteenth century, one of the most common inventory management problems is the newsvendor problem. The newsvendor has received significant attention since Arrow et al. [3] derived the critical fractile solution. In the classic problem, the newsvendor faces an unknown demand for a single product in a single period and makes an ordering decision prior to the realization of demand in order to maximize expected profit. The well-known critical fractile solution expresses the optimal order quantity as the quantity for which the probability that demand exceeds that quantity is equal to a critical fractile whose value depends only on the unit cost and revenue parameters, including revenues from sales and salvaging, and procurement, shortage, and/or expediting costs.

Over the years, numerous variations on the original model have been introduced. Such extensions include price dependent demand, various buyer and supplier pricing policies including quantity discounts, multiple product models and multiple echelon systems. Relevant to this work include extensions to the newsvendor that include alternative objective functions, risk-aversion and market selection which will be described in further detail below. Both Khouja [18] and Qin et al. [28] provide comprehensive reviews of the single period newsvendor problem and existing extensions.

After the introduction of the modern newsvendor problem came the realization that decision makers cared about alternative objectives in addition to maximizing expected profit. Subsequently, in the late 1970s and early 1980s researchers began proposing alternative objectives for the standard newsvendor problem. Kabak and Schiff [17] solved the classic problem under a "satisficing" objective function with the goal of maximizing the probability of achieving some set profit target. Ismail and Louderback [15] proposed the satisficing objective as an alternative measure to study the risk-reward tradeoff, the risk-reward tradeoff of course dating back to the novel portfolio optimization theory derived by Markowitz [23]. Lau [20] extended the analysis regarding the satisficing objective and added an additional objective of maximizing the expected utility to consider the mean-standard deviation of profit tradeoff.

Following the stock market crash of 1987, there was increased interest regarding risk preferences in the finance industry which eventually spread to operations management. One of the byproducts of this increased interest in risk was the development of the Value-at-Risk risk measure by J.P. Morgan who published its methodology in a 1994 RiskMetrics™ report [29]. Value-at-Risk defines a threshold such that the probability losses exceed this threshold in a given time period is less than a predetermined level. To this day, Value-at-Risk is widely used by banks, securities firms, and other trading organizations.

Around the same time as the development of Value-at-Risk, Eeckhoudt et al. [11] published their seminal results regarding a risk-averse newsvendor. Using an expected utility objective function related to profit their work indicated that the optimal order quantity for the risk-averse newsvendor is less than that of the risk-neutral newsvendor. Agrawal and Seshadri [1] incorporated a pricing decision into a similar problem that also maximized the expected utility associated with the newsvendor's profit. Zhou and Zhau [43] studied a newsvendor who maximized expected utility subject to a service-level constraint.

Rather than measure risk-aversion with an expected utility function, some authors incorporate quantitative risk measures such as Value-at-Risk explicitly. Jammernegg and Kischka [16] studied a newsvendor who aimed to maximize expected profit subject to a constraint on service and loss. Gan et al. [13] considered a single newsvendor maximizing expected profit subject to a Value-at-Risk constraint. Özler et al. [25] consider the multi-product version of that problem. Chiu and Choi [8] study the joint stocking and pricing decisions for a newsvendor who optimizes its Value-at-Risk.

Despite its widespread use, Value-at-Risk has several mathematical limitations in that it lacks subadditivity and convexity and is therefore not a coherent risk measure. Artzner et al. [4] defined four attributes associated with coherent risk-measures and first introduced the idea of "tail conditional expectation" as an alternative to Value-at-Risk that more accurately captures the magnitude of the smallest profits. Rockafellar and Uryasev [30] coined the term Conditional Value-at-Risk and provided an analytical expression to compute it when the underlying probability distribution is continuous. Pflug [26] proved that Conditional Value-at-Risk as presented by

Rockafellar and Uryasev [30] is in fact a coherent risk measure.

As with expected utility functions and Value-at-Risk, the use of coherent risk measures spread to the newsvendor problem. Ahmed et al. [2] derived a general form for the structure of the optimal solution of the newsvendor problem under coherent risk measures. Choi and Ruszczyński [9] derived a mean-risk model for a risk-averse newsvendor problem with general coherent measures of risk.

Additionally, several works consider the risk-averse newsvendor using a Conditional Value-at-Risk criterion. Chen et al. [7] consider a risk-averse newsvendor with price dependent demand who adopts a Conditional Value-at-Risk performance measure. The authors provide sufficient conditions for the optimal ordering and pricing policy under two types of demand models. Xu [39] extends these results to include a newsvendor with emergency procurement. Gotoh and Takano [14], Xu and Chen [40], and Xu and Li [41, 42] evaluate a newsvendor optimizing a linear combination of expected profit and Conditional Value-at-Risk.

This dissertation focuses on the risk-aversion for the so-called *Selective Newsvendor Problem* using Value-at-Risk and Conditional Value-at-Risk risk measures. Introduced by Taaffe et al. [36], the selective newsvendor procures and delivers a good within a single selling season to a number of different markets. The price for the good is market dependent and each market has an independent demand distribution. The risk-neutral version of this problem maximizes expected profit and has an optimal solution that satisfies an intuitively appealing ranking scheme.

Bakal et al. [5] integrated a pricing decision into the market selection decision and described an optimal market-specific pricing policy. Strinka et al. [34] studied a class of selective newsvendor problems in which a decision maker must choose the most profitable combination of customizations from a set of raw materials which can be customized shortly before satisfying demand.

As in this dissertation, the aforementioned authors assumed that the demands of different markets are independent and normally distributed. Taaffe et al. [37] consider a version of this problem in which demand for an individual market is defined by a Bernoulli distribution such that the amount ordered is "all or nothing." This normality assumption is relaxed by Strinka and Romeijn [33] who develop approximation algorithms for selective newsvendor type problems.

As in the standard newsvendor analysis there has been an increasing shift towards incorporating the risk preferences of the selective newsvendor type decision maker. Taaffe et al. [35] used a mean-variance approach to solve a risk-averse version of the newsvendor problem with market selection. Chahar and Taaffe [6] utilized a Conditional Value-at-Risk criterion for a selective newsvendor who faces "all or nothing" type demand orders and provided analytical results showing how order selection is affected by the competing expected profit and Conditional Value-at-Risk objectives. This dissertation studies the risk-averse *Selective Newsvendor Problem* in which market demands are independent and normally distributed by utilizing Conditional Value-at-Risk and Value-at-Risk, two common financial risk measures.

1.3 Dissertation Outline

The remainder of this dissertation is structured as follows. In Chapter II we review the standard notation associated with the traditional newsvendor and introduce the notation that will identify the selective newsvendor problem throughout this document. We study a selective newsvendor facing a service level type constraint which results in a straightforward extension to the original model in which the optimal market selection can be found in a small sorted set of candidate solutions originally characterized by the solution to the risk-neutral model.

Chapter III considers the risk-averse selective newsvendor problem in which the objective is to maximize a convex combination of expected profit and Conditional Value-at-Risk. As detailed in the previous section, this problem has been studied for the newsvendor without market selection. We review those results and provide an alternate characterization of the optimal order policy to facilitate our analysis of the selective case. We show that the optimal solution to the weighted sum optimization also can be found in the same set as the risk-neutral model. These results allow us to create a concave envelope of the mean-Conditional Value-at-Risk Pareto efficient frontier.

Identifying the true Pareto efficient frontier for the mean-Conditional Value-at-Risk problem is the focus of Chapter IV. We first introduce several approximations to the frontier and then use said approximations to describe a branch and bound procedure that efficiently finds the Pareto efficient frontier. We provide some computational results to illustrate the intricacies of our algorithm.

In Chapter V we shift our analysis to a selective newsvendor utilizing a Value-at-Risk criterion. We establish that a selective newsvendor optimizing Value-at-Risk has an optimal solution in the set of ranked candidate solutions. We then consider the class of optimization problems associated with maximizing expected profit subject to a Value-at-Risk type constraint. Using the branch and bound algorithms described in the previous chapter we can find the mean-Value-at-Risk Pareto efficient frontier.

Finally, in Chapter VI we provide some concluding remarks and also suggest several potential areas for future study.

CHAPTER II

The Selective Newsvendor with Service Level Constraints

2.1 Introduction

As mentioned in Chapter I, this dissertation focuses on the *Selective Newsvendor Problem* in which a firm integrates their procurement and market selection decisions for a single period selling season. The *Selective Newsvendor Problem* was introduced to address the numerous planning issues that arise when a firm who places an order with a supplier prior to observing demand. The firm's demand which directs the procurement policy depends on a number of things, in this dissertation we consider demand that is affected by the markets it chooses to supply. Taaffe et al. [36] first described this model as a nonlinear and integer optimization problem and described a solution method to find both the procurement quantity and the market selection decision which maximized expected profit. For the remainder of the dissertation, we will refer to this as the risk-neutral version of the *Selective Newsvendor Problem*.

For many years, researchers have assumed that decision makers in inventory problems are indeed risk-neutral and their ultimate goal is to maximize expected profit. In practice, however, this is not always the case. Schweitzer and Cachon [32] investigated this discrepancy for the standard newsvendor problem and provided a list of potential sources for deviation from the original inventory policy. There are numerous preferences which result in firms choosing alternative inventory quantities. A new firm may want to ensure that all potential demand is met and be particularly averse to stockouts. A firm may particularly dislike disposing of excess inventory because it creates an image of wastefulness. In recent years there has been an increased discussion regarding the risk preferences of the newsvendor. A risk-averse newsvendor is willing to incur lower expected profits for a guarantee on demand. Chapters III, IV, and V deal specifically with a risk-averse version of *Selective Newsvendor Problem.*

In this chapter, we consider a selective newsvendor who wants to ensure a certain proportion of demand is met while maximizing expected profit. Our results are consistent with the risk-neutral selective newsvendor in that the optimal collection of markets can be found among a small number of candidate solutions. This straightforward extension provides the basis for all future chapters. The remainder of this chapter is structured as follows: Section 2.2 reviews the notation for the single and selective newsvendor that will be used throughout the dissertation, Section 2.3 extends the selective newsvendor problem to include a service level constraint, Section 2.4 illustrates the results and Section 2.5 provides some concluding remarks.

2.2 Notation and Problem Description

We first review the classic newsvendor problem and detail the notation that will be used throughout the dissertation. Consider a traditional newsvendor who orders items at unit cost c and earns unit revenue r for each item sold. The newsvendor must determine how much to order prior to the start of the selling season. At the end of the selling season, excess inventory has a unit salvage value of v and any excess demand is expedited at a unit cost of e. In other words, all demand is satisfied, either from inventory or through expediting. To avoid unrealistic and uninteresting situations we assume that $r, e > c > v$. Note that $e \leq r$ leads to a situation where units that are expedited yield a net profit of $r - e$, while $e > r$ means that satisfying demand by expediting results in a net loss of $e - r$ per unit expedited (we could, alternatively, interpret this as a case where demand can only be satisfied from inventory, with unsatisfied demand subject to a shortage cost of $e - r$).

The newsvendor faces random demand D with cumulative distribution function (c.d.f.) F and mean $E[D] = \mu$. For convenience, we assume that F has support $\mathbb{R}$ and is both invertible and differentiable (which is, for example, the case if F is the normal distribution), although results in this chapter can be generalized to situations where this is not the case. Moreover, where convenient we will let $\bar{F} = 1 - F$. We also assume that there is a fixed cost, S, associated with this newsvendor. Although somewhat trivial in the analysis of the single newsvendor, the idea of a fixed cost will have a significant role in the analysis of the selective newsvendor. Denoting the order quantity by Q, the newsvendor's profit function can be expressed as

$$\begin{aligned} \pi(Q; D) &= rD + v(Q - D)^+ - e(D - Q)^+ - cQ - S \\ &= (r - v)D - (c - v)Q - (e - v)(D - Q)^+ - S. \end{aligned} \tag{2.1}$$

A traditional (i.e., risk-neutral) newsvendor would maximize expected profit, given by

$$\begin{aligned} P(Q) &\equiv E[\pi(Q; D)] \\ &= (r - v)\mu - (c - v)Q - (e - v) \int_Q^\infty (x - Q)\, dF(x) - S \end{aligned} \tag{2.2}$$

which is known to be a concave function of Q and yield the optimal order quantity $Q^* = F^{-1}(\rho)$, where $\rho \equiv \frac{e-c}{e-v}$ is referred to as the critical fractile. One interpretation of this solution is that it provides the smallest supply quantity to guarantee that all demand will be met with probability at least $100\rho\%$. Accordingly, the profit maximizing solution results in a service level $100\rho\%$. This service level should not be confused with the fraction of demand actually satisfied by on-hand inventory which is referred to as the fill-rate.

The selective newsvendor problem incorporates a market selection component into the standard newsvendor problem. Prior to the selling season, the firm decides both which of m potential markets to serve and the total quantity to procure from the supplier. Consistent with current trends in inventory management, the inventory is centrally pooled for all markets (or customers). The market selection decisions are represented by a binary vector $y = (y_1, \ldots, y_m)^\top$ with $y_i = 1$ denoting that market i is selected and $y_i = 0$ that it is not $(i = 1, \ldots, m)$, while the order quantity is denoted by Q. The vector of market demands is denoted by $D = (D_1, \ldots, D_m)^\top$, where D_i is the demand in market i $(i = 1, \ldots, m)$. The total demand served by the firm is then $D_y = D^\top y$, and we denote its c.d.f. by F_y. As in the case of the traditional newsvendor we assume that F_y is invertible for all $y \in \{0,1\}^m$. We assume that market demands are statistically independent. There are several real world applications in which this might be the case. Consider a specialty store with stores with sufficient geographic separation or a firm that markets to specific market segments which each have their own demand distribution.

We assume each unit is sold at a per unit revenue cost, r. Since inventory is pooled, we have a market independent unit ordering cost c, unit salvage value v, and unit expediting cost e. The expediting cost insures that we meet all demand regardless of our initial inventory level. We assume that there is a high cost, quick response supplier from which we can obtain additional units throughout the selling season. In addition to the cost and revenue parameters, we include a fixed cost S_i

associated with serving market i ($i = 1, \ldots, m$). Letting $S = (S_1, \ldots, S_m)^\top$ and $S_y = S^\top y$, we denote the profit and expected profit by

$$\begin{aligned} \pi(Q; D_y) &= rD_y + v(Q - D_y)^+ - e(D_y - Q)^+ - cQ - S_y \\ &= (r - v)D_y - (c - v)Q - (e - v)(D_y - Q)^+ - S_y. \end{aligned} \tag{2.3}$$

and

$$\begin{aligned} P(Q, y) &\equiv E[\pi(Q; D_y)] \\ &= (r - v)\mu_y - (c - v)Q - (e - v)\int_Q^\infty (x - Q)\, dF_y(x) - S_y. \end{aligned} \tag{2.4}$$

In the remainder of this chapter and throughout the dissertation we will follow Taaffe et al. [36, 35] and assume that $D \sim n(\mu, \Sigma)$ with $\mu = (\mu_1, \ldots, \mu_m)^\top$ and $\Sigma = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_m^2)$ (where n denotes the multivariate normal distribution). It then follows immediately that $D_y \sim n(\mu_y, \sigma_y^2)$ with $\mu_y = \mu^\top y$ and $\sigma_y^2 = y^\top \Sigma y = \mathbf{1}^\top \Sigma y$ where $\mathbf{1} = (1, \ldots, 1)^m \in \mathbb{R}^m$, and the last equality follows from the fact that the vector y is binary.

This assumption of normality allows us to introduce a key result in which we reparameterize the expression given by (2.4). Note that we can write the order quantity as a function of the market selection decision

$$Q_y = F_y^{-1}(\beta) = \mu_y + \Phi^{-1}(\beta)\sigma_y \tag{2.5}$$

for some $\beta \in (0, 1)$, where Φ is the c.d.f. of a standard normal random variable. Using this expression, we can show that

$$\int_{F_y^{-1}(\beta)}^\infty \left(x - F_y^{-1}(\beta)\right) \, dF_y(x) = \Lambda(\beta)\sigma_y$$

where $\Lambda(\beta) \equiv \int_{\Phi^{-1}(\beta)}^\infty \left(z - \Phi^{-1}(\beta)\right) \, d\Phi(z)$ denotes the loss function corresponding to the standard normal distribution. Incorporating these results into (2.4) yields

$$\begin{aligned} P(F_y^{-1}(\beta), y) &= (r-v)\mu_y - (c-v)\left(\mu_y + \Phi^{-1}(\beta)\sigma_y\right) - (e-v)\Lambda(\beta)\sigma_y - S_y \\ &= \sum_{i=1}^{m} \left((r-c)\mu_i - S_i\right) y_i - (c-v)\Phi^{-1}(\beta)\sigma_y - \Lambda(\beta)\sigma_y. \end{aligned}$$

For convenience, we define $\xi_i = (r-c)\mu_i - S_i$ $(i = 1, \ldots, m)$, $\xi = (\xi_1, \ldots, \xi_m)^\top$ and $\xi_y = \xi^\top y$. Then we reformulate the corresponding expected profit function as a function of β and y as

$$P\left(Q_y, y\right) = P\left(F_y^{-1}(\beta), y\right) = \xi_y - K(\beta)\sigma_y \tag{2.6}$$

where

$$K(\beta) = (c-v)\Phi^{-1}(\beta) + (e-v)\Lambda(\beta).$$

For a fixed β, maximizing expected profit results in a problem of the form

$$\begin{array}{rccc} \max & \xi_y & - & K(\beta)\sigma_y \\ \text{subject to:} & y & \in & \{0,1\}^m \end{array} \tag{SNP}$$

which has an intuitive optimal solution which we will now describe. Without loss of generality, assume that markets are ordered in nonincreasing order of the ratio ξ_i/σ_i^2, i.e.,

$$1 \le i < j \le m \quad \Rightarrow \quad \frac{\xi_i}{\sigma_i^2} \ge \frac{\xi_j}{\sigma_j^2}.$$

There exists a market selection that maximizes the function given by (2.6) with the property that, if market ℓ is selected, markets $1, \ldots, \ell - 1$ are selected as well. Defining $y^{(\ell)} = \left(y_1^{(\ell)}, \ldots, y_m^{(\ell)}\right)^\top$ with

$$y_i^{(\ell)} = \begin{cases} 1 & \text{for } i = 1, \ldots, \ell \\ 0 & \text{for } i = \ell + 1, \ldots, m \end{cases}$$

yields the set

$$Y = \left\{y^{(\ell)} : \ell = 0, \ldots, m\right\} \tag{2.7}$$

which contains an optimal solution to the selective newsvendor problem with independent and normally distributed demands. This reduces the number of potential market selection vectors under consideration from 2^m to just $|Y| = m + 1$.

If the optimal order quantity satisfies (2.5) with β a constant independent of y then expected profit is maximized by one of the solutions in Y. Taafe et al. [36] demonstrated and used this result with $\beta = \rho$ for the risk-neutral selective newsvendor problem, where $\rho = \frac{e-c}{e-v}$ and is equivalent to the critical fractile derived in the standard newsvendor problem. In the next section, we will show that a similar result holds for a selective newsvendor with a service level constraint.

2.3 The Selective Newsvendor with a Service Constraint

Consider a selective newsvendor with the objective to maximize expected profits subject to a lower bound on service level. For this problem, we define service level as the probability of meeting demand with current levels of inventory. Define $\kappa \in (0, 1)$ as the service level. We introduce a constraint that requires $\Pr(D_y \leq Q) = F_y(Q)$ to be at least κ. The optimization problem we study is

$$\begin{array}{rrcl} \max & P(Q, y) & & \\ \text{subject to:} & F_y(Q) & \geq & \kappa \\ & Q & \in & \mathbb{R} \\ & y & \in & \{0, 1\}^m. \end{array}$$

The assumption that F_y is invertible for all $y \in \{0, 1\}^m$ allows us to rewrite the service level constraint as

$$Q_y \geq F_y^{-1}(\kappa).$$

For a fixed y, $P(Q, y)$ is concave in Q and its unconstrained optimum is $Q_y = F_y^{-1}(\rho)$ where ρ is the well-known critical fractile solution. If $F_y^{-1}(\rho) \geq F_y^{-1}(\kappa)$, or equivalently, $\rho \geq \kappa$, then the unconstrained solution remains optimal. Alternatively, when $\rho < \kappa$ the constraint is binding.

Thus, for all $\kappa \leq \rho$, the risk-neutral order quantity remains the optimum. When $\kappa > \rho$ we must set $Q_y = F_y^{-1}(\kappa)$ to satisfy the constraint. Thus, the optimal order quantity for a selective newsvendor with a service level constraint is

$$Q_y^* = \begin{cases} F_y^{-1}(\rho) & \text{if } \kappa \leq \rho \\ F_y^{-1}(\kappa) & \text{if } \kappa > \rho. \end{cases}$$

In either case, the optimal order quantity can be expressed as (2.5). Thus we can rewrite (SNP(ϵ)) as

$$\begin{array}{rrcl} \max & \xi_y & - & K(\beta)\sigma_y \\ \text{subject to:} & \beta & = & \max(\kappa, \rho) \\ & y & \in & \{0, 1\}^m. \end{array}$$

Since β does not depend on y, this problem is equivalent to the problem described by (SNP) and therefore the optimal market selection solution is one of the candidate solutions in Y.

2.4 Illustrative Results

The previous section showed that at any given service level, $\kappa \in (0, 1)$, that one of the market selections in Y will maximize expected profit. This result has some broader implications. Mainly, if a selective newsvendor is satisfying a certain selection of markets there will be a corresponding range of ϵ service levels for which this market selection yields the maximum expected profit. Similarly, the selective newsvendor can make market selection decisions based on these results.

To illustrate these concepts we provide two five-market examples in which the markets are sorted such that $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \ldots \frac{\xi_5}{\sigma_5^2}$ and $\rho = \frac{2}{5}$. For each example, we vary κ from 0 to 1 and plot the corresponding expected profit for each of the sorted solutions including the sorted solution $y = \{\mathbf{0}\}$.

In Figure 2.1, the risk-neutral profit maximizing market selection $y = [1\ 1\ 1\ 1\ 1]$ is optimal for all values of $\kappa \in (0, 1)$. Thus no incentive exists for the selective newsvendor to not operate in all five markets. This is not always the case. In Figure 2.2, the risk-neutral optimum $y = [1\ 1\ 1\ 1\ 1]$ remains optimal only for $\kappa \in (0, 0.75)$. The dark black line shows how the expected profit varies as κ changes. Market selections $y = [1\ 1\ 1\ 1\ 0]$, $y = [1\ 1\ 1\ 0\ 0]$, $y = [1\ 1\ 0\ 0\ 0]$, and $y = [1\ 0\ 0\ 0\ 0]$ are all optimal at some point. In this example, as the service level increases the markets are removed from consideration in the reverse order in which they were added. That is, the market with the largest expected return to variance ratio is removed first.

2.5 Conclusion

In this chapter we studied the *Selective Newsvendor Problem* with a service level constraint and showed that there exists an efficient solution method based on an intuitive ranking scheme. We showed two examples of how the choice of market selection changes (or does not) for various service levels. The idea of the optimal market selection changing as the decision maker's preferences change inspired the work in the next several chapters in which we interpret the decision maker's preferences as risk-aversion. In Chapter III we study a risk-averse selective newsvendor who optimizes a weighted expected profit-CVaR function. In Chapters IV and V we find the Pareto efficient frontiers for selective newsvendors using CVaR and VaR risk measures, respectively.

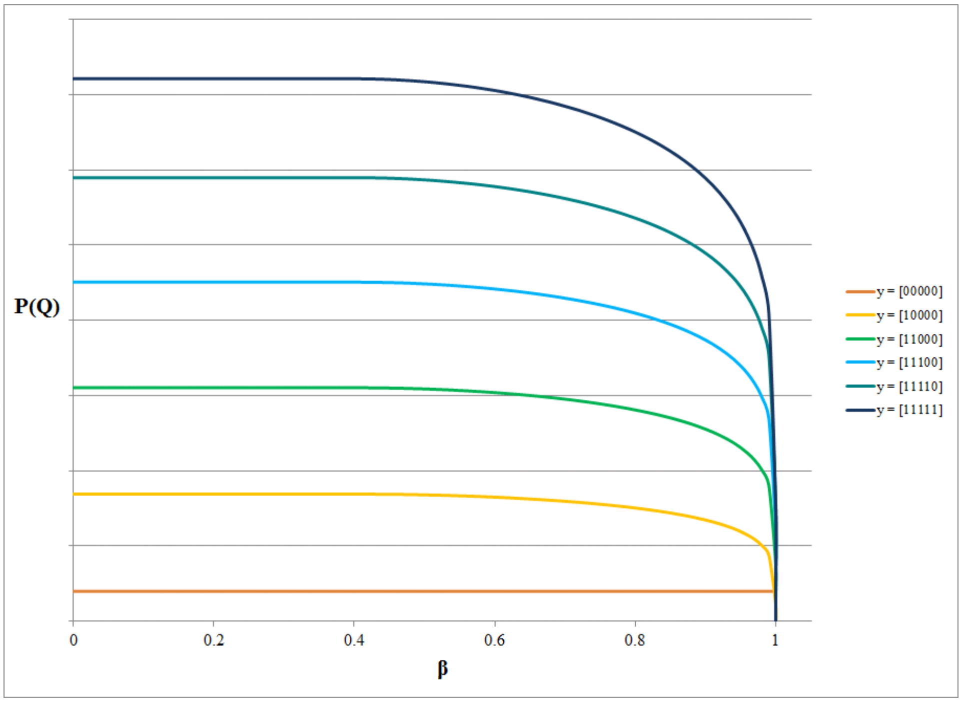

Figure 2.1: Example 1: Expected profit as a function of β for a 5 market example with $\rho = 0.4$.

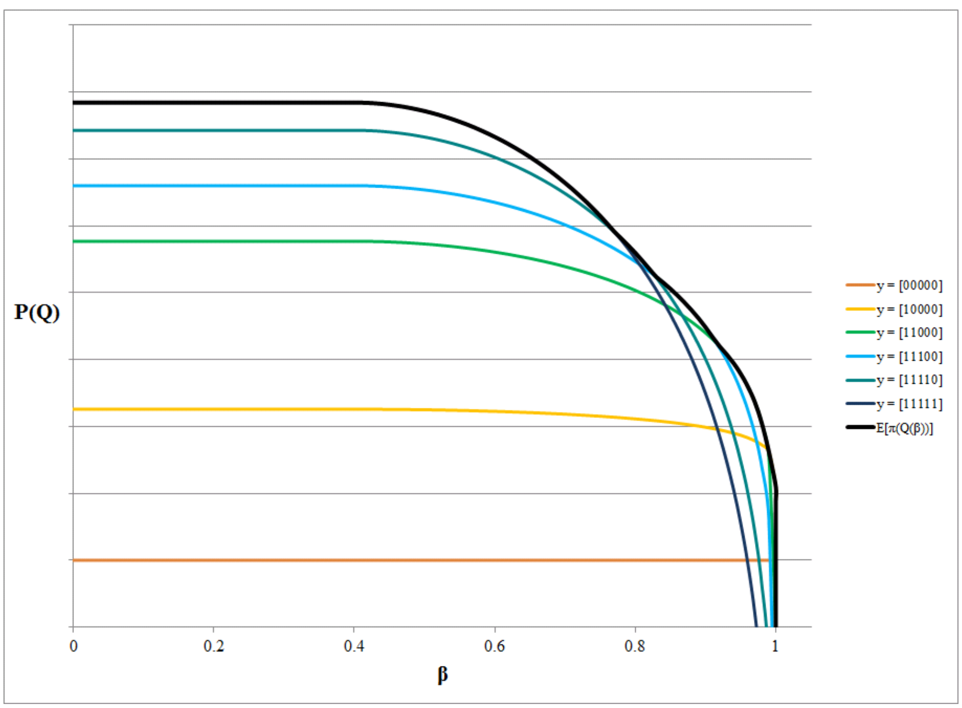

Figure 2.2: Example 2: Expected profit as a function of β for a 5 market example with $\rho = 0.4$.

CHAPTER III

Risk-Averse Selective Newsvendor Problems with a CVaR Risk Measure

3.1 Introduction

In this chapter we study a risk-averse version of the *Selective Newsvendor Problem* (SNP). It is known that a risk-neutral selective newsvendor maximizes expected profits according to the optimal critical fractile which is independent of the set of selected markets, and is then used to show that the optimal collection of markets can be found among a small number of candidate solutions. In Chapter II we described reasons why the decision maker in newsvendor problems does not always choose to use the profit maximizing optimal order quantity. We considered a specific case in which the decision maker wanted to enforce a service level constraint and showed that the risk-neutral selective newsvendor's optimal market selection policy extends to that of a selective newsvendor with a service constraint.

In this chapter, we consider a decision maker who is risk-averse. Risk-aversion is typically ascribed to an investor who is reluctant to invest in an opportunity with an uncertain payoff rather than an opportunity with a more certain, but possibly lower, expected payoff. With regards to investments, a risk-averse investor would frequently choose to invest in bonds and government based securities as opposed to the stock market. The application of risk-aversion has since been extended to the study of decision makers in operations research. Specifically, there has been a recent increase in analyzing the risk preferences of the decision makers in newsvendor problems (see, e.g., Eeckhoudt et al. [11] and Artzner et al. [4]).

Initially expected utility functions were the most common method of measuring a decision maker's risk preferences. The development of specific risk measures within the finance industry has also led to increased research in applying these risk measures to inventory management problems. Any problem has numerous risks associated with it and risk measures were introduce to mathematically assign a functional expression to evaluate said risk. An example would be Value-at-Risk, which is a threshold value for which there is a specified probability that losses exceed that threshold. There is also Conditional Value-at-Risk which is the expected value of the losses that exceeded the Value-at-Risk threshold. There are other risk measures but this dissertation will focus on these two.

We specifically want to study the tradeoffs that exist between the risk-neutral selective newsvendor policies and the risk-averse selective newsvendor policies. To study this tradeoff, in this chapter we analyze the risk-averse *Selective Newsvendor Problem* (SNP) with the objective is to maximize a convex combination of expected profit and Conditional Value-at-Risk (CVaR).

The single newsvendor version of this problem has been studied by Xu and Li [41], Gotoh and Takano [14] and Xu and Chen [40]. We recap their analysis and provide an alternate derivation and characterization of the optimal ordering policy in the shortage cost case that proves valuable when analyzing the selective newsvendor problem. We also provide additional managerial insights by formally analyzing the behavior of the optimal order quantity as a function of the problem parameters. In particular, we formalize the previously made observation that the optimal order quantity is not necessarily nonincreasing in the level of risk aversion.

Extending the single newsvendor results to the risk-averse *Selective Newsvendor Problem* allows us to extend the results of the risk-neutral SNP to some, but not all, risk-averse SNPs in the problem class in this chapter. The outline of this chapter is as follows. In Section 3.2 we review the case of a single risk-averse newsvendor and derive some new results, while in Section 3.3 we study the risk-averse selective newsvendor. Section 3.5 offers concluding remarks and directions for future research.

3.2 A Risk-Averse Newsvendor

In this section, we first review results from the literature regarding the risk-averse single newsvendor who optimizes a weighted sum of expected profit and Conditional Value-at-Risk (CVaR) associated with the optimal profit distribution. We then provide a new derivation of some existing results which is tailored towards facilitating the analysis of the risk-averse selective newsvendor problem in Section 3.3 of this chapter. As a byproduct, we are able to formally analyze the effects of risk preferences on the optimal order quantity in the single newsvendor problem, leading to some new managerial insights.

The traditional risk-neutral newsvendor was described in Chapter II with profit and expected profit given by (2.1) and (2.2), respectively. The optimal order quantity when maximizing expected profit is equal to the inverse of the well known critical fractile, mathematically, $Q^* = F^{-1}(\rho)$.

In this chapter, we consider a risk-averse newsvendor who uses the risk measure Conditional Value-at-Risk at some level, say $1-\alpha$ (CVaR_α). In particular, CVaR_α is the average profit in the left $(1-\alpha)$ tail of the profit distribution. When the profit distribution is absolutely continuous CVaR_α is equal to the expected profit conditional on it being no more than the Value-at-Risk at level $1-\alpha$ (VaR_α), which is defined implicitly through $\Pr\left(\pi(Q;D) \leq \text{VaR}_\alpha\right) = 1-\alpha$. The CVaR_α for a given order quantity (see, e.g., Rockafellar and Uryasev [30]) is given by

$$C_\alpha(Q) = \max_{\theta \in \mathbb{R}} \; C_\alpha(Q,\theta)$$

where

$$C_\alpha(Q,\theta) = \theta - \frac{1}{1-\alpha} E\left[(\theta - \pi(Q;D))^+\right]$$

and is a concave function of Q and θ (see, e.g., Pflug [27]). Note that $\alpha \in [0,1)$ reflects the degree of risk aversion, with larger values of α representing a higher degree of risk aversion. In particular, it can be shown that $C_0(Q) = P(Q)$, so that $\alpha = 0$ corresponds to a risk neutral newsvendor.

To examine the tradeoff between expected profit and CVaR, we maximize a weighted sum of expected profit and CVaR, which itself is a coherent risk measure (see, e.g., Artzner et al. [4] and Choi and Ruszczyński [9]). Define

$$G_\alpha(Q) \equiv \lambda P(Q) + (1-\lambda) C_\alpha(Q)$$

where $\lambda \in [0,1]$ is a weight parameter that characterizes the relative level of importance of the two criteria. Clearly, when $\lambda = 1$ this problem reduces to the traditional risk-neutral newsvendor problem, while when $\lambda = 0$ the newsvendor only optimizes CVaR_α. The optimization problem under consideration is

$$\begin{array}{lc} \max & G_\alpha(Q) \\ \text{subject to:} & Q \in \mathbb{R}. \end{array} \tag{P$_1$}$$

In order to further analyze our optimization problem it will be convenient to distinguish between the cases $e \leq r$ and $e > r$. The former leads to a situation where units that are expedited yield a net profit of $r - e$, while the latter means that satisfying demand by expediting results in a net loss of $e - r$ per unit expedited (we could, alternatively, interpret this as a case where demand can only be satisfied from inventory, with unsatisfied demand subject to a shortage cost of $e - r$). Section 3.2.1 describes the optimal order quantity as a function of λ for the case where $e \le r$, while Section 3.2.2 deals with the case where $e > r$.

3.2.1 Case 1: unit expediting cost does not exceed unit revenue

The first case under consideration is characterized by $e \le r$, which implies that satisfying demand for items through expediting remains profitable, although of course less so than satisfying demand from inventory. Xu and Li [41] study this problem in detail, and in this section we will summarize their results since we will build on those in the remainder of this paper. The following theorem derives the optimal order quantity as a function of λ.

Theorem 3.1 (cf. Xu and Li [41]). *The optimal order quantity for a risk-averse newsvendor problem with $e \le r$ who solves* (P$_1$)) *is*

$$Q_\alpha^*(\lambda) = \begin{cases} F^{-1}\left(\frac{\rho(1-\alpha)}{1-\lambda\alpha}\right) & \text{for } 0 \le \lambda \le \min\left\{1, \frac{1-\rho}{\alpha}\right\} \\ F^{-1}\left(1 - \frac{1-\rho}{\lambda}\right) & \text{for } \frac{1-\rho}{\alpha} < \lambda \le 1. \end{cases}$$

Xu and Li [41] analyzed the effect of risk aversion on the optimal order quantity by studying the behavior of $Q_\alpha^*(\lambda)$ both as a function of λ and α. Note that, for a given risk measure (i.e., a given value of α), increasing the value of λ decreases the relative importance of risk as compared to expected profit, and hence corresponds to a decreasing level of risk aversion. Similarly, for a given relative importance weight λ, increasing the value of α increases the importance of uncertainty, and hence corresponds to an increasing level of risk aversion. These relationships lead to the following corollary.

Corollary 3.2. *When $e \leq r$, the optimal order quantity $Q^*_\alpha(\lambda)$ is (i) nondecreasing in λ; and (ii) nonincreasing in α.*

Proof. Consider the optimal order quantity as a function of λ as given in Theorem 3.1. It immediately follows that $Q^*_\alpha(\lambda)$ is a continuous function of λ. In addition, since F^{-1} is a nondecreasing function, it is easy to see that $Q^*_\alpha(\lambda)$ is nondecreasing in λ as well. Next, consider the optimal order quantity $Q^*_\alpha(\lambda)$ for fixed $\lambda \in [0, 1]$ as a function of $\alpha \in [0, 1)$ by slightly rewriting the given expression for $Q^*_\alpha(\lambda)$ as

$$Q^*_\alpha(\lambda) = \begin{cases} F^{-1}\left(1 - \frac{1-\rho}{\lambda}\right) & \text{for } 0 \leq \alpha < \min\left\{1, \frac{1-\rho}{\lambda}\right\} \\ F^{-1}\left(\frac{\rho(1-\alpha)}{1-\lambda\alpha}\right) & \text{for } \frac{1-\rho}{\lambda} \leq \alpha < 1. \end{cases}$$

It follows that the optimal order quantity is both continuous and nonincreasing in α. □

These results show that, if $e \leq r$, the optimal order quantity is nonincreasing in the level of risk aversion – a result that does, perhaps surprisingly, not extend to the case where $e > r$ (see Section 3.2.2). Figures 3.1 and 3.2 illustrate the behavior of the order quantity as a function of both λ and α.

3.2.2 Case 2: unit expediting cost exceeds unit revenue

The second case is characterized by $e > r$, which implies that satisfying demand for items through expediting is costly. This is relevant whenever there either is a contractual obligation to satisfy all demands or when substantial other costs are associated with not satisfying demand. As mentioned earlier, the optimal order

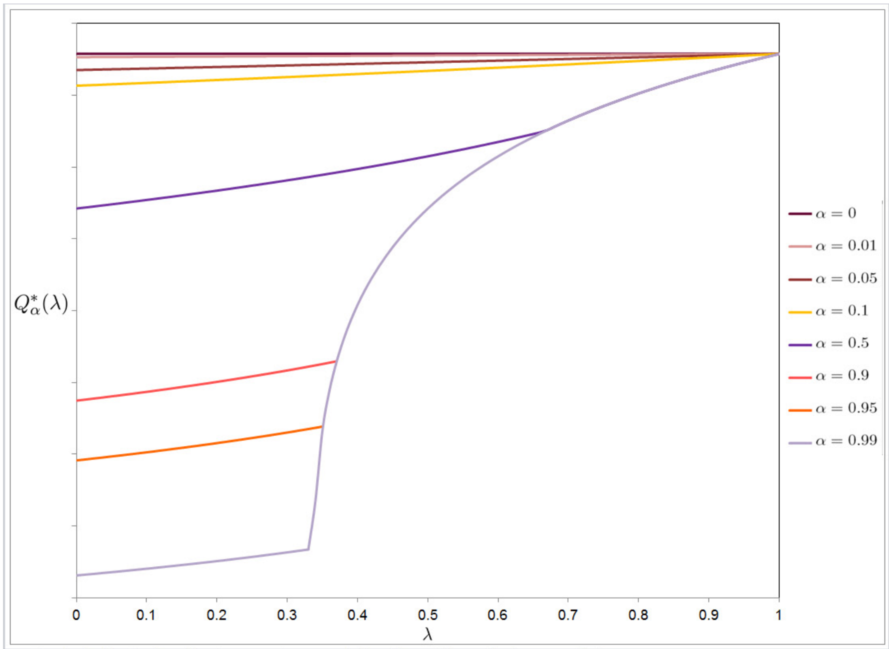

Figure 3.1: Optimal order quantity as a function of λ for different values of α (case 1: $e \leq r$).

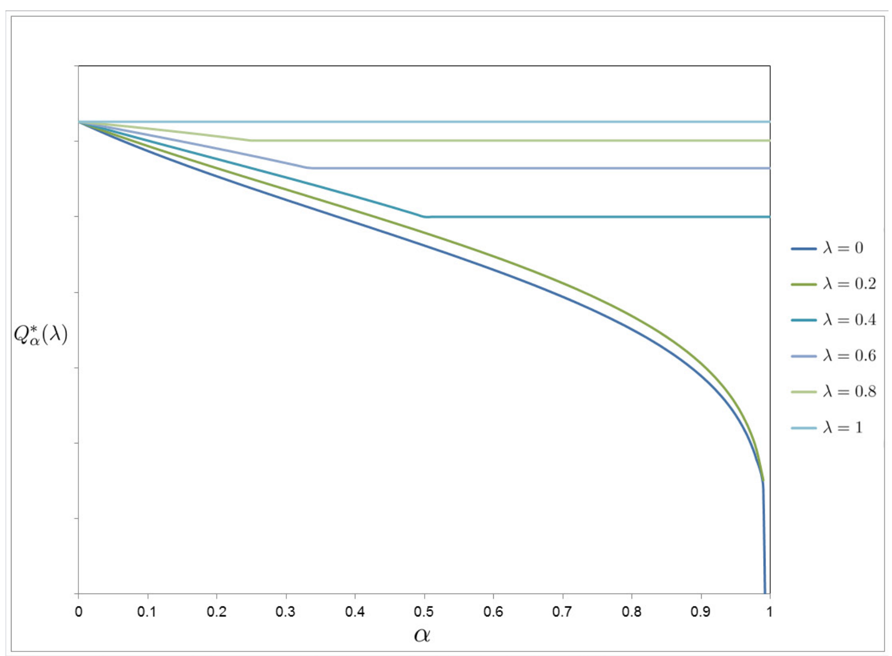

Figure 3.2: Optimal order quantity as a function of α for different values of λ (case 1: $e \leq r$).

quantity for this case has been studied in the literature (see Gotoh and Takano [14] and Xu and Chen [40]). However, we provide an alternative characterization and derivation that will allow us to generalize our model to incorporate market selection. The following theorem derives the optimal order quantity as a function of λ.

Theorem 3.3. *The optimal order quantity for a risk-averse newsvendor problem with $e \leq r$ who solves* (P_1)) *for $\lambda \in [0,1]$ is*

$$Q^*_\alpha(\lambda) = \left(\frac{r-v}{e-v}\right) F^{-1}\left(u_\alpha(\lambda)\right) + \left(\frac{e-r}{e-v}\right) F^{-1}\left(u_\alpha(\lambda)+\alpha\right)$$

where

(i) $u_\alpha(0) = \rho(1-\alpha)$*;*

(ii) for $\lambda > 0$: $u_\alpha(\lambda)$ is the unique solution to

$$\left(\frac{r-v}{e-v}\right) F^{-1}(u) + \left(\frac{e-r}{e-v}\right) F^{-1}(u+\alpha) = F^{-1}\left(\frac{\rho(1-\alpha)-(1-\lambda)u}{\lambda(1-\alpha)}\right).$$

Proof. It is easy to see that we can reformulate the objective function by explicitly incorporating the optimization problem that defines CVaR_α, so that the problem becomes

$$\begin{array}{lll} \max & G_\alpha(Q,\theta) & \\ \text{subject to:} & Q,\theta \in \mathbb{R}. & \end{array} \qquad (\mathrm{P}_1)$$

where

$$G_\alpha(Q,\theta) \equiv \lambda P(Q) + (1-\lambda) C_\alpha(Q,\theta)$$

which is an unconstrained optimization problem with concave objective function. Thus, it suffices to analyze the first order conditions. We first optimize $C_\alpha(Q,\theta)$ over θ for fixed Q. Recall that

$$C_\alpha(Q,\theta) = \theta - \frac{1}{1-\alpha} E[(\theta - \pi(Q;D))^+]$$

where the second term can be expressed as

$$
\begin{aligned}
&E\left[(\theta - \pi(Q))^+\right] \\
&= \int_{-\infty}^{Q} (\theta - (r-v)x + (c-v)Q)^+ \, dF(x) + \int_{Q}^{\infty} (\theta - (r-e)x - (e-c)Q)^+ \, dF(x).
\end{aligned}
$$

When $e > r$, this expression becomes

$$
E\left[(\theta - \pi(Q))^+\right] = \begin{cases} \int_{-\infty}^{\frac{\theta+(c-v)Q}{r-v}} (\theta - (r-v)x + (c-v)Q) \, dF(x) + & \\ \int_{Q}^{\frac{(e-c)Q-\theta}{e-c}} (\theta - (r-e)x - (e-c)Q) \, dF(x) & \text{if } \theta \le (r-c)Q \\ \int_{-\infty}^{Q} (\theta - (r-v)x + (c-v)Q) \, dF(x) + & \\ \int_{Q}^{\infty} (\theta - (r-e)x - (e-c)Q) \, dF(x) & \text{if } \theta > (r-c)Q. \end{cases}
$$

Taking the partial derivative of $C_\alpha(Q,\theta)$ with respect to θ yields:

$$
\frac{\partial C_\alpha(Q,\theta)}{\partial \theta} = \begin{cases} 1 - \frac{1}{1-\alpha}\left(F\left(\frac{\theta+(c-v)Q}{r-v}\right) - \bar{F}\left(\frac{(e-c)Q-\theta}{e-r}\right)\right) & \text{if } \theta < (r-c)Q \\ 1 - \frac{1}{1-\alpha} & \text{if } \theta > (r-c)Q. \end{cases}
$$

When $\theta = (r-c)Q$ both expressions are equal to

$$
1 - \frac{1}{1-\alpha} \tag{3.1}
$$

so that C_α is continuously differentiable with respect to θ for all Q. Since the derivative $\left(1 - \frac{1}{1-\alpha}\right)$ is negative we may restrict ourselves to values $\theta < (r-c)Q$. Moreover,

$$
\lim_{\theta \to -\infty} \frac{\partial C_\alpha(Q,\theta)}{\partial \theta} = 1
$$

so that a solution for θ exists. Setting the first term in $\frac{\partial C_\alpha(Q,\theta)}{\partial \theta}$ equal to zero immediately implies the first order condition

$$
F\left(\frac{\theta + (c-v)Q}{r-v}\right) - F\left(\frac{(e-c)Q - \theta}{e-r}\right) = \alpha \tag{3.2}
$$

for θ. Since, in general, an analytical solution to the above expression does not exist, we incorporate this first-order condition by reparametrizing the problem as one with a new decision variable, $u \in (0, 1-\alpha)$, by defining

$$u = F\left(\frac{\theta + (c-v)Q}{r-v}\right) = F\left(\frac{(e-c)Q-\theta}{e-r}\right) + \alpha$$

or, equivalently,

$$Q(u) = \left(\frac{r-v}{e-v}\right) F^{-1}(u) + \left(\frac{e-r}{e-v}\right) F^{-1}(u+\alpha)$$

$$\theta(u) = (e-c)\left(\frac{r-v}{e-v}\right) F^{-1}(u) - (c-v)\left(\frac{e-r}{e-v}\right) F^{-1}(u+\alpha). \tag{3.3}$$

Then, with a slight abuse of notation, let

$$\tilde{P}(u) = (r-v)\mu - (c-v)Q(u) - (e-v)\int_{Q(u)}^{\infty} (x - Q(u))\ dF(x)$$

and

$$\begin{aligned}
\tilde{C}_\alpha(u) &= \theta(u) - \frac{1}{1-\alpha} E\left[(\theta(u) - \pi(Q(u); D))^+\right] \\
&= \theta(u) - \frac{1}{1-\alpha}(r-v)\int_{-\infty}^{\frac{\theta(u)+(c-v)Q(u)}{r-v}} \left(\frac{\theta(u)+(c-v)Q(u)}{r-v} - x\right) dF(x) \\
&\qquad - \frac{1}{1-\alpha}(e-r)\int_{\frac{(e-c)Q(u)-\theta(u)}{e-r}}^{\infty} \left(x - \frac{(e-c)Q(u)-\theta(u)}{e-r}\right) dF(x) \\
&= \theta(u) - \frac{1}{1-\alpha}(r-v)\int_{-\infty}^{F^{-1}(u)} \left(F^{-1}(u) - x\right) dF(x) \\
&\qquad - \frac{1}{1-\alpha}(e-r)\int_{F^{-1}(u+\alpha)}^{\infty} \left(x - F^{-1}(u+\alpha)\right) dF(x).
\end{aligned} \tag{3.4}$$

The derivatives of these functions are:

$$\begin{aligned}
\tilde{P}'(u) &= -(c-v)Q'(u) + (e-v)\left(1 - F\left(Q(u)\right)\right) Q'(u) \\
&= (e-v)\left(\rho - F\left(Q(u)\right)\right) Q'(u)
\end{aligned}$$

and

$$\begin{aligned}
\tilde{C}'_\alpha(u) &= \theta'(u) - \frac{1}{1-\alpha}\left((r-v)(F^{-1})'(u)u - (e-r)(1-u-\alpha)(F^{-1})'(u+\alpha)\right) \\
&= (e-v)\left(\rho - \frac{u}{1-\alpha}\right) Q'(u).
\end{aligned}$$

Similarly, we can reformulate the original optimization problem in terms of u:

$$\begin{aligned} \max \quad G_\alpha(u) &\equiv \lambda\tilde{P}(u) + (1-\lambda)\tilde{C}_\alpha(u) \\ \text{subject to:} \quad u &\in (0, 1-\alpha). \end{aligned}$$

Thus, the first order condition is

$$\begin{aligned} G'_\alpha(u) &= \lambda\tilde{P}'(u) + (1-\lambda)\tilde{C}'(u) \\ &= (e-v)Q'(u)\left(\left(\rho - \frac{1-\lambda}{1-\alpha}u\right) - \lambda F\left(Q(u)\right)\right). \end{aligned}$$

We break this problem into two cases:

(i) $\lambda = 0$:

In this case the first order optimality condition is $u = \rho(1-\alpha)$, which is the desired result.

(ii) $\lambda > 0$:

In this case the first order optimality condition is

$$F^{-1}\left(\frac{\rho(1-\alpha) - (1-\lambda)u}{\lambda(1-\alpha)}\right) = Q(u).$$

The desired result now follows provided that this equation always has a unique solution for u. Note that the function Q is strictly increasing, and

$$\lim_{u \downarrow 0} Q(u) = -\infty \text{ and } \lim_{u \uparrow 1-\alpha} Q(u) = \infty.$$

Furthermore,

$$F^{-1}\left(\frac{\rho(1-\alpha) - (1-\lambda)u}{\lambda(1-\alpha)}\right)$$

is a nonincreasing (possibly extended real-valued) function of u for any $\lambda \in (0, 1]$ which means that a unique solution indeed exists. $\square$

As in Case 1, we are interested in the effect of risk aversion on the optimal order quantity. We start by analyzing the behavior of the implicitly defined function $u_\alpha(\lambda)$

as a function of both λ and α. The former allows us to formally characterize the behavior of $Q^*_\alpha(\lambda)$ as a function of λ. The latter will show that the behavior of $Q^*_\alpha(\lambda)$ as a function of α depends on the shape of the function F^{-1} and will be done experimentally.

Theorem 3.4. *The implicitly defined function $u_\alpha(\lambda)$ is*

(i) nonincreasing in α;

(ii) nonincreasing in λ when $\alpha \leq \bar{\alpha}$ and nondecreasing in λ when $\alpha \geq \bar{\alpha}$, where $\bar{\alpha}$ satisfies

$$\gamma F^{-1}\left((1-\bar{\alpha})\rho\right) + (1-\gamma)F^{-1}\left((1-\bar{\alpha})\rho + \bar{\alpha}\right) = F^{-1}(\rho).$$

Proof. Letting $\gamma = (r-v)/(e-v) \in (0,1]$ define the function

$$H_\alpha(u,\lambda) = \gamma F^{-1}(u) + (1-\gamma)F^{-1}(u+\alpha) - F^{-1}\left(\frac{\rho(1-\alpha)-(1-\lambda)u}{\lambda(1-\alpha)}\right)$$

so that $u_\alpha(\lambda)$ is defined by $H_\alpha(u,\lambda) = 0$. From Theorem 3.3 we know that this function exists and is unique. Moreover, since $H_\alpha(u,\lambda)$ is differentiable in u, λ, and α, the Implicit Function Theorem (see, e.g., Rudin [31]) can be used to determine the derivatives of $u_\alpha(\lambda)$ with respect to α and λ.

(i) The derivative of $u_\alpha(\lambda)$ with respect to α is

$$\frac{\partial u'_\alpha(\lambda)}{\partial \alpha} = -\frac{\frac{\partial H_\alpha(u,\lambda)}{\partial \alpha}}{\frac{\partial H_\alpha(u,\lambda)}{\partial u}}$$

where

$$\frac{\partial H_\alpha(u,\lambda)}{\partial \alpha} = (1-\gamma)(F^{-1})'(u+\alpha) + (F^{-1})'\left(\frac{\rho(1-\alpha)-(1-\lambda)u}{\lambda(1-\alpha)}\right)\cdot\frac{(1-\lambda)u}{\lambda(1-\alpha)^2}.$$

Since F^{-1} is increasing the desired result follows.

(ii) Similarly to (i), the derivative of $u_\alpha(\lambda)$ with respect to λ is

$$u'_\alpha(\lambda) = -\frac{\frac{\partial H_\alpha(u,\lambda)}{\partial \lambda}}{\frac{\partial H_\alpha(u,\lambda)}{\partial u}}$$

where

$$\frac{\partial H_\alpha(u,\lambda)}{\partial \lambda} = -(F^{-1})'\left(\frac{\rho(1-\alpha)-(1-\lambda)u}{\lambda(1-\alpha)}\right)\cdot\left(\frac{u}{\lambda^2(1-\alpha)} - \frac{\rho}{\lambda^2}\right)$$

and

$$\frac{\partial H_\alpha(u,\lambda)}{\partial u} = \gamma(F^{-1})'(u) + (1-\gamma)(F^{-1})'(u+\alpha) + (F^{-1})'\left(\frac{\rho(1-\alpha)-(1-\lambda)u}{\lambda(1-\alpha)}\right)\cdot\frac{(1-\lambda)u}{\lambda(1-\alpha)}.$$

Since F^{-1} is increasing it immediately follows that $u'_\alpha(\lambda) \geq 0$ if and only if $u \geq (1-\alpha)\rho = u_\alpha(0)$. It is easy to see that the function $u_\alpha(\lambda) = (1-\alpha)\rho$ for all λ when $\alpha = \bar{\alpha}$. Furthermore, it then follows from (i) that the function $u_\alpha(\lambda)$ is nondecreasing in λ for smaller values of α while it is nonincreasing for larger values. □

The following corollary now immediately follows from Theorem 3.4 the expression for the optimal order quantity in Theorem 3.3:

Corollary 3.5. *The optimal order quantity $Q^*_\alpha(\lambda)$ is (i) monotone in λ; and (ii) not necessarily monotone in α.*

Proof. This follows immediately by noting that

$$Q^{*\prime}_\alpha(\lambda) = \big(\gamma(F^{-1})'(u_\alpha(\lambda)) + (1-\gamma)(F^{-1})'(u_\alpha(\lambda)+\alpha)\big)\, u'_\alpha(\lambda)$$

which proves (ii) since the derivatives of $Q^*_\alpha(\lambda)$ and $u_\alpha(\lambda)$ with respect to λ have the same sign; and

$$\frac{\partial Q^*_\alpha(\lambda)}{\partial \alpha} = \gamma(F^{-1})'(u_\alpha(\lambda))\frac{\partial u_\alpha(\lambda)}{\partial \alpha} + (1-\gamma)(F^{-1})'(u_\alpha(\lambda)+\alpha)\left(\frac{\partial u_\alpha(\lambda)}{\partial \alpha} + 1\right)$$

which implies that the sign of the derivative of $Q^*_\alpha(\lambda)$ with respect to α depends on the function F^{-1}. $\square$

Figures 3.3 and 3.4 illustrate the behavior of the order quantity as a function of both λ and α for an example where demands are normally distributed. Figure 3.3 shows that, indeed, the order quantity is monotone in λ, but could be either nondecreasing or nonincreasing depending on the value of α. Figure 3.4 shows that the order quantity is not necessarily monotone in α. Overall, this confirms other studies by showing that the optimal order quantity is not necessarily nonincreasing in the level of risk aversion (see, e.g., Xu and Chen [40]).

3.3 A Risk-Averse Selective Newsvendor

The risk-neutral expected profit maximizing selective newsvendor problem introduced by Taaffe et al. [36] is described in Chapter II. Extending the notation, we denote the CVaR and objective functions by

$$\begin{aligned} C_\alpha(Q,y) &= \max_{\theta\in\mathbb{R}} \; C_\alpha(Q,y,\theta) \\ G_\alpha(Q,y) &= \lambda P(Q,y) + (1-\lambda) C_\alpha(Q,y) \end{aligned}$$

where

$$C_\alpha(Q,y,\theta) = \theta - \frac{1}{1-\alpha} E\left[(\theta - \pi(Q;D_y))^+\right].$$

and $P(Q,y)$ is given by (2.4). Our optimization problem is then

$$\begin{array}{llcl} \max & G_\alpha(Q,y) & & \\ \text{subject to:} & Q & \in & \mathbb{R} \\ & y & \in & \{0,1\}^m \end{array} \qquad (\mathrm{P}_2)$$

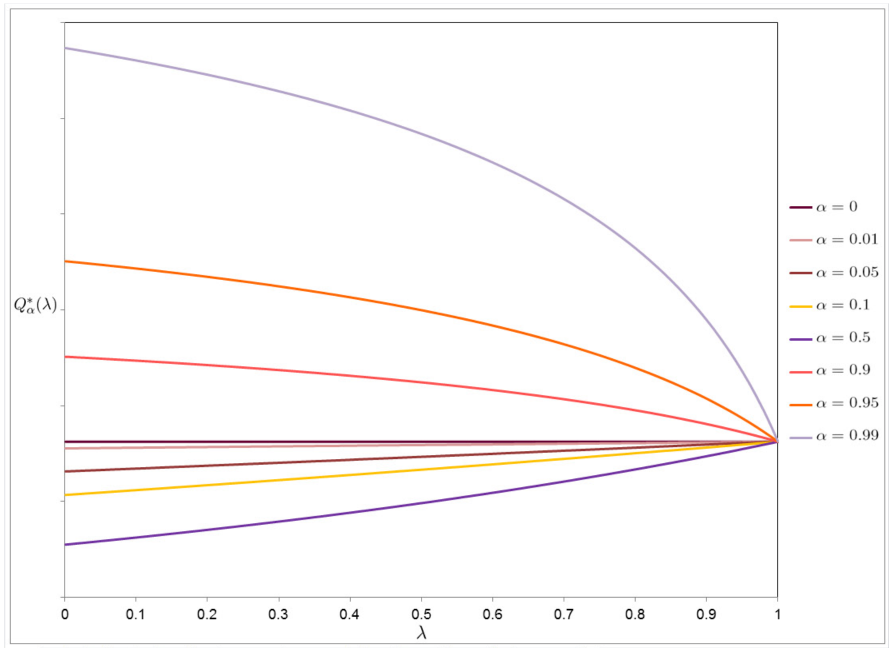

Figure 3.3: Optimal order quantity as a function of λ for different values of α (case 2: $e > r$).

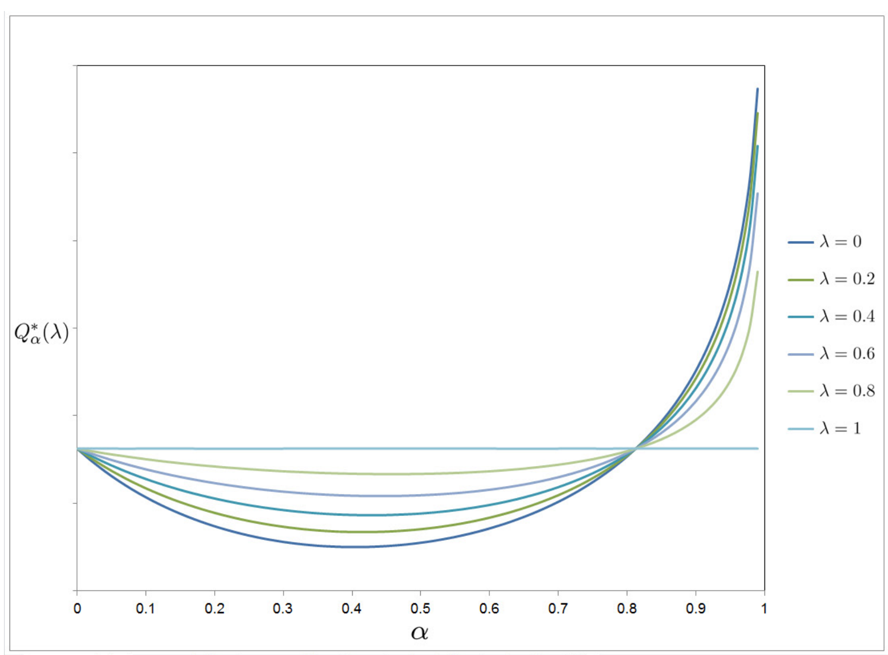

Figure 3.4: Optimal order quantity as a function of α for different values of λ (case 2: $e > r$).

where

$$G_\alpha(Q, y) \equiv \lambda P(Q, y) + (1 - \lambda) C_\alpha(Q, y).$$

It is clear that, for a fixed market selection vector y, the problem simply reduces to the one studied in Section 3.2. As we have shown previously, when demands are normally distributed and the order quantity as a function of y is given by (2.5), then expected profit can be written as (2.6). Then, there exists a market selection in the set Y (defined by (2.7)) that maximizes expected profit with the property that, if market ℓ is selected, markets $1, \ldots, \ell - 1$ are selected as well. As we will show below, a similar result holds for the risk-averse selective newsvendor when $e \leq r$ but not when $e > r$. Nevertheless, we are able to show that Y contains an optimal solution to the risk-averse selective newsvendor problem in either case.

3.3.1 Case 1: unit expediting cost does not exceed unit revenue

We begin by showing that when $e \leq r$, C_α exhibits a similar structure to (2.6).

Lemma 3.6. *Suppose that $e \leq r$ and Q_y is of the form* (2.5)*. Then*

$$C_\alpha\left(F_y^{-1}(\beta), y\right) = \xi_y - L_\alpha(\beta)\sigma_y$$

where $L_\alpha(\beta)$ only depends on β and the problem parameters.

Proof. Xu and Li [41, equation (7)] show that, for fixed selection vector y, the optimal value of θ as a function of the order quantity Q is given by

$$\theta_y^*(Q) = \begin{cases} (r - v)F_y^{-1}(1 - \alpha) - (c - v)Q & \text{if } F_y(Q) \geq 1 - \alpha \\ (r - e)F_y^{-1}(1 - \alpha) + (e - c)Q & \text{if } F_y(Q) < 1 - \alpha. \end{cases}$$

Moreover, consider the following term from the expression for CVaR:

$$
\begin{aligned}
&E\left[(\theta-\pi(Q;D_y))^+\right] \\
&= \int_{-\infty}^{Q} (\theta-(r-v)x+(c-v)Q)^+ \; dF_y(x) + \int_{Q}^{\infty} (\theta-(r-e)x-(e-c)Q)^+ \; dF_y(x)
\end{aligned}
$$

which, since $e \le r$,

$$
= \begin{cases}
\int_{-\infty}^{\frac{\theta+(c-v)Q}{r-v}} (\theta-(r-v)x+(c-v)Q) \; dF_y(x) & \text{if } \theta \le (r-c)Q \\
\int_{-\infty}^{Q} (\theta-(r-v)x+(c-v)Q) \; dF_y(x) + & \\
\int_{Q}^{\frac{\theta-(e-c)Q}{r-e}} (\theta-(r-e)x-(e-c)Q) \; dF_y(x) & \text{if } \theta > (r-c)Q.
\end{cases}
$$

Incorporating $\theta_y^*(Q)$ and $E\left[(\theta-\pi(Q;D_y))^+\right]$ into the expression for $C_\alpha(Q,y)$:

$$
C_\alpha(Q,y) = \begin{cases}
(r-e)F_y^{-1}(1-\alpha) + (e-c)Q - S_y - & \\
\frac{1}{1-\alpha}(r-e)\int_{-\infty}^{F_y^{-1}(1-\alpha)} \left(F_y^{-1}(1-\alpha)-x\right) \, dF_y(x) - & \\
\frac{1}{1-\alpha}(e-v)\int_{-\infty}^{Q} (Q-x) \; dF_y(x) & \text{if } Q \le F_y^{-1}(1-\alpha) \\
(r-v)F_y^{-1}(1-\alpha) - (c-v)Q - S_y - & \\
\frac{1}{1-\alpha}\cdot(r-v)\int_{-\infty}^{F_y^{-1}(1-\alpha)} \left(F_y^{-1}(1-\alpha)-x\right) \, dF_y(x) & \text{if } Q > F_y^{-1}(1-\alpha).
\end{cases}
$$

This then yields that

$$
C_\alpha\left(F_y^{-1}(\beta), y\right) = \xi_y - L_\alpha(\beta)\sigma_y
$$

where

$$
L_\alpha(\beta) = \begin{cases}
\frac{1}{1-\alpha}(r-e)\left(\Phi^{-1}(1-\alpha)+\Lambda(1-\alpha)\right) + & \\
\frac{1}{1-\alpha}(e-v)\left(\Phi^{-1}(\beta)+\Lambda(\beta)\right) & \\
-(r-e)\Phi^{-1}(1-\alpha) - (e-c)\Phi^{-1}(\beta) & \text{if } \beta \le 1-\alpha \\
\frac{1}{1-\alpha}(r-v)\left(\Phi^{-1}(1-\alpha)+\Lambda(1-\alpha)\right) & \\
-(r-v)\Phi^{-1}(1-\alpha) + (c-v)\Phi^{-1}(\beta) & \text{if } \beta > 1-\alpha.
\end{cases}
$$

Since $L_\alpha(\beta)$ is indeed only a function of β and the problem parameters we obtain the desired result. $\square$

From Chapter II we can see that, when $e \leq r$, the optimal order quantity is of the form (2.5) with β replaced by a value only depending on the problem parameters and the weighting factor λ for the selective newsvendor who optimizes the risk measure $G_\alpha(Q, y)$. Thus we know that for a fixed β, the problem (P_1) can be solved using a sorting solution. The following theorem summarizes these properties to show that an optimal solution to the risk-averse selective newsvendor problem can be found among the same set of candidate solutions as given in the risk-neutral selective newsvendor problem.

Theorem 3.7. *If $e \leq r$, the set Y contains an optimal solution to the risk-averse selective newsvendor problem. That is, an optimal solution exists such that if we select market ℓ, we also select markets $1, 2, \ldots, \ell - 1$.*

Proof. From Theorem 3.1 we know that $Q_y^* = F_y^{-1}(\beta)$ where

$$\beta = \begin{cases} \frac{\rho(1-\alpha)}{1-\lambda\alpha} & \text{for } 0 \leq \lambda \leq \min\left\{1, \frac{1-\rho}{\alpha}\right\} \\ 1 - \frac{1-\rho}{\lambda} & \text{for } \frac{1-\rho}{\alpha} < \lambda \leq 1 \end{cases}$$

depends only on the problem parameters. Combining (2.6) and Lemma 3.6 we obtain that

$$G_\alpha\left(F_y^{-1}(\beta), y\right) = \xi_y - \left(\lambda K(\beta) + (1-\lambda) L_\alpha(\beta)\right)\sigma_y. \tag{3.5}$$

This means that, for each $\lambda \in [0, 1]$, the coefficient of σ_y in the objective (3.5) is a constant independent of y, which yields the desired result. □

Thus, when $e \leq r$ we need only solve $m + 1$ optimization problems as opposed to enumerating all 2^m possibilities. Furthermore, it means that the set Y which includes the market selection that is optimal for the risk-neutral selective newsvendor also contains the the market selection that is optimal for the risk-averse selective newsvendor. In fact, these market selections may frequently be the same. In Section 3.4 we provide an example in which this is true.

3.3.2 Case 2: unit expediting cost exceeds unit revenue

Returning to the case in which unit expediting cost exceeds unit revenue such that satisfying items via expediting is especially costly. From Theorem 3.3 we can see that the situation is more complicated when $e > r$. In particular, the second condition of the theorem suggests that the optimal value of β in (2.5) may in fact depend on the selection vector y. Nevertheless an optimal solution to the risk-averse selective newsvendor problem can be found in the set Y and the proof is described below.

Theorem 3.8. *If $e > r$, the set Y contains an optimal solution to the risk-averse selective newsvendor problem. That is, an optimal solution exists such that if we select market ℓ, we also select markets* $1, 2, \ldots, \ell - 1$.

Proof. Using a reparametrization of the optimization problem in terms of u as in the proof of Theorem 3.3 we obtain

$$Q_y(u) = \mu_y + q(u)\sigma_y = F_y^{-1}\left(\Phi(q(u))\right)$$

where

$$q(u) = \left(\frac{r-v}{e-v}\right)\Phi^{-1}(u) + \left(\frac{e-r}{e-v}\right)\Phi^{-1}(u+\alpha).$$

We can then write the expected profit function in terms of u and y as follows:

$$\begin{aligned}\tilde{P}(u,y) &= (r-v)\mu_y - (c-v)Q_y(u) - (e-v)\int_{Q_y(u)}^{\infty} (x - Q_y(u))\, dF_y(x) - S_y \\ &= \xi_y - \left((c-v)q(u) + (e-v)\Lambda\left(\Phi(q(u))\right)\right)\sigma_y.\end{aligned}$$

Similarly, by substituting equation (3.3) into equation (3.4) we can write the CVaR function in terms of u and y as well:

$$\begin{aligned}
\tilde{C}_\alpha(u,y) = (e-c)\left(\frac{r-v}{e-v}\right)F_y^{-1}(u) - (c-v)\left(\frac{e-r}{e-v}\right)F_y^{-1}(u+\alpha)\\
-\frac{1}{1-\alpha}(r-v)\int_{-\infty}^{F_y^{-1}(u)}\left(F_y^{-1}(u)-x\right)\,dF_y(x)\\
-\frac{1}{1-\alpha}(e-r)\int_{F_y^{-1}(u+\alpha)}^{\infty}\left(x-F_y^{-1}(u+\alpha)\right)\,dF_y(x) - S_y\\
= \xi_y - \left((r-v)\left(\frac{1}{1-\alpha}-\frac{e-c}{e-v}\right)\Phi^{-1}(u) + (c-v)\left(\frac{e-r}{e-v}\right)\Phi^{-1}(u+\alpha)\right.\\
\left. + \frac{1}{1-\alpha}(r-v)\Lambda(u) + \frac{1}{1-\alpha}(e-r)\Lambda(u+\alpha)\right)\sigma_y.
\end{aligned}$$

In summary, this means that we can write

$$\tilde{P}(u,y) = \xi_y - \tilde{K}(u)\sigma_y$$

$$\tilde{C}_\alpha(u,y) = \xi_y - \tilde{L}_\alpha(u)\sigma_y$$

where

$$\tilde{K}(u) = (c-v)q(u) + (e-v)\Lambda\left(\Phi\left(q(u)\right)\right)$$

and

$$\begin{aligned}
\tilde{L}_\alpha(u) = (r-v)\left(\frac{1}{1-\alpha}-\frac{e-c}{e-v}\right)\Phi^{-1}(u) + (c-v)\left(\frac{e-r}{e-v}\right)\Phi^{-1}(u+\alpha)\\
+\frac{1}{1-\alpha}(r-v)\Lambda(u) + \frac{1}{1-\alpha}(e-r)\Lambda(u+\alpha)
\end{aligned}$$

which are both independent of y. We can then reformulate (P_2) as

$$\begin{array}{lrcl}
\max & \tilde{G}_\alpha(u,y) & \equiv & \lambda\tilde{P}(u,y) + (1-\lambda)\tilde{C}_\alpha(u,y)\\
\text{subject to:} & u & \in & (0, 1-\alpha)\\
 & y & \in & \{0,1\}^m.
\end{array}$$

Notice that for a fixed value of u, say u^*, the optimization problem is simply

$$\begin{array}{lrcl}
\max & \tilde{G}_\alpha(u^*,y) & \equiv & \lambda\tilde{P}(u^*,y) + (1-\lambda)\tilde{C}_\alpha(u^*,y)\\
\text{subject to:} & y & \in & \{0,1\}^m.
\end{array}$$

This means that, for each fixed u, we only need to consider selection vectors in Y, yielding the desired result. $\square$

In Section 3.3 we derived the optimal procurement and market selection policy for a risk-averse selective newsvendor who optimizes a weighted sum of expected profit and Conditional Value-at-Risk. In the next section we provide some illustrative results.

3.4 Computational Results

The results in Section 3.3 show that it is sufficient to use the results from Section 3.2 to solve a single-market risk-averse newsvendor problem corresponding to each of the $m+1$ market selection vectors in Y. In fact, since one of them ($y^{(0)} = \mathbf{0}$) leads to a trivial optimization problem, we have only m nontrivial problems to solve.

We illustrate the results that can be obtained using the methods developed in this paper on two problem instances with five markets in which the markets are ranked

$$\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2} \geq \frac{\xi_4}{\sigma_4^2} \geq \frac{\xi_5}{\sigma_5^2}.$$

See Figures 3.5 and 3.6. For both instances we started by drawing the tradeoff curves between expected profit and CVaR_α for each of the $m+1$ solutions in Y. Note that, for clarity, we have extended the curves with horizontal and vertical lines, corresponding to the maximum attainable expected profit and CVaR_α for the corresponding y (which, except for the points corresponding to the solutions maximizing expected profit or CVaR_α, these are not actually Pareto efficient).

Now view that all solutions that are both Pareto efficient for a given market selection vector in Y and are on the concave envelope of these tradeoff curves are optimal solutions to (P_2) for some value of λ. In Figures 3.5 and 3.6 these solutions are indicated by the thick black curves. These curves are connected via dashed lines to complete the concave envelope. Note that the two endpoints of each of these connecting dashed lines are alternative optimal solutions corresponding to the value of λ that corresponds to the slope of the dashed line. These curves can be used to efficiently assess the tradeoff between expected profit and risk in the presence of market selection decisions. In the first example (Figure 3.5), for all optimal solutions to instances of (P_2) it is optimal to select all markets. On the other hand, in the second example (Figure 3.6), four out of six candidate market selections in Y are optimal for some values of $\lambda \in [0, 1]$.

In both of these examples, the optimal market selection that maximizes expected profit is $y = [1\ 1\ 1\ 1]$. A common incorrect assumption would be that increasing the number of units sold by operating in more markets will increase the expected profit. This is not always the case and is made particularly apparent by the choice of the fixed cost for operating in a given market. If the difference between the net revenue, $\xi_i = (r - c)\mu_i$, for a given market and its fixed cost, S_i, is small then it is likely that that specific market will never be profitable and will not increase expected profit despite increasing the number of units sold. Of course each problem is parameter specific and requires individual analysis.

Returning to the two examples depicted by Figures 3.5 and 3.6, note that the supported points may represent the entire Pareto efficient frontier (Figure 3.5) but often the supported points do not necessarily show all possible market selection solutions that comprise the Pareto efficient frontier (Figure 3.6). Specifically, as shown by Figure 3.6, there are gaps between the concave envelope of the Pareto efficient frontier and the curves given by the sorted solution.

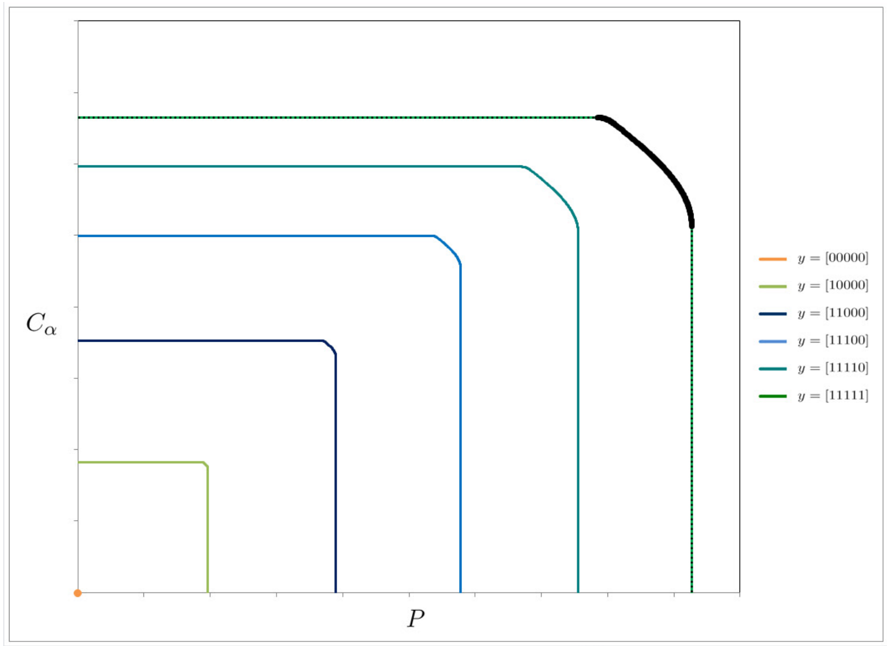

Figure 3.5: Optimal solutions for $m = 5$ for $\lambda \in [0, 1]$ (thick black curve).

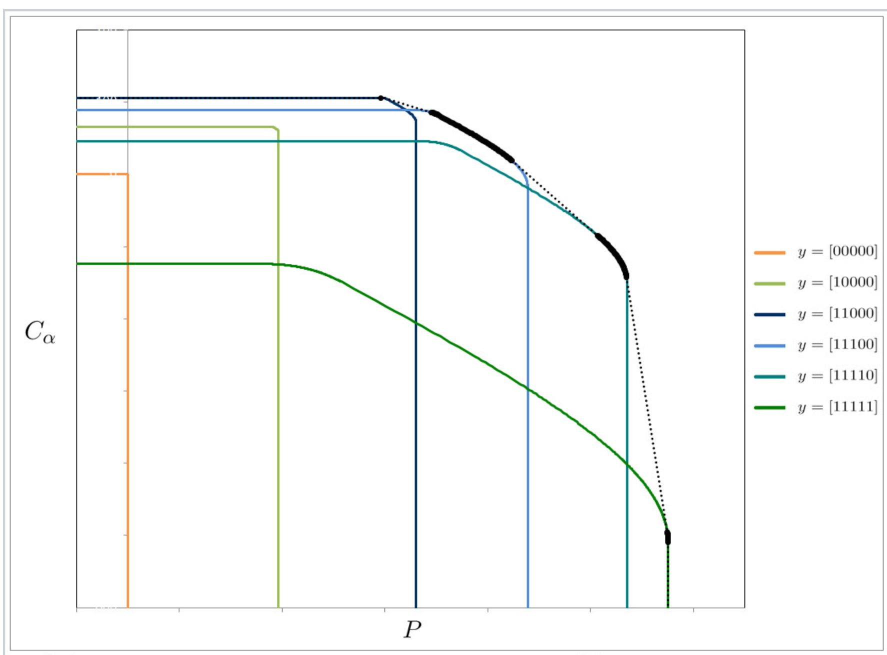

Figure 3.6: Optimal solutions for $m = 5$ for $\lambda \in [0, 1]$ (thick black curves).

3.5 Conclusion

In this chapter we studied a risk-averse newsvendor who uses Conditional Value-at-Risk to measure risk. We derived the optimal order quantity when the newsvendor optimizes a weighted sum of expected profit and risk, where we distinguished between the case in which the newsvendor can expedite orders at cost in excess of or below unit revenue. We provided insights into the behavior of the optimal order quantity as a function of the level of risk aversion. We used these results to analyze the optimization problem faced by a selective newsvendor who faces independent and normally distributed market demands. We show that the optimal market selection is one of a small and intuitive set of candidate selection vectors. This result implies that the optimal solutions to the optimization problem as a function of the relative weight of risk and expected profit can efficiently be determined.

Due to the nonconvexity of the selective newsvendor problem the techniques used in this chapter do not characterize the entire Pareto frontier of solutions that are efficient with respect to both profit and risk. We address this in the next chapter.

CHAPTER IV

Identifying the Mean vs. CVaR Pareto Efficient Frontier for a Selective Newsvendor

4.1 Introduction

As detailed in the previous three chapters, there has been a recent shift to consider the decision maker's risk preferences when analyzing inventory management problems. Eeckhoudt et al. [11] were among the first to study a risk-averse newsvendor and examine the difference between the optimal order quantity that maximizes the risk-neutral, expected profit maximizing newsvendor and the optimal order quantity for a risk-averse, expected utility maximizing newsvendor. With the advent of specific risk measures such as Value-at-Risk and Conditional Value-at-Risk (see, e.g., Artzner et al. [4], Rockafellar and Uryasev [30], and Choi and Ruszczyński [9]), there has been a shift away from using expected utility functions to analyze risk preferences towards using well defined functional risk measures.

This chapter uses the risk measure Conditional Value-at-Risk (CVaR) to assess the decision maker's risk-averseness in a *Selective Newsvendor Problem*. The selective newsvendor is a firm which, prior to the realization of demand, must determine its inventory procurement policy and choose in which markets the good will be sold. Also recall that we defined CVaR as the average profit in the left $(1 - \alpha)$ tail of the profit distribution.

In Chapter III we studied the risk-averse selective newsvendor with a CVaR criterion by maximizing a weighted sum of expected profit and CVaR to represent the tradeoff between expected profit and CVaR. For a given market selection, solving the tradeoff problem yields a Pareto efficient frontier balancing risk (CVaR) and reward (expected profit). By definition, Pareto efficiency implies that neither measure can improve without acting as a detriment to the other. In the financial engineering literature, the Pareto efficient frontier is the solution to the optimization problem set forth in the seminal work by Markowitz [23].

We showed that an optimal solution to the selective version of the tradeoff problem yields a concave envelope of the mean-CVaR Pareto efficient frontier with supported points on the set of so-called sorted solutions. The sorted solutions refers to the collection of candidate market selection solutions that contains the optimal market selection decision for both the risk-neutral selective newsvendor and the risk-averse selective newsvendor who utilizes a CVaR criterion. This set lists the candidate solutions according to an intuitive ranking scheme.

In this chapter, we study a bicriteria optimization problem and develop specific algorithms for identifying the expected profit-CVaR Pareto efficient frontier and evaluate how market selection decisions change with risk-preferences. Initially, we provide four approximations to the true efficient frontier. Then we use these approximations to identify a branch and bound procedure. The remainder of this chapter is structured as follows: Section 4.2 describes the bicriteria problem used to identify the mean-CVaR Pareto efficient frontier and Section 4.3 provides four approximations to the true Pareto efficient frontier. Section 4.4 uses these approximations to establish a branch and bound solution procedure to identify the frontier and includes some illustrative results. Section 4.6 offers concluding remarks.

4.2 A Bicriteria Selective Newsvendor Problem

In this section, we analyze the multi-objective optimization problem faced by the selective newsvendor who has the twin objectives of optimizing expected profit and Conditional Value-at-Risk (CVaR). Multi-objective optimization problems describe the class of optimization problems in which two or more competing objectives are simultaneously optimized. A recent survey of continuous nonlinear multi-objective optimization solution methods is given by Marler and Arora [24]. Recall that Taaffe et al. [36] solved the risk-neutral version of the *Selective Newsvendor Problem* by maximizing expected profit. In Chapter III we provided results for a risk-averse selective newsvendor whose sole objective was to maximize CVaR. Additionally, we analyzed a selective newsvendor who maximized a weighted sum of expected profit and CVaR.

Consider the selective newsvendor with two objectives: (i) maximizing expected profit and (ii) maximizing Conditional Value-at-Risk at level $(1-\alpha)$ associated with profit (CVaR_α). Specifically, CVaR_α is the average profit in the left $(1-\alpha)$ tail of the profit distribution. As shown in Chapter III, the order quantities that maximize CVaR_α can be smaller or larger than the risk-neutral optimal order quantity based on the problem parameters. The concept of Pareto efficiency has long been used to evaluate tradeoffs between competing objectives, that is, the selective newsvendor cannot increase its expected profit without incurring additional risk nor can it decrease its risk exposure without a willingness to accept smaller payoffs.

Using the notation provided in Chapter II, we can identify the Pareto efficient solutions for the *Selective Newsvendor Problem* by maximizing $\text{CVaR}_\alpha(Q, y)$ subject to a constraint on $P(Q, y)$ given by a lower bound, B. Define $B : [\underline{B}, \overline{B}] \rightarrow \mathbb{R}$

where $\overline{B}$ is the unconstrained maximum value of expected profit $P(Q, y)$ and $\underline{B}$ is the expected profit corresponding to the unconstrained maximum value of $C_\alpha(Q, y)$). Thus, we want to solve the following constrained optimization problem for all values of B:

$$\begin{array}{llcl} \max & C_\alpha(Q,y) & & \\ \text{subject to:} & P(Q,y) & \geq & B \\ & Q & \in & \mathbb{R} \\ & y & \in & \{0,1\}^m \end{array} \tag{C(B)}$$

where

$$C_\alpha(Q, y) = \max_{\theta \in \mathbb{R}} \left\{ \theta - \frac{1}{1-\alpha} E[(\theta - \pi(Q; D_y))^+] \right\}$$

and

$$P(Q, y) = (r - v)\mu_y - (c - v)Q - (e - v) \int_Q^\infty (x - Q) dF_y(x) - S_y.$$

The Pareto frontier is then given by the optimal value function of this class of optimization problems, which we will denote by $\mathcal{F}(B)$.

Recall that for a fixed y, $P(Q, y)$ is concave in Q. Therefore, solving (C(B)) for a fixed market selection is straightforward. Simply use a search algorithm to find a Q that makes $P(Q, y) = B$ and compute the corresponding $\text{CVaR}_\alpha(Q, y)$. Repeating this procedure for all possible B would result in a tradeoff curve for a given y. Clearly enumerating all possible market combinations, finding their tradeoff curves, and taking the upper envelope of said curves is inefficient as the number of potential solutions increases exponentially with each additional market.

As discussed in Chapter II, a selective newsvendor who maximizes $P(Q, y)$ exhibits a special solution structure. That is, there exists an optimal market selection solution with the property that, if market ℓ is selected, markets $1, ..., \ell - 1$ are selected as well, provided the markets are indexed in nonincreasing order of ξ_i where

$\xi_i = \frac{(r-c)\mu_i - S_i}{\sigma_i^2}$, $(i = 1, ..., m)$. The set of candidate optimal solutions is described as the set Y.

This solution structure holds for a selective newsvendor who maximizes $C_\alpha(Q, y)$ as well as a selective newsvendor who maximizes a weighted sum of $C_\alpha(Q, y)$ and $P(Q, y)$(see Chapter III). Unfortunately, this property does not extend to the problem (C(B)) which makes identifying the Pareto frontier $\mathcal{F}(B)$ challenging. The remainder of this chapter focuses on methods to efficiently identify $\mathcal{F}(B)$.

4.3 Approximating the Pareto Efficient Frontier

The non-linear integer constrained optimization problem described by (C(B)) that yields the Pareto efficient frontier $\mathcal{F}(B)$ is difficult to solve. The goal of this section and the next are to provide methods for approximating and identifying the expected profit-CVaR Pareto efficient frontier over all possible y values.

We begin first by describing four methods which approximate the true Pareto efficient frontier. The approximations vary in difficulty and proximity to the actual frontier. The first approximation creates a lower bound while the latter three create upper bounds of varying tightness. Eventually we use these bounds to describe a branch and bound type approach to solving (C(B)).

To facilitate our analysis, we provide a simple three market example where the markets are sorted such that $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2}$. The three potential markets lead to eight possible combinations. For each market selection possibility, we computed the Pareto efficient frontier and the results are shown in Figure 4.1. As you can see, the Pareto efficient frontier is comprised of the sorted solutions $y = [1\ 0\ 0], y = [1\ 1\ 0], y = [1\ 1\ 1]$ and the non-sorted solution $y = [1\ 0\ 1]$. This example will be used throughout the

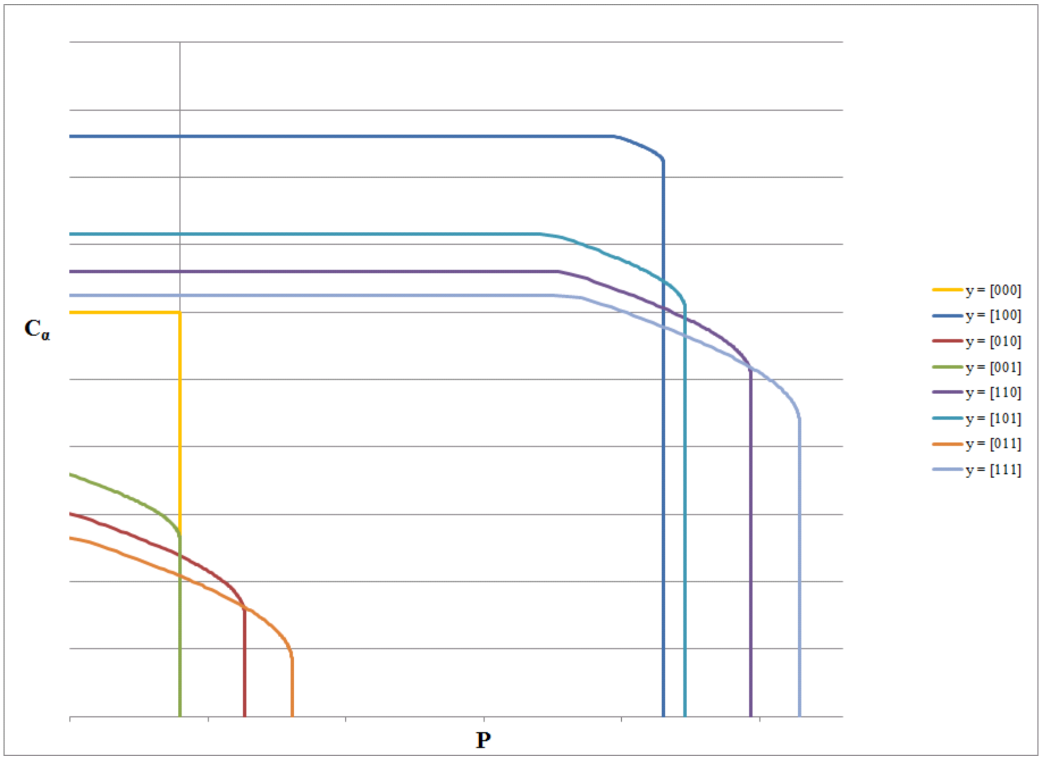

Figure 4.1: Three market example with $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2}$.

section.

4.3.1 Set of the Sorted Solutions

The previously defined set Y contains a set of indexed market selection solutions such that if market ℓ is included, markets $1, \ldots, \ell - 1$ are also included. When these candidate solutions are sorted in nondecreasing order of $\frac{\xi_i}{\sigma_i^2}$ this set includes optimal market selections that maximize $P(Q, y)$ as well as $C_\alpha(Q, y)$. We can use these optimums to define a range of B values for which $\mathcal{F}(B)$ exists. The lower bound, denoted $\underline{B}$ will correspond to the expected profit associated with the maximum value of $C_\alpha(Q, y)$ and the upper bound $\overline{B}$ is equivalent to the maximum of $P(Q, y)$. Thus $\mathcal{F}(B)$ will be bookended by candidate solutions found in Y.

We introduce our first approximation as the outer envelope of the Pareto efficient

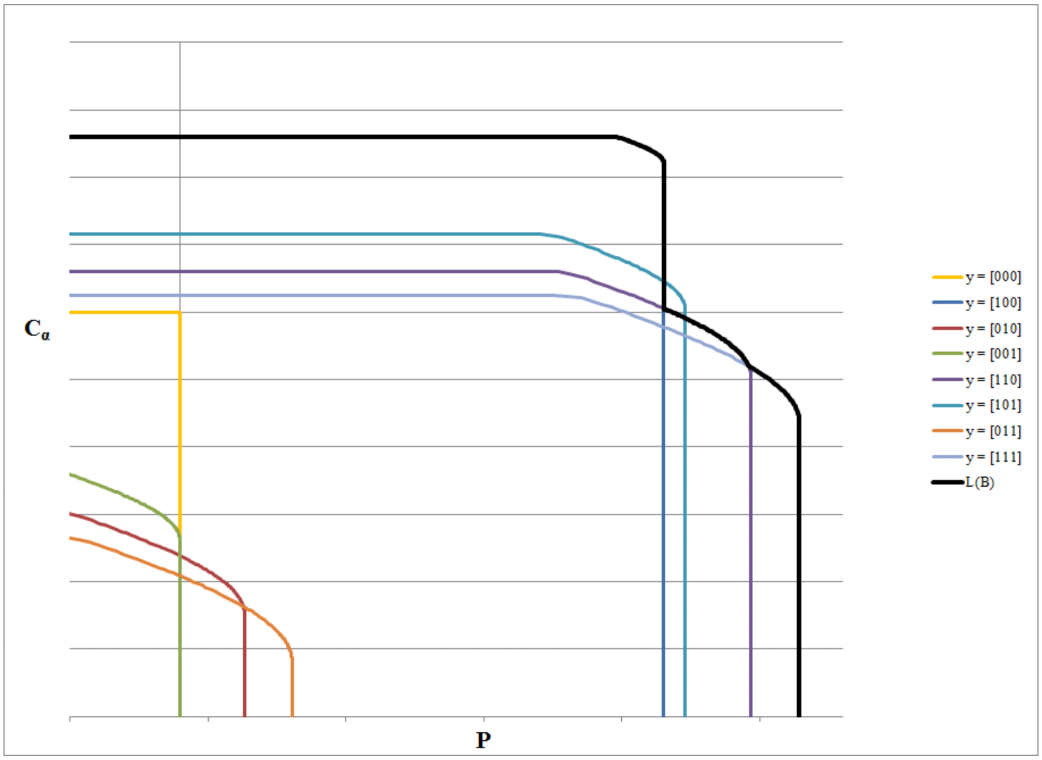

Figure 4.2: Lower bound, $\mathcal{L}(B)$, for a three market example.

frontiers from the $m + 1$ sorted solutions. Calling this approximation $\mathcal{L}(B)$, we know it is equivalent to the true efficient frontier at $\underline{B}$ and $\overline{B}$. Naturally, the main problem associated with using the sorted solutions to approximate the frontier is the possibility of overlooking any non-sorted solutions that are part of $\mathcal{F}(B)$.

In our three market example, the markets under consideration would be $y = [0\ 0\ 0], y = [1\ 0\ 0], y = [1\ 1\ 0]$, and $y = [1\ 1\ 1]$. As seen in Figure 4.2, the concave envelope represented by the thick black line is constructed from three of the potential candidate solutions. The market selection $y = [1\ 0\ 1]$ is part of $\mathcal{F}(B)$ but is overlooked.

4.3.2 Continuous Relaxation of the Binary Constraint

In the next three sections we consider relaxations of the problem (C(B)). Since we are solving a maximization problem, the relaxations will all yield upper bounds for the true Pareto efficient frontier. The simplest and most straightforward is to relax the requirement that participation in a given market be binary.

$$\begin{array}{lrcl} \max & C_\alpha(Q,y) & & \\ \text{subject to:} & P(Q,y) & \geq & B \\ & Q & \in & \mathbb{R} \\ & y & \in & [0,1]^m \end{array} \qquad (\mathrm{R}^1(B))$$

From the original description of the *Selective Newsvendor Problem* in Chapter II we know that $D_y \sim n(\mu_y, \sigma_y^2)$ with $\mu_y = \mu^\top y$ and $\sigma_y^2 = y^\top \Sigma y$. The relaxation above is a convex optimization problem (maximizing a concave function over a convex feasible region) for a given value of B. The class of optimization problems solved over all possible values of B results in an upper bound which we denote $\mathcal{U}^1(B)$. Solving this relaxation for the three market example described in the previous section yields the upper bound depicted in Figure 4.3.

4.3.3 Lagrange Relaxation

Consider a Lagrangian relaxation of the expected profit constraint in (C(B)). Denote the optimal value function of this problem by $\mathcal{U}^2(B)$. It is well-known that the function $\mathcal{U}^2(B)$ is the concave envelope to $\mathcal{F}(B)$ and therefore can be equivalently identified by solving a selective newsvendor problem with a weighted expected profit-CVaR objective function as described by Chapter III.

For a fixed y, the objective function is concave in Q yielding an expression for

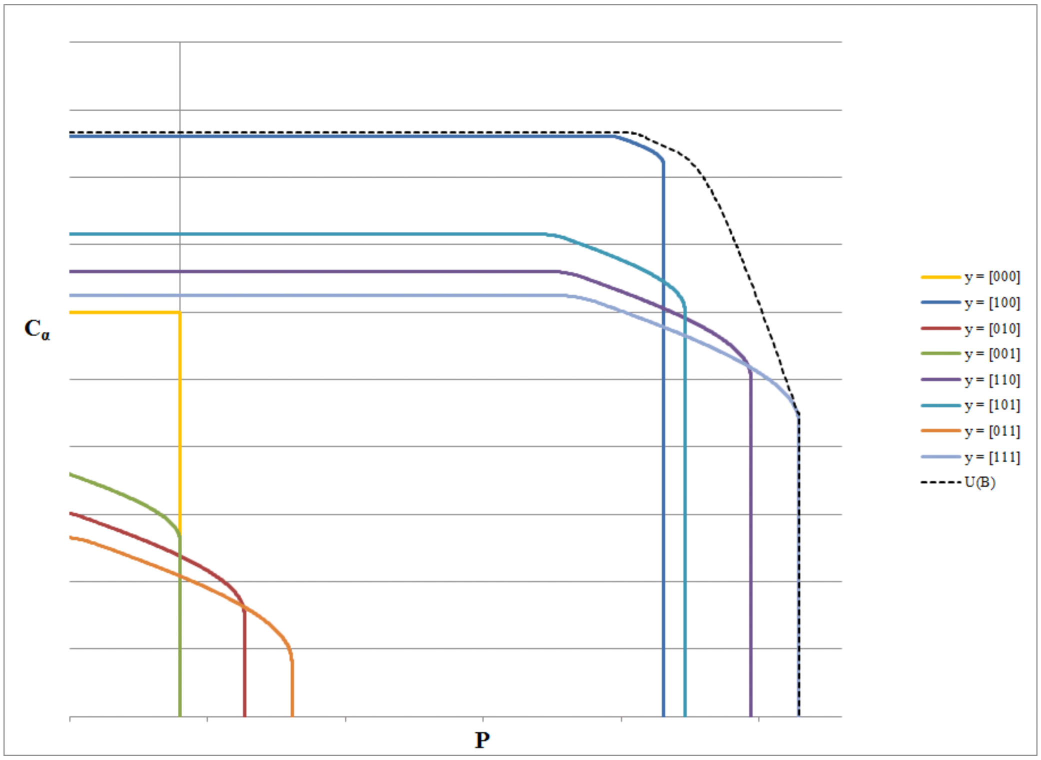

Figure 4.3: Upper bound, $\mathcal{U}^1(B)$, for a three market example.

Q as a function of λ. Using these expressions, we showed that the optimal market selection for a selective newsvendor solving the above problem is one of the small and intuitive set defined by Y. Since one of the optimal market selections is simply $y = \{\mathbf{0}\}$, we only have m nontrivial optimization problems to solve and the integral sorted solutions create a concave envelope, $\mathcal{U}^2(B)$, for $\mathcal{F}(B)$.

This approximation yielded by solving for $\mathcal{U}^2(B)$ is useful because it is straightforward to compute and can be solved quickly and efficiently. Unfortunately, it has some significant shortfalls. As with the lower bound in the previous section, $\mathcal{U}^2(B)$ will of course overlook any non-sorted solutions that are on the efficient frontier. Additionally, there is no guarantee that $\mathcal{U}^1(B)$ will identify all of the sorted solutions that are on the efficient frontier.

Solving this problem for our simple three market example highlights this issue.

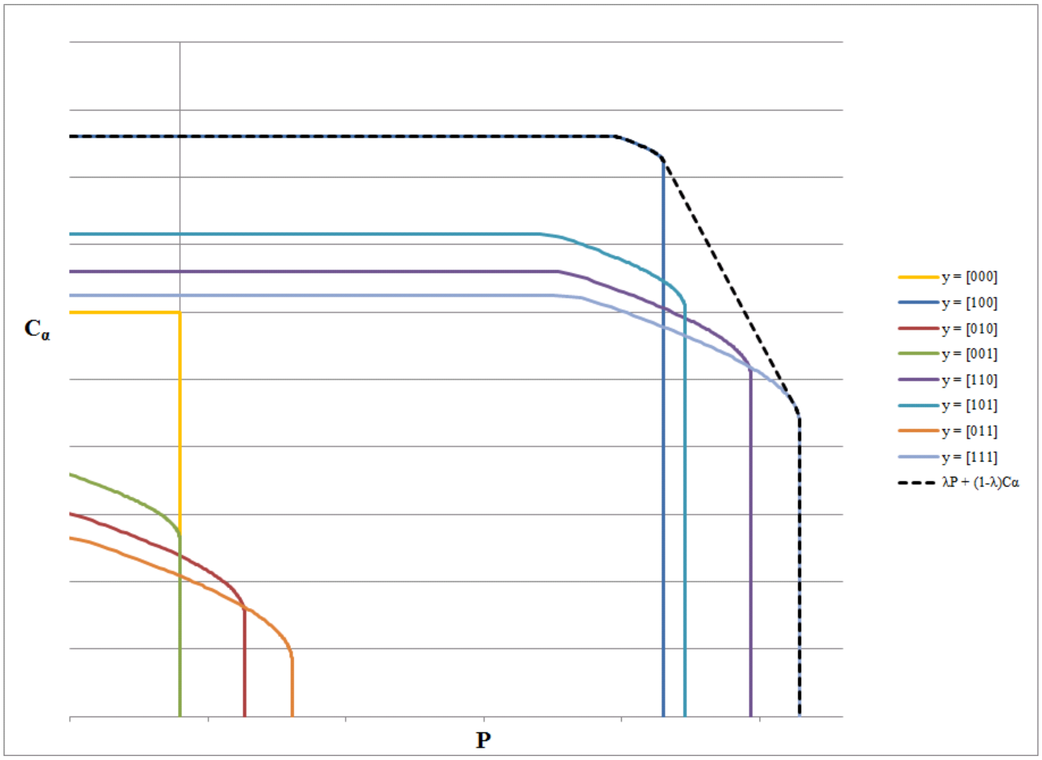

Figure 4.4: Upper bound, $\mathcal{U}^2(B)$, for a three market example.

You can see from Figure 4.4 that the concave envelope provided by $\mathcal{U}^2(B)$ fails to identify the sorted solution $y = [110]$ in addition to overlooking the non-sorted solution $y = [101]$.

4.3.4 An Alternative Continuous Relaxation of the Binary Constraint

For this approximation, we again relax the integer constraint on market selection. Recall that in the original description of the model we stated that the variance associated with D_y is $\sigma_y^2 = y^\top \Sigma y = \mathbf{1}^\top \Sigma y$ with the second equality following from the fact that y is binary. In this model, we relax the binary requirement on y while maintaining that the variance is $\sigma_y^2 = \mathbf{1}^\top \Sigma y$. Since we are explicitly utilizing the assumption of normally and independently distributed market demands, $D_y \sim n(\mu_y, \sigma_y^2)$, it will be useful to preface our analysis with a reparameterization of the problem similar to that used in Chapter III. Define the order quantity as a function of the variables y and w such that

$$Q_y(w) = \mu_y + H_\alpha(w)\sigma_y$$

where

$$H_\alpha(w) = \begin{cases} \Phi^{-1}(w) & \text{for } e \leq r \\ \left(\frac{r-v}{e-v}\right)\Phi^{-1}(w) + \left(\frac{e-r}{e-v}\right)\Phi^{-1}(w+\alpha) & \text{for } e > r. \end{cases}$$

As shown in the previous chapter, the case when unit expediting cost exceeds unit revenue is considerably more complicated. Incorporating this expression into the expected profit function yields

$$\tilde{P}_\alpha(w, y) \equiv P(Q_{y,\alpha}(w), y) = \xi_y - K_\alpha(w)\sigma_y$$

where

$$K_\alpha(w) = (c-v)H_\alpha(w) + (e-v)\Lambda(\Phi(H_\alpha)).$$

and Φ is the c.d.f. of the standard normal distribution and $\Lambda(w) = \int_{\Phi^{-1}(w)}^{\infty}(z - \Phi^{-1}(w))d\Phi(z)$. We also reparametize the expression for CVaR_α such that

$$\tilde{C}_\alpha(w, y) \equiv P(Q_{y,\alpha}(w), y) = \xi_y - L_\alpha(w)\sigma_y$$

where

$$L_\alpha(w) = \begin{cases} (c-v)H_\alpha(w) + & \\ (\frac{1}{1-\alpha} - 1)[(r-e)\Phi^{-1}(1-\alpha) + (e-v)\Phi^{-1}(w)] + & \\ \frac{1}{1-\alpha}[(r-e)\Lambda(1-\alpha) + (e-v)\Lambda(w)] & \text{if } w \leq 1-\alpha \\ (c-v)H_\alpha(w) + & \\ (\frac{1}{1-\alpha} - 1)(r-v)\Phi^{-1}(1-\alpha) + \frac{1}{1-\alpha}(r-v)\Lambda(1-\alpha) & \text{if } w > 1-\alpha \end{cases}$$

for $e \leq r$, and

$$L_\alpha(w) = (c-v)H_\alpha + (\frac{1}{1-\alpha} - 1)(r-v)\Phi^{-1}(w) + \frac{1}{1-\alpha}[(r-v)\Lambda(w) + (e-r)\Lambda(w+\alpha)]$$

for $e > r$.

Therefore, the proposed relaxation is

$$\begin{array}{rrcl} \max & \xi_y - L_\alpha(w)\sigma_y & & \\ \text{subject to:} & \xi_y - K_\alpha(w)\sigma_y & \geq & B \\ & w & \in & (0, \bar{w}) \\ & y & \in & [0,1]^m \end{array} \qquad (\mathrm{R}^2(B))$$

where $\bar{w} = 1$ if $e \leq r$ and $\bar{w} = 1-\alpha$ if $e > r$. We will further simplify this optimization problem by explicitly characterizing a set of potentially optimal solutions.

Lemma 4.1. *There exists an optimal solution to* $(\mathrm{R}(B))$ *in the set*

$$\bar{Y} = \{\mathbf{0}\} \cup \left\{ y^{(\ell)}(t) : t \in (0,1], \ell = 1, \dots, m \right\}$$

where

$$y_i^{(\ell)}(t) = \begin{cases} 1 & \text{for } i = 1, \dots, \ell - 1 \\ t & \text{for } i = \ell \\ 0 & \text{for } i = \ell + 1, \dots, m \end{cases}$$

for some $\ell = 1, ..., m$ *and some* $t \in (0, 1]$.

Proof. Introduce a nonnegative decision variable $z = \sigma_y^2$ to the problem $(\mathrm{R}(B))$. Consider the following optimization problem for fixed z and w:

$$\begin{array}{rrcl} \max & \xi_y - L_\alpha(w)\sigma_y & & \\ \text{subject to:} & \xi_y - K_\alpha(w)\sigma_y & \geq & B \\ & \sigma_y^2 & = & z \\ & y & \in & [0,1]^m. \end{array}$$

Since the second term in the objective function is a constant for fixed w and z, it follows that $\xi_y - K_\alpha(w)\sqrt{z} \geq B$ is simply a feasibility constraint. We can solve the optimization problem without it and then verify whether or not the corresponding solution satisfies the constraint, that is the solution will either be optimal or infeasible. The remaining problem

$$\begin{array}{rrcl} \max & \xi_y - L_\alpha(w)\sigma_y & & \\ \text{subject to:} & \sigma_y^2 & = & z \\ & y & \in & [0,1]^m \end{array}$$

is a continuous knapsack problem. The solution algorithm requires ordering the markets in nonincreasing values of ξ_i and finding the index j such that

$$\sum_{i=1}^{j-1} \sigma_i^2 < z \quad \text{and} \quad \sum_{i=1}^{j} \sigma_i^2 \geq z.$$

Defining $\bar{z} = z - \sum_{i=1}^{j-1} \sigma_i^2$, the optimal solution y is given by

$$y_i = \begin{cases} 1 & \text{for } i = 1, ..., j-1 \\ \frac{\bar{z}}{\sigma_j^2} & \text{for } i = j \\ 0 & \text{for } i = j+1, ..., n \end{cases}$$

which yields the desired result. $\square$

Note that the binary solutions $y_i^{(\ell)} \in Y$ defined earlier now correspond to $y^{(\ell)}(1), \ell = 1, ..., m$. We now reformulate (R($B$)) by explicitly restricting ourselves to $y \in \bar{Y}$. In fact, we will replace (R(B)) by a collection of m optimization problems (one for each value of ℓ) in the decision variables w and t only. For convenience, we define the following functions

$$\xi_\ell(t) = \sum_{i=1}^{\ell-1} \xi_i + \xi_\ell t \qquad \text{for } t \in (0,1]$$

$$s_\ell(t) = \sqrt{\sum_{i=1}^{\ell-1} \sigma_i^2 + \sigma_\ell^2 t} \qquad \text{for } t \in (0,1]$$

and consider the following class of optimization problems:

$$\begin{array}{rrcl} \max & \xi_\ell(t) - L_\alpha(w)s_\ell(t) & & \\ \text{subject to:} & \xi_\ell(t) - K_\alpha(w)s_\ell(t) & \geq & B \\ & w & \in & (0,\bar{w}) \\ & t & \in & (0,1]. \end{array} \tag{$R_\ell(B)$}$$

We will use the KKT conditions for ($\mathrm{R}_\ell(B)$) to characterize candidate solutions to this problem. Noting that the constraints $w \in (0,\bar{w})$ and $t > 0$ cannot be binding, the KKT conditions for this problem can be written as:

$$\begin{aligned} (\eta+1)\xi_\ell'(t) - (L_\alpha(w) + \eta K_\alpha(w))s_\ell'(t) &= \delta \\ (L_\alpha'(w) + \eta K_\alpha'(w))s_\ell(t) &= 0 \\ \xi_\ell(t) - \eta K_\alpha(w)s_\ell(t) &\geq B \\ t &\in (0,1] \\ w &\in (0,\bar{w}) \\ \eta[B - \xi_\ell(t) + K_\alpha(w)s_\ell(t)] &= 0 \\ \delta(t-1) &= 0 \\ \delta, \eta &\geq 0. \end{aligned}$$

Since $s_\ell(t) > 0$ for all $t \in (0,1]$ we can use the second condition to write η as a function of w:

$$\eta(w) = -\frac{L_\alpha'(w)}{K_\alpha'(w)}.$$

Observe that

$$K_\alpha'(w) = (e-v)\left[\Phi(H_\alpha(w)) - \rho\right] H_\alpha'(w)$$

and

$$L_\alpha'(w) = (e-v)\left(\min\{\frac{w}{1-\alpha}, 1\} - \rho\right) H_\alpha'(w).$$

It is easy to see that $H_\alpha(w)' > 0$ for all w. Therefore, provided that $w \neq H_\alpha^{-1}(\Phi^{-1}(\rho))$, we have

$$\eta(w) = \frac{\min\{\frac{w}{1-\alpha}, 1\} - \rho}{\rho - \Phi(H_\alpha(w))}.$$

We can now identify a set of solutions that contains all Pareto efficient solutions to the bicriteria optimization problem ($\mathrm{R}_\ell(B)$) for some value of B. In particular, consider a fixed value of $\ell = 1, ..., m$ and $w \in (0, \bar{w})$ such that

$$(1-\alpha)\rho \leq w \leq H_\alpha^{-1}(\Phi^{-1}(\rho)) \quad \text{or} \quad H_\alpha^{-1}(\Phi^{-1}(\rho)) < w \leq (1-\alpha)\rho.$$

Incorporating $\eta(w)$ into the third KKT condition yields an expression of δ as a function of w and t:

$$\delta(w,t) = (\eta(w)+1)\eta_\ell'(t) - (L_\alpha(w) + \eta(w)K_\alpha(w))s_\ell'(t).$$

We then distinguish between integral and fractional KKT solutions:

- *Integral selection vector:* From the condition $\delta(t-1) = 0$ we know that $\delta > 0$ implies $t = 1$. Hence we obtain an integral KKT solution if and only if $\delta(w,1) > 0$.

- *Fractional selection vector:* There is at most one value of $t \in (0,1]$ that yields $\delta(w,t) = 0$ for a given w. Solving for t yields a fractional KKT solution. Specifically, there exists no value of t for which $\delta(w,t) = 0$ if $\eta(w)K_\alpha(w) < L_\alpha(w)$. Otherwise, the following value of t sets $\delta(w,t) = 0$:
$$t = \frac{\left([\eta(w)K_\alpha(w) - L_\alpha(w)]\frac{\sigma_\ell^2}{2\eta(w)\xi_\ell}\right)^2 - \sum_{i=1}^{\ell-1}\sigma_i^2}{\sigma_\ell^2}$$
which yields a KKT solution provided that $t \in (0,1]$.

In either case, the point (P, C) given by

$$P = \xi_\ell(t) - K_\alpha(w)s_\ell(t)$$

$$C = \xi_\ell(t) - L_\alpha(w)s_\ell(t)$$

is potentially on the Pareto efficient frontier. In addition, for a fixed value of $\ell = 1, ..., m$ we obtain a possibly Pareto efficient solution with $w = H_\alpha^{-1}(\Phi^{-1}(\rho))$. Note that in that case, the feasible region of ($\mathrm{R}_\ell(B)$) will often be a singleton so that the KKT conditions do not apply. We also consider the candidate solution $y = \mathbf{0}$.

Once we have all potential Pareto efficient candidate solutions, we must take the concave envelope of them to find an upper bound to our original problem, ($\mathrm{C}(B)$). Since we are maximizing, our relaxation is by definition an upper bound. Let's call this upper bound, $\mathcal{U}^2(B)$. Here are the steps required to convert these points into a frontier:

1. Considering all n candidate points, sort the coordinates (P, C) in decreasing order according to the P values such that

$$(P_1, C_1) \geq (P_2, C_2) \geq \ldots \geq (P_n, C_n) \quad \text{where } P_1 \geq P_2 \geq \ldots P_n.$$

2. Starting with the second sorted point, set

$$C_i' = \begin{cases} C_i & \text{if } C_i > C_{i-1} \\ C_{i-1} & \text{if } C_i \leq C_{i-1} \end{cases}$$

3. The upper bound $\mathcal{U}^3(B)$ will consist of the points (P_1, C_1) and (P_i, C_i') for every $i = 2, ..., n$.

Figure 4.5 graphically depicts what a set of potential KKT conditions could look like and Figure 4.6 provides the corresponding upper bound associated with those points.

Obviously, calculating the potential KKT candidate solutions is computationally expensive. However, this relaxation identifies all sorted solutions that exist on the Pareto frontier as well as fractional candidate solutions which can provide insight

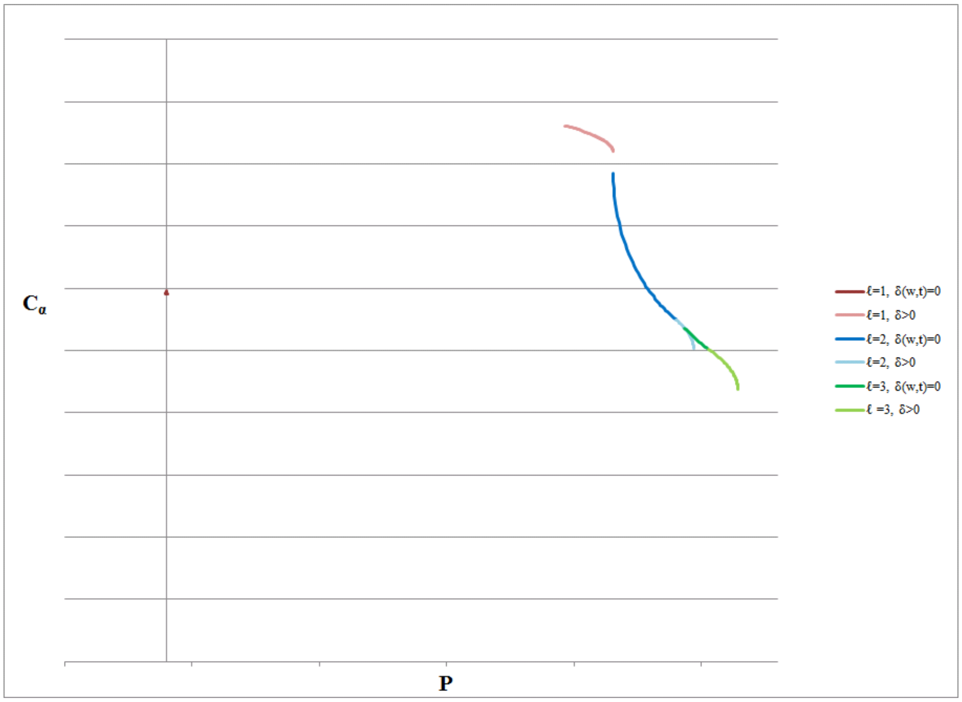

Figure 4.5: Potential Pareto candidate solutions for a three market example.

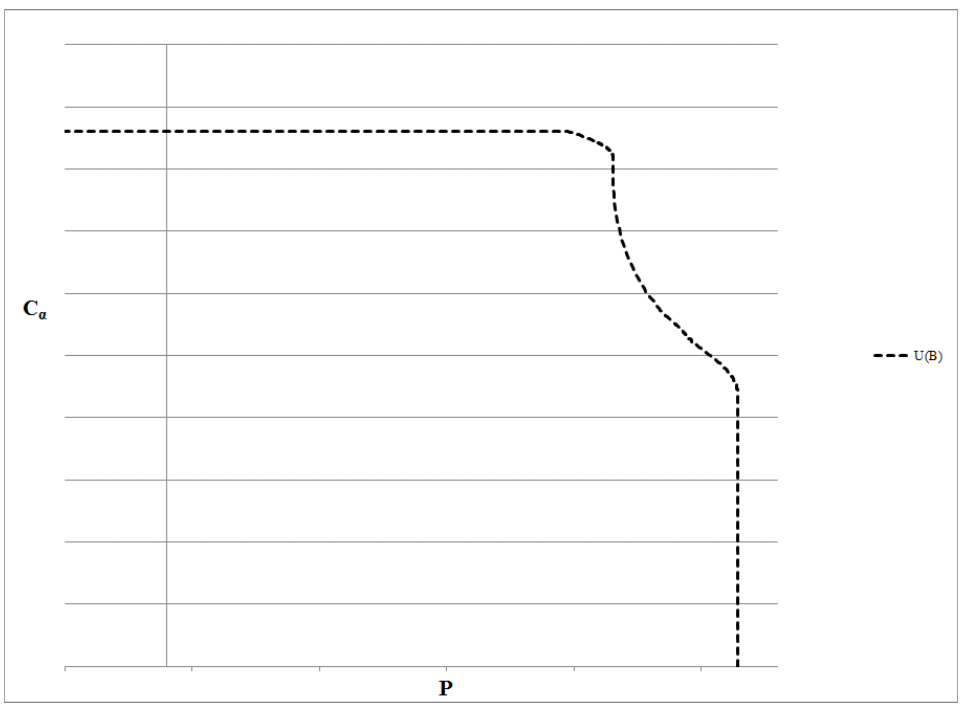

Figure 4.6: The concave envelope $\mathcal{U}^3(B)$ for a three market example.

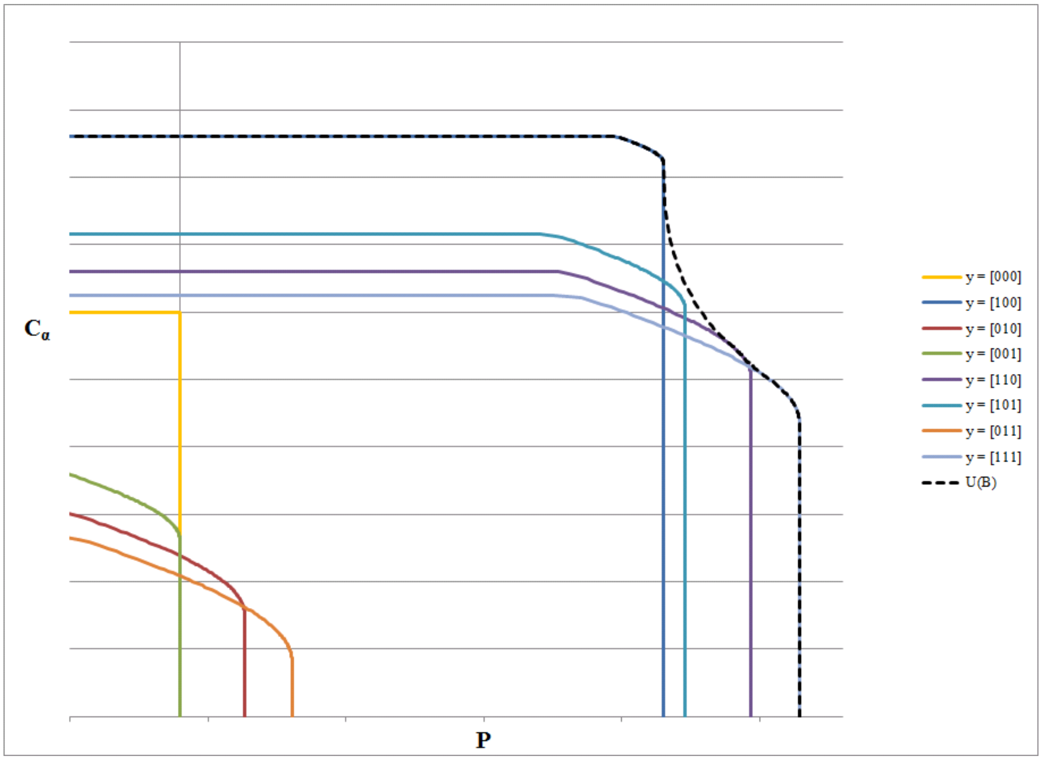

Figure 4.7: Upper bound, $\mathcal{U}^3(B)$, for a three market example.

into which other market selections to include. Returning to the original three market example introduced at the beginning of this section, we create the corresponding upper bound using the KKT conditions to find candidate Pareto optimal solutions. This is depicted in Figure 4.7. In this specific example, the KKT solutions identified the three sorted solutions that comprise the true efficient frontier and yielded fractional results for $\ell = 2$ and $\ell = 3$.

4.3.5 Relationship Between the Approximations

In this section, we provided three approximations to the true Pareto efficient frontier, $\mathcal{F}(B)$, each with its own merits. In the next section we will use these approximations to create a branch-and-bound scheme to find the true frontier. Prior to that analysis, we introduce the following lemma identifying the relationship of the approximations to the Pareto efficient frontier as well as to each other.

Lemma 4.2. *The value functions of the three approximations to* $\mathrm{C}(B)$ satisfy

$$\mathcal{U}^1(B) \geq \mathcal{U}^2(B) \geq \mathcal{U}^3(B) \geq \mathcal{F}(B) \geq \mathcal{L}(B)$$

for all B.

Proof. As shown in Chapter III, $\mathcal{U}^2(B)$ is the concave envelope of $\mathcal{F}(B)$. The convexity of $\mathrm{R}^1(B)$ implies that $\mathcal{U}^1(B)$ is concave, the first inequality follows immediately. Now consider the weighted sum optimization problem with $\sigma_y^2 = \mathbf{1}^\top \Sigma y$. Since a binary optimal solution to its continuous relaxation exists so that relaxing the binary constraint does not change the optimal solution. The second inequality then immediately follows from the fact that the continuous relaxation of the weighted sum optimization problem is precisely the Lagrange relaxation of $\mathrm{R}^2(B)$ with respect to the expected profit constraint. $\mathcal{U}^1(B), \mathcal{U}^2(B)$ and $\mathcal{U}^3(B)$ are solutions to relaxations thus the third inequality. The final inequality follows from the fact that $\mathcal{L}(B)$ is at best equivalent to $\mathcal{F}(B)$ when the true efficient frontier is comprised only of sorted solutions. $\square$

Graphical depictions of these approximations for our three problem example are shown by Figure 4.8 and Figure 4.9 .

The existence of these approximations, three that provide upper bounds of varying tightness and one lower bound, directed our research in the next section. We utilize these bounds to develop a branch and bound procedure, a well-known approach to constrained optimization problems, to find the true mean-CVaR Pareto efficient frontier.

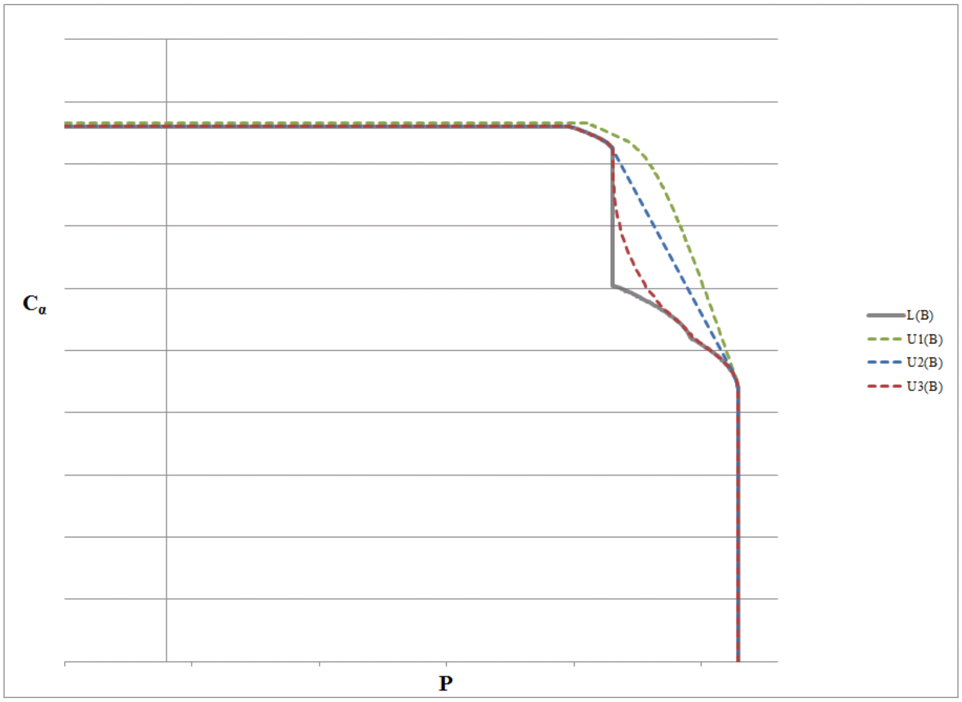

Figure 4.8: Four approximations to $\mathcal{F}(B)$ for a three market example.

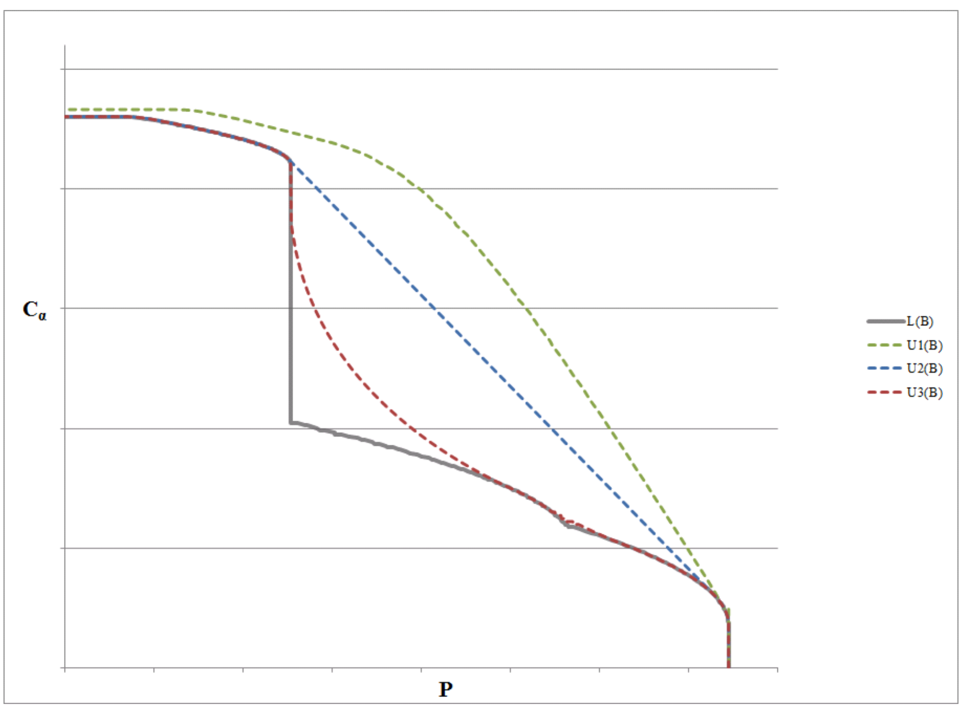

Figure 4.9: Close-up of four approximations to $\mathcal{F}(B)$ for a three market example.

4.4 Identifying the Pareto Efficient Frontier

A widely-used approach to solving constrained optimization problems is the branch and bound technique (see e.g., Land and Doig [19], Little et al. [22], and Lawler and Wood [21]). The idea behind branch and bound is a cleverly structured enumeration procedure that searches the solution space for the optimum while only examining a fraction of all feasible solutions. Each iteration of branch and bound requires the selection of a node, determining the bounds at that node, and then branching the problem into smaller and smaller subsets. The initial node, usually called the root node, is simply the set of all feasible solutions. The bounding function is determined and compared to the current best solution. If it can be established that the subset of feasible solutions cannot contain the optimal solution, the whole subset is discarded. Otherwise, the branching procedure is employed and the nodes created are added to the pool of unexplored nodes. From Section 4.3, we know that we have upper and lower bounding functions on $\mathcal{F}(B)$. In this section, we will use these bounds to strategically determine the true efficient frontier for (C(B)). First, we introduce some convenient notation.

4.4.1 Notation

The Pareto frontier, $\mathcal{F}(B)$, is the optimal value function of the problem $(C(B))$ and is defined $\mathcal{F}(B) : [\underline{B}, \overline{B}] \to \mathbb{R}$ where $\overline{B}$ is the unconstrained maximum value of expected profit $P(Q, y)$ and $\underline{B}$ is the expected profit corresponding to the unconstrained maximum value of $C_\alpha(Q, y)$). For a fixed $y \in \{0, 1\}^m$ we have an efficient way of determining the corresponding frontier, $\mathcal{F}_y(B)$. Thus,

$$\mathcal{F}(B) = \max_{y \in \{0,1\}^m} \mathcal{F}_y(B).$$

At any point during the branch and bound algorithm we require an incumbent (current best) solution characterized by a number of binary market selection vectors. In Section 4.3 we defined such a solution as $\mathcal{L}(B)$ which was the outer envelope of the market selection vectors in Y. Specifically, when Y is the set of all sorted solutions and $\mathcal{L}(B)$ is defined as the collection of sorted solutions that lay on the efficient frontier.

For a given node in the branch and bound algorithm, we denote the set of potential market selection solutions used to determine the current best solution as Ω. This set of vectors provides the following lower bound for $\mathcal{F}(B)$:

$$\mathcal{L}_\Omega(B) = \max_{y \in \Omega} \mathcal{F}_y(B).$$

Notice that each node in the branch and bound procedure has two disjoint subsets of $\{1, ..., m\}$: I_0 and I_1. The former contains the markets that are forced out of the solution while the latter contains the markets that are forced into the solution in the current branch and bound subtree. For example $I_0 = 2$ implies that market 2 will never be in the optimal solution for that subset. Alternatively $I_1 = 2$ means that market 2 will always be in the optimal solution for that branch.

Mathematically speaking, the current subtree considers an optimization problem of the form:

$$\begin{array}{lrcll}
\max & C_\alpha(Q, y) & & & \\
\text{subject to:} & P(Q, y) & \geq & B & \\
 & Q & \in & \mathbb{R} & \\
 & y_i & = & 0 & \text{for } i \in I_0 \\
 & y_i & = & 1 & \text{for } i \in I_1 \\
 & y & \in & \{0, 1\}^m. &
\end{array} \qquad (C_{I_0, I_1}(B))$$

Assume that, for given subsets I_0 and I_1, we can find an upper bound $\mathcal{U}_{I_0,I_1} : [\underline{B}, \overline{B}] \to \mathbb{R} \cup \{-\infty\}$. Section 4.3 described solution methods for finding three such upper bounds which will be the primary focus of this dissertation.

For each iteration of the problem, it is important to construct Ω in a meaningful way. Since we have shown that the two approximations we consider for upper bounds in this chapter have a sorting element, we narrow our attention to a set of "restricted sorted solutions." By this, we mean that although I_0 and I_1 will fix certain markets as $y_i = 0$ or $y_i = 1$ respectively, we add the remaining potential market selections in nondecreasing order of ξ_i. Consider for example a simple three market scenario in which the markets are ranked $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2}$. If $I_0 = 2$, then the y vectors that should be included in Ω are $y = [0\ 0\ 0], y = [1\ 0\ 0]$, and $y = [1\ 0\ 1]$. Alternatively, if $I_1 = 2$, then the y vectors that should be included in Ω are $y = [0\ 1\ 0], y = [1\ 1\ 0]$, and $y = [1\ 1\ 1]$. We define these "restricted sorted solutions" by $Y(I_0, I_1)$.

4.4.2 Branch-and-Bound Algorithms

In this section we detail a branch and bound algorithm that can be used to find the expected profit-CVaR Pareto efficient frontier, $\mathcal{F}(B)$. Analogous to a standard integer linear programming problem in which the branching procedure involves setting certain variables to either 0 or 1, we make our decision based on whether or not to operate in a given market. Thus, we can branch the problem into two separate subsets, one with $y_i = 0$ and the other with $y_i = 1$. Using our notation from the previous section, when $y_i = 0$, $I_0 = \{i\}$ and similarly when $y_i = 1$, $I_1 = \{i\}$.

We will denote each node in the algorithm as $N(I_0, I_1)$, specifically identifying which markets can and cannot be part of the solution. The initialization of the procedure begins with I_0 and I_1 both as empty sets. Naturally, the root node is

$N(I_0 = \emptyset, I_1 = \emptyset)$. For each node there exists an incumbent (current best) solution denoted Ω. Initially, we set $\Omega = Y(I_0 = \emptyset, I_1 = \emptyset)$ so that $\Omega = Y$ the original set of candidate solutions. Our initial lower bound, $\mathcal{L}_\Omega(B)$ is the set of the outermost sorted solutions.

After an iteration of the branch and bound procedure, Ω will be updated by adding the "restricted sorted solutions." This allows us to strategically include non-sorted solutions and determine if any non-sorted solution can improve the existing lower bound. Clearly, the goal is not to eventually have $\Omega = |m + 1|$ so we must systematically eliminate nodes and their descendants.

We require a method to prune some of the nodes so that we do not have to enumerate every possible market combination to find $\mathcal{F}(B)$. We do this by finding a corresponding upper bound $\mathcal{U}_{I_0,I_1}(B)$. The upper bound could be any of the potential relaxations outlined in Section 4.3. If $\mathcal{U}_{I_0,I_1}(B) \leq \mathcal{L}_\Omega(B)$ then we can prune that node and all its descendants as we know that it will never yield a solution better than our current best solution.

However, if there is room for improvement, that is if $\mathcal{U}_{I_0,I_1}(B) > \mathcal{L}_\Omega(B)$, then we divide the node under consideration into two separate nodes each with their own solution space, $N(I_0 \cup \{i\}, I_1)$ and $N(I_0, I_1 \cup \{i\})$. The choice of market to branch on will likely depend on the set of values of $B \in [\underline{B}, \overline{B}]$ for which the upper bound at a given node exceeds the current lower bound. We continue until all nodes are explored. A generic description of this approach is provided by Algorithm 1.

The branch and bound approach described in Algorithm 1 is straightforward and relatively simple to implement, but does have its drawbacks. First, the decision of which upper bound to utilize drastically affects the procedure. Refer again to Figure 4.8 which can be interpreted as the result of analyzing the root node for the various

Algorithm 1 Branch and Bound I

Step 0. Set $I_0 = I_1 = \emptyset$, $\Omega = \emptyset$.

Step 1. Choose an unexplored node, $N(I_0, I_1)$.

Step 2. Define the incumbent, $\Omega = \Omega \cup Y(I_0, I_1)$.

Step 3. Determine the current (global) lower bound, $\mathcal{L}_\Omega$, and the current (local) upper bound, $\mathcal{U}_{I_0,I_1}$.

Step 4. If $\mathcal{U}_{I_0,I_1}(B) \leq \mathcal{L}_\Omega(B)$ for all $B \in [\underline{B}, \overline{B}]$ then the current subtree cannot improve the lower bound and the node may be pruned. Otherwise, select a market $i \in \{1, ..., m\} \setminus (I_0, I_1)$ to branch on, then create the nodes $N(I_0 \cup \{i\}, I_1)$ and $N(I_0, I_1 \cup \{i\})$.

Step 5. If there are any unexplored nodes left, return to Step 1. Otherwise, $\mathcal{F}(B) = L_\Omega$.

---

upper bounds. Note that if $\mathcal{U}^1(B)$ is selected, the only place where the lower bound is equivalent to the upper bound is when $y = [1\ 1\ 1]$. Thus, from the initial node an appropriate market selection to branch on could be either y_1, y_2, or y_3. This could result in going through almost all possibilities before narrowing in on $\mathcal{F}(B)$. If $\mathcal{U}^2(B)$ is selected the decision maker is more likely to select y_1 or y_2 but there still might be several iterations. In this example, the upper bound given by $\mathcal{U}^3(B)$ would be the best fit because the KKT solutions yield fractional values of y_2. Having a fractional candidate solution is ideal for implementing a branch and bound scheme but finding $\mathcal{U}^3(B)$ is computationally expensive.

Another drawback to this approach is that there are no specifications regarding whether the nodes be branched breadth first or depth first. Look at Figure 4.1. With the exception of $y = [1\ 0\ 1]$, the remaining market selections are so far below the initial lower bound that they should never be up for consideration. The branch and bound algorithm could spend time delving deeper into each possible node when a breadth approach would yield quicker results.

Finally, an additional limitation to this approach is that adding to the pool of unexplored nodes two at a time may be especially time consuming, particularly for large problem instances.

With these shortcomings in mind, we specify an alternative branch and bound algorithm in which we explicitly incorporate the values of B for which gaps exist between the lower bound and the upper bound. The goal is to quickly target the potential areas for improvement by enumerating the areas in which gaps between the upper and lower bounds exist and strategically identifying which markets to use for branching decisions.

Define $\mathcal{B}$ as an interval of B values at any node for which there exists a gap between the current lower bound and the current upper bound. A node in the tree will be denoted by $N(I_0, I_1, \mathcal{B}_k)$. The initialization of Algorithm 2 remains essentially the same as Algorithm 1. The root node is $N(I_0 = \emptyset, I_1 = \emptyset, \mathcal{B}_k)$ and $\mathcal{B}$ where defined over all values of B, $\mathcal{B}_k : [\underline{B}, \overline{B}] \to \mathbb{R} \cup \{-\infty\}$.

Using a similar "restricted sorting" scheme as before, we identify the potential markets under consideration. Given that a market is fixed as in or out, we add additional markets in order of nondecreasing $\frac{\xi_i}{\sigma_i^2}$. An important effect of restricting ourselves to each $\mathcal{B}_k$ is that we can eliminate any element of Ω that is not defined over $\mathcal{B}_k$. We define this set as $Y(I_0, I_1, \mathcal{B}_k)$.

The critical difference between the algorithms is the population of the nodes. Algorithm 2 divides the problem into k subsets of $\mathcal{B}$ where k is the number of gaps between the upper and lower bound so we can define $\mathcal{B}_k$ as the collection of intervals in $\mathcal{B}$ for which there exist gaps between the upper and lower bound. For each k subset, two nodes are added. Thus each iteration of the branch and bound algorithm results in the addition of $2k$ more unexplored nodes. This algorithm works best in circumstances where breadth is required over depth. The details of the algorithm are provided by Algorithm 2.

Algorithm 2 Branch and Bound II

Step 0. Set $I_0 = I_1 = \emptyset$, $\Omega = \emptyset$.

Step 1. Choose an unexplored node, $N(I_0, I_1, \mathcal{B}_k)$.

Step 2. Define the incumbent, $\Omega = \Omega \cup Y(I_0, I_1)$.

Step 3. Determine the current (global) lower bound, $\mathcal{L}_\Omega : \mathcal{B} \to \mathbb{R} \cup \{-\infty\}$, and the current (local) upper bound, $\mathcal{U}_{I_0,I_1} : \mathcal{B} \to \mathbb{R} \cup \{-\infty\}$.

Step 4. If $\mathcal{U}_{I_0,I_1}(B) \leq \mathcal{L}_\Omega(B)$ for all $B \in \mathcal{B}$ then the current subtree cannot improve the lower bound on $\mathcal{B}$ and the node may be pruned. Otherwise, select a market $i \in \{1, ..., m\} \setminus (I_0, I_1)$ to branch on and let $\mathcal{B}_k (k = 1, ..., K)$ be the collection of intervals in $\mathcal{B}$ on which $\mathcal{U}_{I_0,I_1}(B) > \mathcal{L}_\Omega(B)$. Then create the nodes $N(I_0 \cup \{i\}, I_1, \mathcal{B}_k)$ and $N(I_0, I_1 \cup \{i\}, \mathcal{B}_k)$ for $(k = 1, ..., K)$.

Step 5. If there are any unexplored nodes left, return to Step 1. Otherwise, $\mathcal{F}(B) = L_\Omega$.

Recall the three market example as described in Section 4.3. For each of the three potential upper bounds previously described, we identify the range of values and the number of $\mathcal{B}_k$ intervals required to initialize Algorithm 2. In Figures 4.10, 4.12, and 4.12 each initial lower bound given by the sorted solutions is depicted with a think black curve and the upper bound is given by a dotted black line. The range of values for which there are gaps between the lower bound and upper bound are highlighted by yellow.

When using $\mathcal{U}^1(B)$, clearly the second approach offers no improvement. However, you can see in Figures 4.12 and 4.12 in which we use $\mathcal{U}^2(B)$ and $\mathcal{U}^3(B)$ respectively, narrowing down the window in which market selections are added to Ω would beneficially result in eliminating certain vectors from consideration. Working within a specified interval, $\mathcal{B}_k$, results in a more targeted approach. However as with the previous algorithm, the problem parameters dictate the effectiveness of the algorithm.

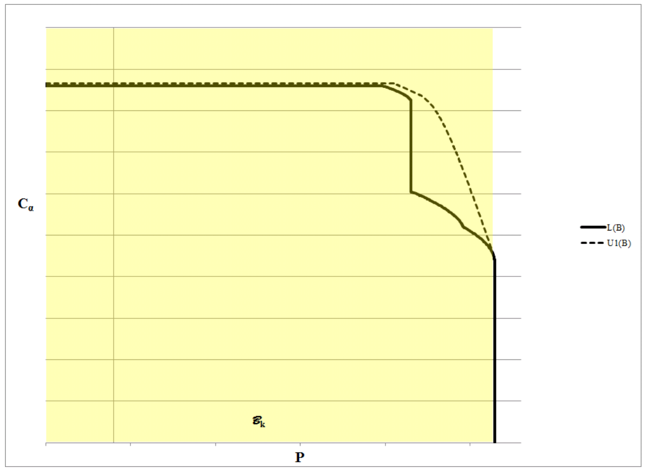

Figure 4.10: Initialization of Algorithm 2 for a three market example with $\mathcal{U}^1(B)$.

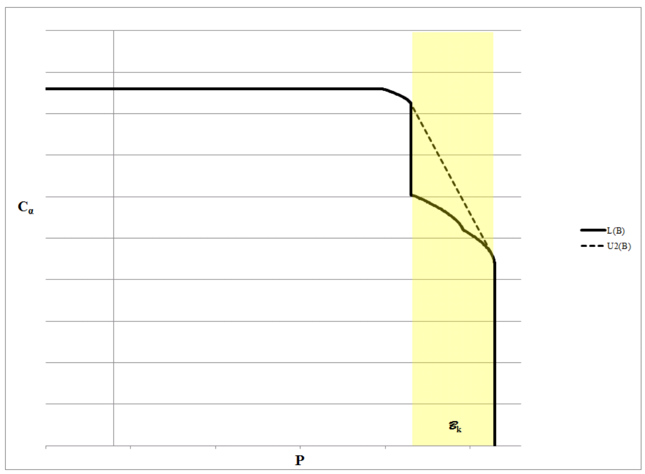

Figure 4.11: Initialization of Algorithm 2 for a three market example with $\mathcal{U}^2(B)$.

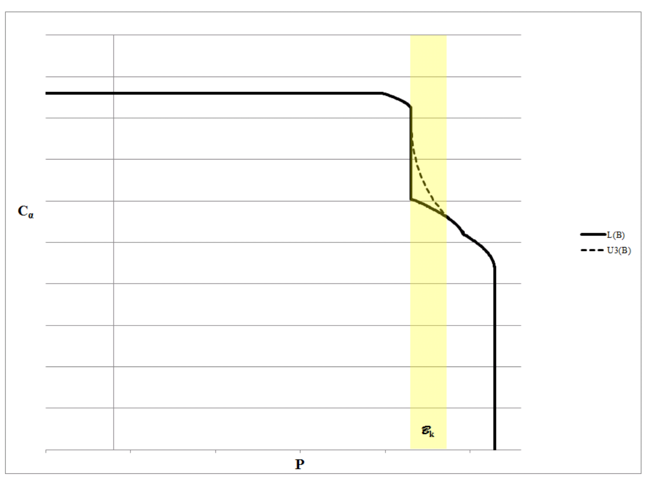

Figure 4.12: Initialization of Algorithm 2 for a three market example with $\mathcal{U}^3(B)$.

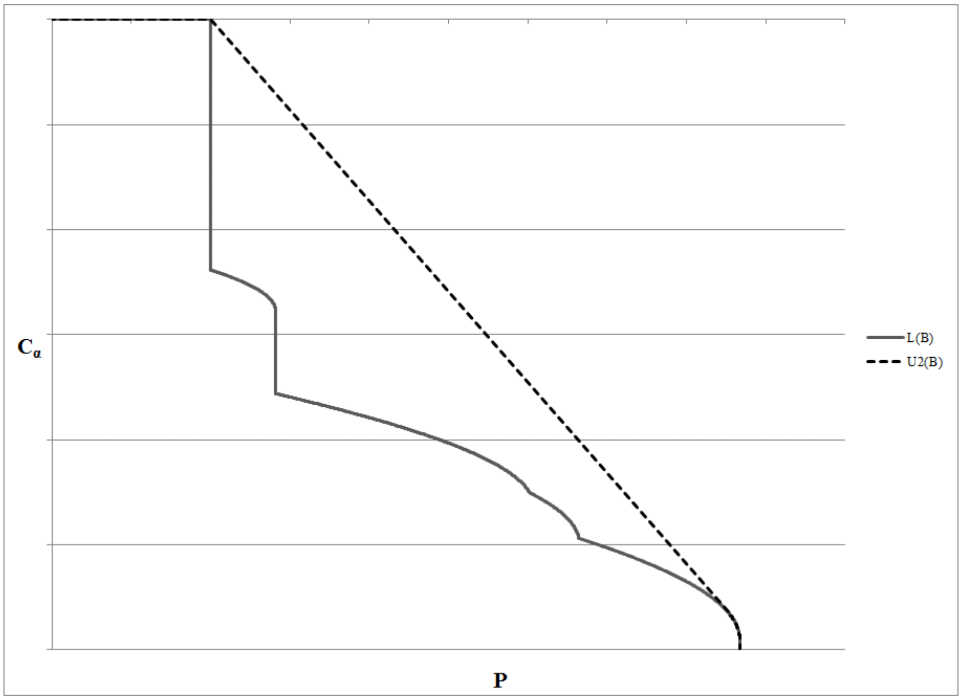

Figure 4.13: $\mathcal{L}(B)$ and $\mathcal{U}^2(B)$ for a four market example.

4.5 Illustrative Example

In this section, we consider a four market example and describe the steps of the branch and bound procedure required to find $\mathcal{F}(B)$. In this example, the markets are conveniently ranked $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2} \geq \frac{\xi_4}{\sigma_4^2}$ and the lower bound, $\mathcal{L}(B)$ is comprised of $y = [0\ 0\ 0\ 0], y = [1\ 0\ 0\ 0], y = [1\ 1\ 0\ 0], y = [1\ 1\ 1\ 0]$, and $y = [1\ 1\ 1\ 1]$. We initially consider the weighted sum tradeoff and compute the upper bound, $\mathcal{U}^2(B)$ for the sorted solutions. This is shown in Figure 4.13.

Since $\mathcal{U}^2(B)$ provides no real direction for the branch and bound algorithm, we solve a relaxation of the problem to find $\mathcal{U}^3(B)$ which is given by Figure 4.14. There are fractional market selection solutions for markets 1, 2 and 3.

Note that the upper bound is equivalent to the lower bound in several sections as shown by Figure 4.14, so we defer to the second branch and bound algorithm and

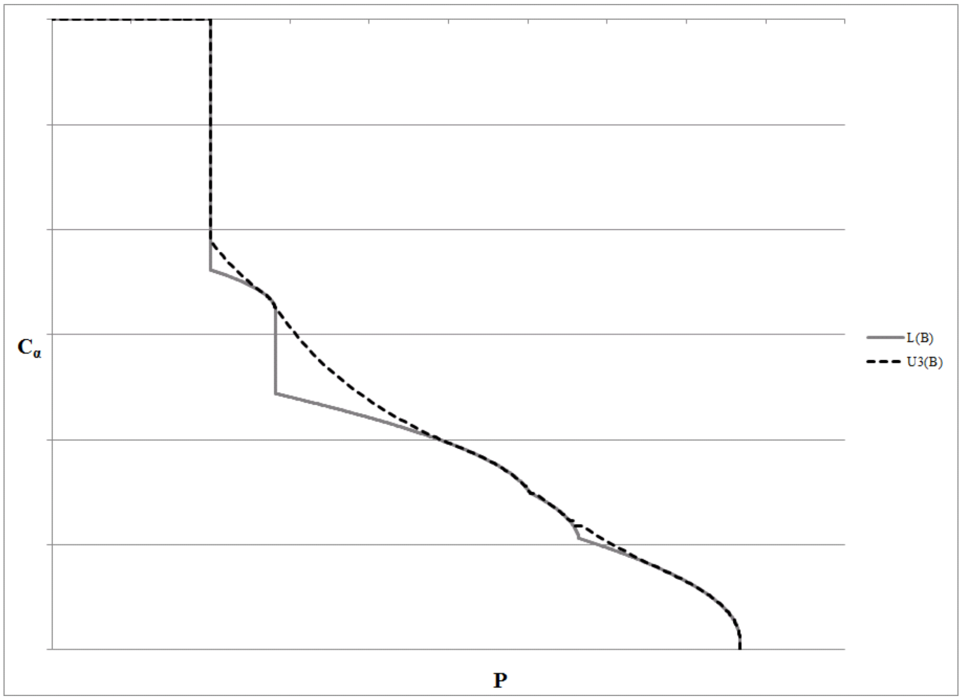

Figure 4.14: $\mathcal{L}(B)$ and $\mathcal{U}^3(B)$ for a four market example.

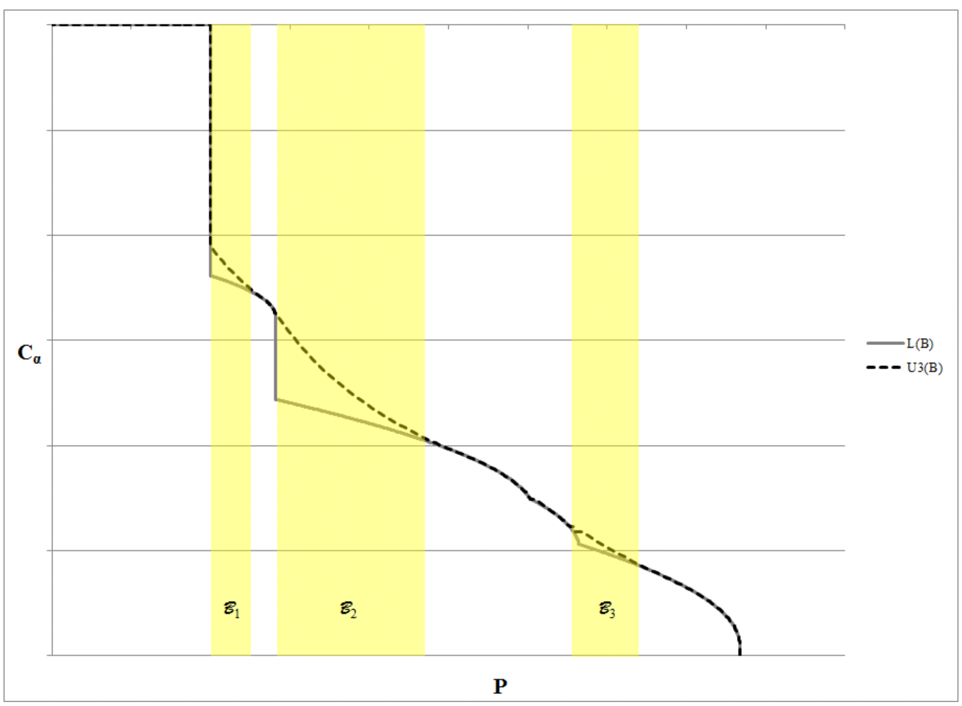

Figure 4.15: Initialization of Algorithm 2 for a four market example with $\mathcal{U}^3(B)$.

identify intervals of $\mathcal{B}$ for which there exists a gap. These intervals are identified in yellow in Figure 4.15. Our results allow us to divide the root node into 6 separate nodes:

$$N(I_0 = \emptyset, I_1 = 1, \mathcal{B}_1)$$

$$N(I_0 = 1, I_1 = \emptyset, \mathcal{B}_1)$$

$$N(I_0 = \emptyset, I_1 = 2, \mathcal{B}_2)$$

$$N(I_0 = 2, I_1 = \emptyset, \mathcal{B}_2)$$

$$N(I_0 = \emptyset, I_1 = 3, \mathcal{B}_3)$$

$$N(I_0 = 3, I_1 = \emptyset, \mathcal{B}_3).$$

Consider the node $N(I_0 = \emptyset, I_1 = 3, \mathcal{B}_3)$ where $\Omega = [0\ 0\ 1\ 0], [1\ 0\ 1\ 0], [1\ 1\ 1\ 0]$, and $[1\ 1\ 1\ 1]$. A simple analysis of their maximum expected profit shows that the solutions $y = [0\ 0\ 1\ 0], [1\ 0\ 1\ 0]$ are not defined on $\mathcal{B}_3$ and the remaining solutions are sorted solutions that already are established as part of $\mathcal{F}(B)$. This node is quickly pruned. Similarly, in the node $N(I_0 = \emptyset, I_1 = 3, \mathcal{B}_3)$, the solution $y = [1\ 1\ 0\ 1]$ is the only possible solution valid for $\mathcal{B}_3$ and is shown to be on $\mathcal{F}(B)$. This quick method of pruning is especially valuable for the market selections closest to the profit maximizing selection.

The remaining nodes are not so straightforward and require a bit more analysis. Consider the two nodes created for the range $\mathcal{B}_2$. To evaluate each node, we compute a corresponding upper bound and compare it to our existing best current solution. In Figure 4.16 we show that the upper bound for $N(I_0 = \emptyset, I_1 = 2, \mathcal{B}_2)$ is equivalent to the current best solution in that range and thus we prune the node. In Figure 4.17 we show that the upper bound for $N(I_0 = 2, I_1 = \emptyset, \mathcal{B}_2)$ exceeds the current local bound and therefore we must continue branching. A similar analysis is done

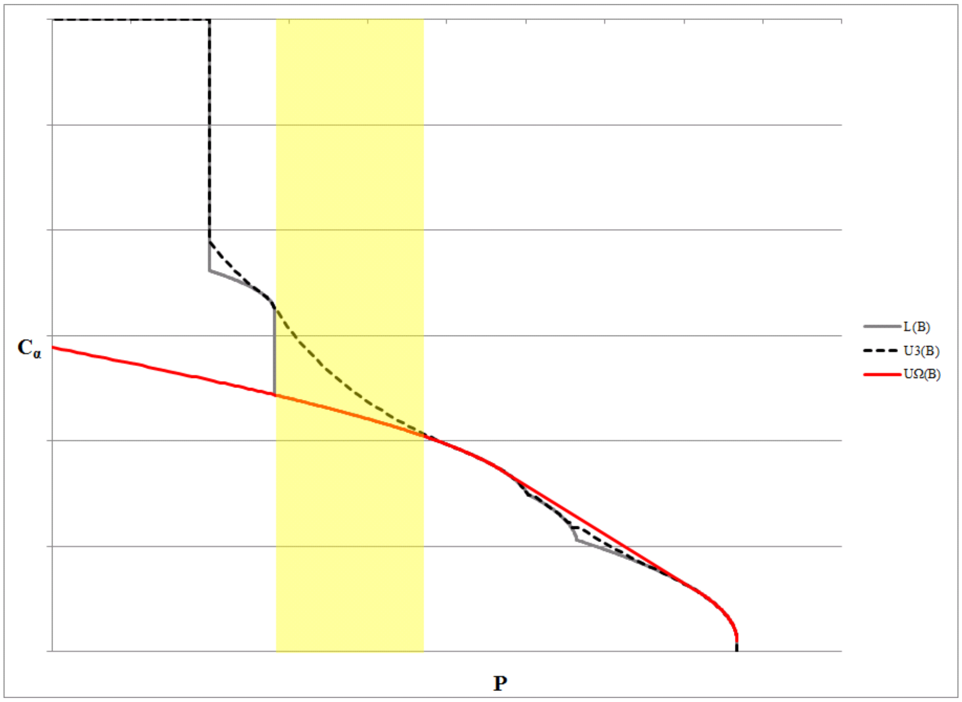

Figure 4.16: $\mathcal{U}_{I_0=\emptyset,I_1=2}(\mathcal{B}_2)$ for $N(I_0=\emptyset, I_1=2, \mathcal{B}_2)$.

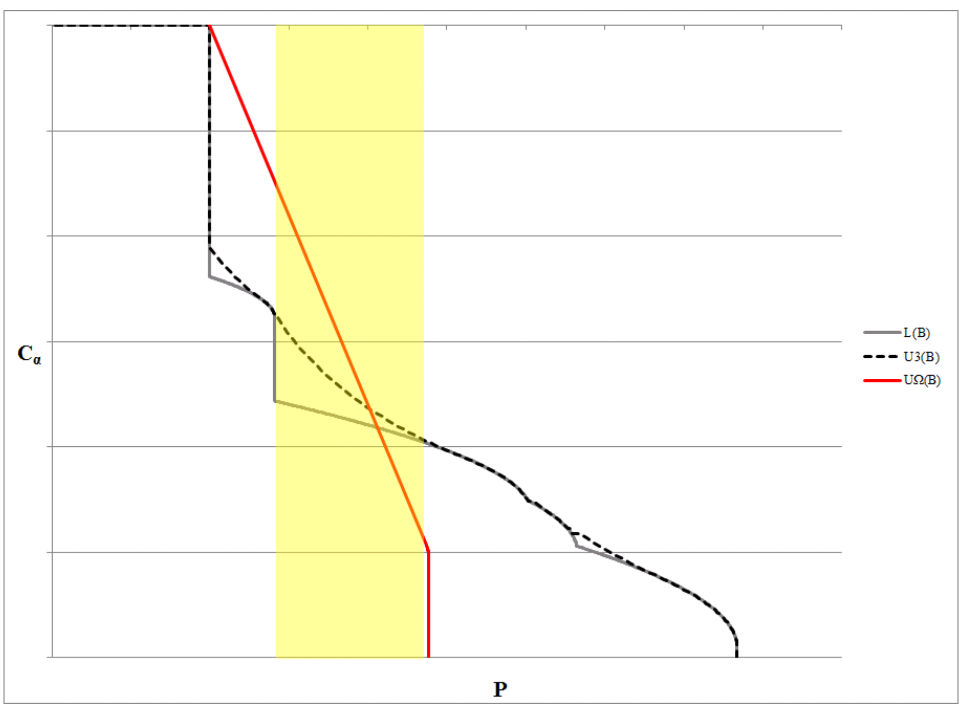

Figure 4.17: $\mathcal{U}_{I_0=\emptyset,I_1=2}(\mathcal{B}_2)$ for $N(I_0=2, I_1=\emptyset, \mathcal{B}_2)$.

for the nodes created in $\mathcal{B}_1$.

After all nodes are explored, we are left with $\mathcal{F}(B)$ that consists of the original sorted solutions as well as market selection vectors $y = [0\ 1\ 0\ 0]$ and $y = [1\ 1\ 0\ 1]$.

4.6 Conclusion

In this chapter we developed a method to identify the expected profit-CVaR Pareto efficient frontier for the *Selective Newsvendor Problem*. We defined four approximations which act as bounds to be used in conjunction with a branch and bound framework commonly used to solve integer programming problems. We discussed the tightness of each approximation and illustrated our discussion with some computational results. We debated the merits of using the various upper bounds in either branch and bound algorithm. Our methods are relatively simple and straightforward to implement and provide interesting results regarding how market selection decisions change as risk preferences are adjusted. In the next chapter we continue this type of analysis by considering the expected profit-VaR Pareto efficient frontier.

CHAPTER V

Risk-Averse Selective Newsvendor Problems with a VaR Risk Measure

5.1 Introduction

The two preceding chapters focused on the risk-averse *Selective Newsvendor Problem* using a Conditional Value-at-Risk (CVaR) risk measure. In this chapter, we shift our focus to the Value-at-Risk (VaR) risk measure. Value-at-Risk (VaR) is a risk assessment technique commonly used in the finance industry. Developed in the late 1980s and early 1990s, VaR measures the potential loss in value of a risky asset or portfolio over a defined period for a given confidence interval. While CVaR is more commonly used in the financial engineering literature because of the mathematical properties associated with it, Value-at-Risk is frequently used in practice in the finance industry. The mathematical properties that make CVaR appealing include monotonicity, subadditivity, homogeneity, and its translational invariance; that is, it is a coherent risk measure. Value-at-Risk lacks the subadditivity property and therefore, in general, is not coherent. Recall that subadditivity is the property of a function such that the sum of the function's values at two given elements is greater than or equal to the value of the function at the sum of the two elements. In fact, VaR is coherent when losses are normally distributed and the portfolio value is a linear function of the asset prices, but that is not the case under consideration, so we assume Value-at-Risk measure is not coherent for the remainder of the chapter.

In this chapter, we consider a risk-averse selective newsvendor utilizing a VaR risk measure. Initially, we evaluate the optimal inventory and market selection policy for a selective newsvendor who only optimizes VaR. As in previous chapters, we derive the optimal procurement policy for a risk-averse newsvendor with a VaR measurement and then show that this order quantity is independent of the market selection decision and the optimal market selection decision is one of a few candidate solutions determined by ranking the net revenue to variance for each potential market.

Next we examine the tradeoff between VaR and expected profit for the selective newsvendor by maximizing the Value-at-Risk function subject to a lower bound constraint on expected profit. Özler et al. [25] and Gan et al. [13] studied a version of this problem for the newsvendor without a selection decision in which they maximized expected profit subject to a Value-at-Risk constraint. The earlier models either lacked expediting (shortage) costs (Özler et al. [25]) or lacked both expediting costs and salvage returns (Gan et al. [13]). Prior to incorporating the market selection decision, we provide the methodology for computing the VaR-expected profit Pareto efficient frontier for a single market or market selection. Analyzing the selective version of the tradeoff problem yields the result that a branch and bound algorithm as demonstrated in Chapter IV can be applied to the VaR case as well. We also provide a relaxation for this problem to identify an upper bound as required by the branch and bound procedure.

This chapter is setup as follows: Section 5.2 studies the selective version of this problem; Section 5.4 analyzes the mean-VaR tradeoff, Section 5.5 details a relaxation that can be used to find an upper bound in the Branch and Bound algorithm and Section 5.7 contains concluding remarks.

5.2 Value-at-Risk for the Selective Newsvendor

In this section we analyze a risk-averse version of the *Selective Newsvendor Problem* in which the decision maker utilizes a Value-at-Risk (VaR) risk measure. We explicitly define the optimization problem required to find the order quantity that optimizes VaR in this context. In the risk-neutral version of this problem, the optimal market selection is included in a set of intuitively ranked candidate solutions. We show that this result extends to a risk-averse selective newsvendor with a VaR criterion. We review the notation regarding the *Selective Newsvendor Problem* to formulate the optimization problem to optimize VaR. As in earlier chapters, we then divide our analysis into two cases based on the parameter values.

5.2.1 Notation and problem formulation

As a portfolio optimization tool, VaR measures the potential loss in value of a risky asset or portfolio over a defined period of time for a given confidence level. For the selective newsvendor, we use VaR to define a threshold value such that the probability expected profit does not exceed this threshold is given by a pre-specified risk level, say $(1-\alpha)$. A byproduct of the analysis completed by Özler et al. [25] and Gan et al. [13] was the order quantity that maximized VaR for two simple single newsvendor problems. Gan et al. [13] considered a newsvendor who purchased inventory for a per unit cost and received a per unit revenue for each item sold. Özler et al. [25] introduced a per unit salvage value to the expected profit calculation for any excess inventory.

Our model, in addition to integrating a market selection component, explicitly ensures that we meet all demand either from inventory or by placing an order with another retailer during the selling season by incorporating a per unit expediting cost for any unmet demand. Using the notation provided in Chapter II, the profit for the selective newsvendor is given by (2.3) and repeated below for convenience

$$\pi(Q; D_y) = (r - v)D_y - (c - v)Q - (e - v)(D_y - Q)^+ - S_y.$$

We employ a Value-at-Risk (VaR_α) measure to define a threshold value, θ, such that the probability expected profit is less than θ is at most $(1 - \alpha)$. Note that $\alpha \in [0, 1)$ reflects the degree of risk aversion, with larger values of α representing a higher degree of risk aversion. To find the order quantity that maximizes VaR_α we solve the following,

$$\begin{array}{rrcl} \max & & \theta & \\ \text{subject to:} & \Psi(\theta, Q, y) & \leq & 1 - \alpha \\ & \theta, Q & \in & \mathbb{R} \\ & y & \in & \{0, 1\}^m. \end{array} \tag{VaR_α}$$

where $\Psi(\theta, Q, y) = \Pr(\pi(Q; D_y) < \theta)$. Using the expression for the selective newsvendor's profit, we find

$$\begin{aligned} \Psi(\theta, Q, y) &= \Pr((r - e)D_y - (c - v)Q + (e - v)\min(Q, D_y) - S_y < \theta) \\ &= \Pr\Big((r - v)D_y - (c - v)Q - S_y < \theta \ \text{ or } \ (r - e)D_y + (e - c)Q - S_y < \theta\Big) \\ &= \begin{cases} \Pr\Big(D_y < \max\left(\frac{\theta + (c-v)Q + S_y}{(r-v)}, \frac{\theta - (e-c)Q + S_y}{(r-e)}\right)\Big) & \text{for } e \leq r \\ \Pr\Big(D_y < \frac{\theta + (c-v)Q + S_y}{(r-v)}\Big) + \Pr\Big(D_y > \frac{\theta - (e-c)Q + S_y}{(r-e)}\Big) & \text{for } e > r \end{cases} \\ &= \begin{cases} \max\left(F_y\left(\frac{\theta + (c-v)Q + S_y}{(r-v)}\right), F_y\left(\frac{\theta - (e-c)Q + S_y}{(r-e)}\right)\right) & \text{for } e \leq r \\ F_y\left(\frac{\theta + (c-v)Q + S_y}{(r-v)}\right) + \bar{F}_y\left(\frac{\theta - (e-c)Q + S_y}{(r-e)}\right) & \text{for } e > r. \end{cases} \end{aligned}$$

From this point forward it will be useful to divide our analysis into two separate cases: one in which the unit expediting cost does not exceed unit revenue ($e \leq r$) and one in which unit expediting cost does exceed unit revenue ($e > r$). As described in previous chapters the former case implies that expediting any unmet orders remains profitable while expediting in the latter case is not profitable but perhaps required by contract.

5.2.2 Case 1: unit expediting cost does not exceed unit revenue

The first case under consideration is characterized by $e \leq r$, which implies that satisfying demand for items through expediting remains profitable, although of course less so than satisfying demand from inventory. We first provide an explicit expression for the optimal order quantity that maximizes (VaR_α) for a fixed market selection, y.

Theorem 5.1. *The optimal order quantity for a risk-averse newsvendor problem with $e \leq r$ who solves* (VaR_α) *for a fixed y is*

$$Q_y^*(\alpha) = F_y^{-1}(1-\alpha).$$

Proof. When $e \leq r$,

$$\Psi(\theta, Q, y) = \max\left(F_y\left(\frac{\theta + (c-v)Q + S_y}{(r-v)}\right), F_y\left(\frac{\theta - (e-c)Q + S_y}{(r-e)}\right)\right).$$

Note that $\left(\frac{\theta+(c-v)Q+S_y}{(r-v)}\right) \leq \left(\frac{\theta-(e-c)Q+S}{(r-e)}\right)$ when $(r-c)Q - S_y \leq \theta$. Thus,

$$\Psi(\theta, Q, y) = \begin{cases} F_y\left(\frac{\theta+(c-v)Q+S_y}{(r-v)}\right) & \text{for } \theta \leq (r-c)Q - S_y \\ F_y\left(\frac{\theta-(e-c)Q+S_y}{(r-e)}\right) & \text{for } \theta \geq (r-c)Q - S_y. \end{cases}$$

One can see that $\theta = (r-c)Q - S_y$, the two expressions are equivalent and $\Psi(\theta, Q, y) = F_y(Q)$. For a fixed y, $\Psi(\theta, Q, y)$ is simply a cumulative distribution function (c.d.f.)

that increases to 1 as θ increases. Thus, the constraint in (VaR_α) is binding. Setting $\Psi(\theta, Q, y) = 1 - \alpha$ and solving for θ yields

$$\theta^*(Q,y) = \begin{cases} (r-v)F_y^{-1}(1-\alpha) - (c-v)Q - S_y & \text{for } F_y(Q) \geq (1-\alpha), \\ (r-e)F_y^{-1}(1-\alpha) + (e-c)Q - S_y & \text{for } F_y(Q) < (1-\alpha). \end{cases} \tag{5.1}$$

Clearly, for a fixed y, $\theta^*(Q, y)$ is increasing for values of $Q \leq F_y^{-1}(1-\alpha)$ and decreasing when $Q \geq F_y^{-1}(1-\alpha)$. So the corresponding order quantity that maximizes the Value-at-Risk, $\theta^*(Q, y)$, is $Q_y^* = F^{-1}(1-\alpha)$. $\square$

Note that risk-averse order quantity that optimizes (VaR_α) may be smaller or larger than the risk-neutral order quantity, $Q_y = F_y^{-1}(\rho)$, depending on the problem parameters. Since D_y follows a normal distribution, we can express the order quantity as a function of y as shown by (2.5). Using this relationship, we can show that when $e \leq r$, the problem (VaR_α) is equivalent to original SNP optimization problem (SNP) described in Chapter II.

Theorem 5.2. *The VaR criterion for a risk-averse selective newsvendor with $e \leq r$ whose order quantity satisfies* (2.5) *is*

$$\theta\left(F_y^{-1}(\beta), y\right) = \xi_y + \vartheta(\beta)\sigma_y$$

where $\vartheta(\beta)$ only depends on β and the problem parameters.

Proof. For a fixed selection vector y, the optimal value of θ as a function of the order quantity, Q, is given by

$$\theta^*(Q,y) = \begin{cases} (r-v)F_y^{-1}(1-\alpha) - (c-v)Q - S_y & \text{for } F_y(Q) \geq (1-\alpha), \\ (r-e)F_y^{-1}(1-\alpha) + (e-c)Q - S_y & \text{for } F_y(Q) < (1-\alpha). \end{cases}$$

When Q is a function of y as described by (2.5),

$$\theta_y^*(\beta) = \begin{cases} (r-c)\mu_y + \left((r-v)\Phi^{-1}(1-\alpha) - (c-v)\Phi^{-1}(u)\right)\sigma_y - S_y & \text{for } \beta \geq 1-\alpha, \\ (r-c)\mu_y + \left((r-e)\Phi^{-1}(1-\alpha) + (e-c)\Phi^{-1}(u)\right)\sigma_y - S_y & \text{for } \beta < 1-\alpha. \end{cases}$$

Thus,

$$\theta\left(F_y^{-1}(\beta), y\right) = \xi_y + \vartheta(\beta)\sigma_y$$

where

$$\vartheta_\alpha(\beta) = \begin{cases} (r-v)\Phi^{-1}(1-\alpha) - (c-v)\Phi^{-1}(\beta) & \text{for } \beta \geq 1-\alpha, \\ (r-e)\Phi^{-1}(1-\alpha) + (e-c)\Phi^{-1}(\beta) & \text{for } \beta < 1-\alpha \end{cases} \tag{5.2}$$

Since $\vartheta_\alpha(\beta)$ is only a function of β and the problem parameters, we achieve the desired result. □

From Theorem 5.1 we can see that, when $e \leq r$, the optimal order quantity is of the form (2.5) with $\beta = 1-\alpha$ which depends only on the problem parameters. Using 5.2 yields allows us to rewrite (VaR_α) as

$$\begin{array}{rrcl} \max & \xi_y & + & \vartheta(\beta)\sigma_y \\ \text{subject to:} & y & \in & \{0,1\}^m \end{array}$$

where $\beta = 1-\alpha$ and $\vartheta(\beta)$ is independent of y. This of course is a variant (SNP) which exhibits the well-known sorting solution structure described in previous chapters. That is, the optimal solution to (SNP) can be found in the set

$$Y = \left\{y^{(\ell)} : \ell = 0, \ldots, m\right\}$$

where

$$y_i^{(\ell)} = \begin{cases} 1 & \text{for } i = 1, \ldots, \ell \\ 0 & \text{for } i = \ell+1, \ldots, m \end{cases}$$

such that if market ℓ is selected, markets $1, \ldots, \ell - 1$ are selected as well.

Note of course that since the second term in the objective function has a positive sign, that if $\vartheta(1-\alpha) \geq 0$ the optimal solution is to operate in all markets.

5.2.3 Case 2: unit expediting cost exceeds unit revenue

The second case that we will consider is characterized by $e > r$, which occurs when satisfying demand for items through expediting is costly, for example, if there either is a contractual obligation to satisfy all demands. Unfortunately, this case does not have an analytical expression for $Q^*_y(\alpha)$ but we will show that the optimal market selection is one of the sorted solutions in set Y.

Theorem 5.3. *A selective newsvendor with $e > r$ who solves* (VaR_α) *has an optimal solution in Y.*

Proof. When $e > r$,

$$\Psi(\theta, Q, y) = F_y\left(\frac{\theta + (c-v)Q + S_y}{(r-v)}\right) + \bar{F}_y\left(\frac{\theta - (e-c)Q + S_y}{(r-e)}\right).$$

For a fixed y, $\left(\frac{\theta+(c-v)Q+S_y}{(r-v)}\right) \leq \left(\frac{\theta-(e-c)Q+S_y}{(r-e)}\right)$ for $(r-c)Q - S_y \leq \theta$. Thus,

$$\Psi(\theta, Q, y) = \begin{cases} F_y\left(\frac{\theta+(c-v)Q+S_y}{(r-v)}\right) + \bar{F}_y\left(\frac{\theta-(e-c)Q+S_y}{(r-e)}\right) & \text{for } \theta \leq (r-c)Q - S_y, \\ 1 & \text{for } \theta \geq (r-c)Q - S_y. \end{cases}$$

Since $\Psi(\theta, Q, y)$ for a given y is increasing to 1 as $\theta \to \infty$, we know that the constraint $\Psi(\theta, Q, y) \leq 1 - \alpha$ is binding for $\alpha \in (0, 1)$. Therefore, to solve (VaR_α) we must find the θ^* and Q^* that satisfy

$$F_y\left(\frac{\theta - (e-c)Q + S_y}{(r-e)}\right) - F_y\left(\frac{\theta + (c-v)Q + S_y}{(r-v)}\right) = \alpha.$$

Since an analytical solution to the above expression does not, in general, exist, we reparameterize the problem with a new decision variable, $u \in (0, 1-\alpha)$, by defining

$$u = F_y\left(\frac{\theta - (e-c)Q + S_y}{(r-e)}\right) - \alpha = F_y\left(\frac{\theta + (c-v)Q + S_y}{(r-v)}\right).$$

Solving for Q and θ yields

$$Q_y(u) = \left(\frac{r-v}{e-v}\right) F_y^{-1}(u) + \left(\frac{e-r}{e-v}\right) F_y^{-1}(u+\alpha),$$

and

$$\theta_y(u) = (e-c)\left(\frac{r-v}{e-v}\right) F_y^{-1}(u) - (c-v)\left(\frac{e-r}{e-v}\right) F_y^{-1}(u+\alpha) - S_y.$$

Thus we can rewrite (VaR_α) as

$$\begin{array}{lrcl} \max & (e-c)\left(\frac{r-v}{e-v}\right) F_y^{-1}(u) & - & (c-v)\left(\frac{e-r}{e-v}\right) F_y^{-1}(u+\alpha) - S_y \\ \text{subject to:} & u & \in & (0, 1-\alpha) \\ & y & \in & \{0,1\}^m. \end{array}$$

Using the normality of D_y, we can show $F_y^{-1}(u)$ satisfies the expression in (2.5) and we rearrange terms so that the problem we are solving becomes

$$\begin{array}{rrcl} \max & \xi_y & + & \vartheta(u)\sigma_y \\ \text{subject to:} & u & \in & (0, 1-\alpha) \\ & y & \in & \{0,1\}^m. \end{array}$$

where

$$\vartheta_\alpha(u) = (e-c)\left(\frac{r-v}{e-v}\right)\Phi^{-1}(u) - (c-v)\left(\frac{e-r}{e-v}\right)\Phi^{-1}(u+\alpha). \tag{5.3}$$

You can see that for a fixed u, the problem is equivalent to the original (SNP) problem described in Chapter II. Thus, although the selection of u may depend on y we actually need only consider market selection vectors in Y to find the optimal solution. $\square$

Therefore, the set Y contains an optimal solution for the risk-neutral selective newsvendor and the risk-averse selective newsvendor who optimizes a VaR or CVaR criterion. In the next section, we consider the tradeoff that exist between expected profit and VaR for a single newsvendor so that we can transition to analysis of the tradeoff problem for the selective newsvendor in Section 5.4.

5.3 Expected Profit-VaR Pareto Efficient Frontier for a Single Newsvendor

To analyze the tradeoff that exists between expected profit and Value-at-Risk for a risk-averse selective newsvendor we first must consider the tradeoff problem for a fixed market selection, or equivalently, a single newsvendor. The idea of a mean-VaR tradeoff was first introduced for the portfolio selection problem (see,e.g., Gaivoronski and Pflug [12] and Tsao [38]). Gan et al. [13] studied a supply chain with a risk-neutral supplier and single retailer with a downside risk constraint. Özler et al. [25] considered a two-product single newsvendor with a downside risk constraint.

To study the mean-VaR tradeoff for a given market selection, we maximize VaR subject to a lower bound constraint on expected profit, B,

$$\begin{array}{lrl} \max & VaR_\alpha(Q) & \\ \text{subject to:} & E[\pi(Q)] & \geq B \\ & Q & \in \mathbb{R} \end{array}$$

where $E[\pi(Q)]$ is given by (2.2) and $VaR_\alpha(Q)$ is described by the optimization problem

$$\begin{array}{lrcl} \max & \theta & & \\ \text{subject to:} & \Psi(\theta, Q) & \leq & 1-\alpha \\ & \theta, Q & \in & \mathbb{R} \end{array} \qquad (\mathrm{VaR}_\alpha(\mathrm{Q}))$$

where $\Psi(\theta, Q) = \Pr(\pi(Q; D) < \theta)$. For a fixed market selection,

$$\Psi(\theta, Q) = \begin{cases} \max\left(F\left(\frac{\theta+(c-v)Q+S}{(r-v)}\right), F\left(\frac{\theta-(e-c)Q+S}{(r-e)}\right)\right) & \text{for } e \leq r \\ F\left(\frac{\theta+(c-v)Q+S}{(r-v)}\right) + \bar{F}\left(\frac{\theta-(e-c)Q+S}{(r-e)}\right) & \text{for } e > r. \end{cases}$$

Note that the solution to (VaR_α(Q)) is the unconstrained solution to the tradeoff problem. As discussed in the previous section, the optimal order quantity for a fixed market selection when $e \leq r$ is $Q^*_\alpha = F^{-1}(1-\alpha)$. When $e > r$, is given by

$$Q^*_\alpha = Q(u^*) = \left(\frac{r-v}{e-v}\right) F^{-1}(u^*) + \left(\frac{e-r}{e-v}\right) F^{-1}(u^* + \alpha)$$

where u^* is the solution to

$$\begin{array}{lrcl} \max & (e-c)\left(\frac{r-v}{e-v}\right) F^{-1}(u) & - & (c-v)\left(\frac{e-r}{e-v}\right) F^{-1}(u+\alpha) - S \\ \text{subject to:} & u & \in & (0, 1-\alpha). \end{array}$$

Returning to the constrained optimization problem under consideration, recall that the expected profit function for the single newsvendor is maximized by $Q^* = F^{-1}(\rho)$. Thus, the problem is infeasible for all values of $B > E[\pi(Q^*)]$. If $E[\pi(Q^*_\alpha)] \geq B$ is true for some value B then the order quantity that maximizes VaR is not only feasible but also optimal.

When values of B fall in the range $[E[\pi(Q^*_\alpha)], E[\pi(Q^*)]]$, we must use a binary search mechanism to find the optimal order quantity that satisfies the constraint. Since the expected profit function is unimodal, we know that the range of values we will be searching over will be either strictly increasing or decreasing based on whether the risk-averse order quantity is greater than or less than the risk-neutral order quantity.

The steps below outline a method for solving the constrained optimization problem to create an expected profit-VaR Pareto efficient frontier for a single newsvendor.

1. Find the optimal risk-averse and risk-neutral order quantities, Q_α^* and Q^*.

2. Identify the range of B values for which the risk-averse order quantity is optimal: $E[\pi(Q_\alpha^*)] \geq B$.

3. Identify the range of B values for which the problem is infeasible: $E[\pi(Q^*)] < B$.

4. For each $B \in [E[\pi(Q_\alpha^*)], E[\pi(Q^*)]]$ find a Q value that solves

$$E[\pi(Q)] = B$$

   by using a binary search algorithm with the initial lower bound given by $\min(Q_\alpha^*, Q^*)$ and the initial upper bound given by $\max(Q_\alpha^*, Q^*)$.

5. Find the corresponding $VaR(Q)$ associated with this value of Q. To find the corresponding $VaR(Q)$, the methods vary based on the parameter values.

   - When the unit expediting cost does not exceed unit revenue ($e \leq r$), compute $VaR(Q)$ using the equation described in Section 5.2.2,

$$VaR(Q) = \begin{cases} (r-v)F^{-1}(1-\alpha) - (c-v)Q - S & \text{for } F(Q) \geq (1-\alpha), \\ (r-e)F^{-1}(1-\alpha) + (e-c)Q - S & \text{for } F(Q) < (1-\alpha). \end{cases}$$

   - When the unit expediting cost exceeds unit revenue ($e > r$), first find the $u \in (0, 1-\alpha)$ that makes

$$Q = \left(\frac{r-v}{e-v}\right) F^{-1}(u^*) + \left(\frac{e-r}{e-v}\right) F^{-1}(u^* + \alpha)$$

     then use the u value to compute the corresponding $VaR(Q(u))$

$$(e-c)\left(\frac{r-v}{e-v}\right) F^{-1}(u) - (c-v)\left(\frac{e-r}{e-v}\right) F^{-1}(u+\alpha) - S.$$

6. Plot the point $(B, VaR(Q))$. Continue for all feasible values of B.

These steps effectively identify an expected profit-VaR Pareto efficient frontier for a single newsvendor. In the next section, we detail how these curves are used to solve the selective version of this problem.

5.4 Expected Profit-VaR Pareto Efficient Frontier for the Selective Newsvendor

In this section, we study the mean-VaR tradeoff for the *Selective Newsvendor Problem* by maximizing VaR subject to a lower bound constraint on expected profit, B,

$$\begin{aligned} \max \quad & VaR_\alpha(Q, y) & \\ \text{subject to:} \quad & E[\pi(Q, y)] & \geq B \\ & Q & \in \mathbb{R} \\ & y & \in \{0,1\}^m. \end{aligned}$$

where $E[\pi(Q)]$ is given by (2.4) and $VaR_\alpha(Q, y)$ is described by the optimization problem (VaR_α). Solving this problem over all possible values of B will result in a Pareto efficient frontier. Since $VaR_\alpha(Q, y)$ is itself an optimization problem, we restructure the optimization problem as

$$\begin{aligned} \max \quad & \theta & \\ \text{subject to:} \quad & \Psi(Q, y) & \leq 1 - \alpha \\ & P(Q, y) & \geq B \\ & Q & \in \mathbb{R} \\ & y & \in \{0,1\}^m. \end{aligned}$$

In Section 5.2, we showed that the $\Psi(Q, y) \leq 1 - \alpha$ constraint is binding and incorporating this constraint leads to defining θ as a function of the market selection y and a fractional $u \in (0, 1)$. Using the expression given in Theorem 5.2 and the reparameterization detailed in Theorem 5.3, we can show that the problem we are trying to solve is

$$\begin{aligned} \max \quad & \xi_y + \vartheta_\alpha(u)\sigma_y \\ \text{subject to:} \quad \xi_y - K_\alpha(u)\sigma_y & \geq B \\ u & \in (0, \bar{u}) \\ y & \in \{0,1\}^m. \end{aligned} \tag{$\theta(B)$}$$

Based on the parameter values, the expressions in ($\theta(B)$) are described below.

- Case 1: unit expediting cost does not exceed unit revenue

  When $e \leq r$, $\vartheta_\alpha(u)$ is given by (5.2) and

  $$K_\alpha(u) = (c - v)\Phi^{-1}(u) + (e - v)\Lambda(u)$$

  where Φ is the c.d.f. of a standard normal random variable and $\Lambda(u)$ is the loss function corresponding to the standard normal distribution. In this case u can take on any fractional amount so that $\bar{u} = 1$.

- Case 2: unit expediting cost exceeds unit revenue

  When $e > r$, $\vartheta_\alpha(u)$ is given by (5.3) and

  $$K_\alpha(u) = (c - v)q(u) + (e - v)\Lambda(\Phi(q(u)))$$

  where Φ is the c.d.f. of a standard normal random variable, $\Lambda(u)$ is the loss function corresponding to the standard normal distribution, and

  $$q(u) = \left(\frac{r - v}{e - v}\right)\Phi^{-1}(u) + \left(\frac{e - r}{e - v}\right)\Phi^{-1}(u + \alpha).$$

  The corresponding $\bar{u} = 1 - \alpha$.

Solving the class of optimization problems described by ($\theta(B)$) over all values of B results in a Pareto efficient frontier which we will denote $\mathcal{T}(B)$. Clearly, this problem has a similar structure to the class of problems described by (C(B)) in

Chapter IV and we can use the equivalent Branch and Bound methods to find the solution, provided that there exists an appropriate upper bound to $(\theta(B))$.

First, define $\mathcal{T}(B) : [\underline{B}, \overline{B}] \rightarrow \mathbb{R}$ where $\overline{B}$ is the unconstrained maximum value of expected profit $P(Q, y)$ and $\underline{B}$ is the expected profit corresponding to the unconstrained maximum value of $VaR_{\alpha}(Q, y))$. For a fixed $y \in \{0, 1\}^m$ there exists an efficient way of determining the corresponding frontier, $\mathcal{T}_y(B)$. Thus,

$$\mathcal{T}(B) = \max_{y \in \{0,1\}^m} \mathcal{T}_y(B).$$

At any point during the branch and bound algorithm we require an incumbent (current best) solution characterized by a number of binary market selection vectors. As in Chapter IV we note that there exists a lower bound to the $(\theta(B))$ which is the concave envelope of all market selections in Y. This initial lower bound is denoted $\mathcal{L}(B)$. Each iteration of the Branch and Bound algorithm yields a different set Ω of potential solutions. Ω initially is comprised of only sorted solutions but additional market selections are added according to our so-called "restricted sorted solution" method in which if certain markets are fixed as either in or out then the remaining markets are still added according to the original sorting scheme. Thus, there exists a lower bound for $\mathcal{T}(B)$:

$$\mathcal{L}_{\Omega}(B) = \max_{y \in \Omega} \mathcal{T}_y(B).$$

Again, we divide each node in the branch and bound procedure into two disjoint subsets of $\{1, ..., m\}$: I_0 and I_1. The former contains the markets that are forced out of the solution while the latter contains the markets that are forced into the solution in the current Branch and Bound subtree. In other words, the current subtree considers an optimization problem of the form:

$$\begin{array}{rrcll}
\max & \xi_y + \vartheta_\alpha(u)\sigma_y & & & \\
\text{subject to:} & \xi_y - K_\alpha(u)\sigma_y & \geq & B & \\
 & u & \in & (0,\bar{u}) & \\
 & y_i & = & 0 & \text{for } i \in I_0 \\
 & y_i & = & 1 & \text{for } i \in I_1 \\
 & y & \in & \{0,1\}^m. &
\end{array} \qquad (\theta_{I_0,I_1}(B))$$

At each node, we again define the set of "restricted sorted solutions" by $Y(I_0, I_1)$. Thus, assuming that there exists an upper bound $\mathcal{U}_{I_0,I_1} : [\underline{B}, \overline{B}] \to \mathbb{R} \cup \{-\infty\}$ for given subsets I_0 and I_1, we can find $\mathcal{T}(B)$ using either branch and bound procedure described by Algorithm 1 or Algorithm 2.

5.5 Upper Bound on the Expected Profit-VaR Pareto Efficient Frontier

A key requirement for the application of a branch and bound algorithm is the existence of an upper bound $\mathcal{U}(B)$ that is the value function solution corresponding to a relaxation of $(\theta(B))$. The non-concavity of ϑ when $e > r$ imposes challenges not met in previous chapters. However, employing the method described in Chapter IV we can find the expected profit-VaR Pareto efficient frontier.

In this section, we relax the requirement that y be binary so that the problem we are solving is

$$\begin{array}{rrl}
\max & \xi_y + \vartheta_\alpha(u)\sigma_y & \\
\text{subject to:} & \xi_y - K_\alpha(u)\sigma_y & \geq B \\
 & u & \in (0,\bar{u}) \\
 & y & \in [0,1]^m.
\end{array} \qquad (\mathrm{R}_\theta(\mathrm{B}))$$

where $\bar{u} = 1$ if $e \le r$ and $\bar{u} = 1 - \alpha$ if $e > r$. We will further simplify this optimization problem by explicitly characterizing a set of potentially optimal solutions. As discussed in Chapter IV the following lemma applies to problems of this form.

Lemma 5.4. *There exists an optimal solution to* $(R_\theta(B))$ *in the set*

$$\bar{Y} = \{\mathbf{0}\} \cup \left\{y^{(\ell)}(t) : t \in (0,1], \ell = 1, \ldots, m\right\}$$

where

$$y_i^{(\ell)}(t) = \begin{cases} 1 & \text{for } i = 1, \ldots, \ell - 1 \\ t & \text{for } i = \ell \\ 0 & \text{for } i = \ell + 1, \ldots, m \end{cases}$$

for some $\ell = 1, ..., m$ *and some* $t \in (0, 1]$.

Proof. Introduce a nonnegative decision variable $z = \sigma_y^2$ to the problem $(R_\theta(B))$ and consider the following optimization problem for fixed u and z:

$$\begin{array}{lrl} \max & \xi_y + \vartheta_\alpha(u)\sqrt{z} & \\ \text{subject to:} & \xi_y & \ge B + K_\alpha(u)\sqrt{z} \\ & \sigma_y^2 & = z \\ & y & \in (0,1)^m. \end{array}$$

Since the second term in the objective function is a constant for fixed u and z, it follows that we can simply solve this optimization problem without the second constraint and verify whether or not the corresponding solution satisfies that constraint. If the optimal solution to the above problem satisfies the constraint, it remains optimal; if not, the problem is infeasible. The desired result follows from the fact that the problem to be solved is a continuous knapsack problem. □

We can now reformulate $(R_\theta(B))$ by explicitly restricting ourselves to $y \in \bar{Y}$. In fact, we will replace $(R_\theta(B)$ by a collection of m optimization problems (one for each value of ℓ) in the decision variables w and t only. For convenience, define the following functions:

$$\xi_\ell(t) = \sum_{i=1}^{\ell-1} \xi_i + \xi_\ell t \qquad \text{for } t \in (0,1]$$

$$s_\ell(t) = \sqrt{\sum_{i=1}^{\ell-1} \sigma_i^2 + \sigma_\ell^2 t} \qquad \text{for } t \in (0,1]$$

and consider the following class of optimization problems:

$$\begin{aligned} \max \quad & \xi_\ell(t) + \vartheta_\alpha(u) s_\ell(t) \\ \text{subject to:} \quad & \xi_\ell(t) - K_\alpha(u) s_\ell(t) \geq B \\ & u \in (0, \bar{u}) \\ & t \in (0,1]. \end{aligned} \qquad (\mathrm{R}_\theta(\mathrm{B,t}))$$

We will use the KKT conditions for $(R_\theta(B,t))$ to characterize candidate solutions to this problem. Noting that the constraints $u \in (0, \bar{u})$ and $t > 0$ cannot be binding, the KKT conditions for this problem can be written as:

$$\begin{aligned} (\eta+1)\xi_\ell'(t) - (\eta K_\alpha(u) + \vartheta_\alpha(u)) s_\ell'(t) &= \delta \\ (\eta K_\alpha'(u) - \vartheta_\alpha'(u)) s_\ell(t) &= 0 \\ \xi_\ell(t) - \eta K_\alpha(w) s_\ell(t) &\geq B \\ t &\in (0,1] \\ u &\in (0, \bar{u}) \\ \eta[B - \xi_\ell(t) + K_\alpha(u) s_\ell(t)] &= 0 \\ \delta(t-1) &= 0 \\ \delta, \eta &\geq 0. \end{aligned}$$

Since $s_\ell(t) > 0$ for all $t \in (0,1]$ we can use the second condition to write η as a function of u, $\eta(u) = \frac{\vartheta_\alpha'(u)}{K_\alpha'(w)}$. This expression is straightforward for the $e \leq r$ case but is significantly harder when $e > r$. We can now identify a set of solutions that contains all Pareto efficient solutions to the optimization problem $(R_\theta(B))$ for some value of B. In particular, consider a fixed value of $\ell = 1, ..., m$ and $u \in (0, \bar{w})$, using the expression for $\eta(w)$ in the first condition yields an expression of δ as a function of w and t:

$$\delta(u,t) = (\eta(u) + 1)\eta_\ell'(t) - (\eta K_\alpha(u) + \vartheta_\alpha(u))s_\ell'(t).$$

We then distinguish between integral and fractional KKT solutions:

1. *Integral selection vector:* From the condition $\delta(t-1) = 0$ we know that $\delta > 0$ implies $t = 1$. Hence we obtain an integral KKT solution if and only if $\delta(w, 1) > 0$.

2. *Fractional selection vector:* There is at most one value of $t \in (0, 1]$ that yields $\delta(w, t) = 0$, which then yields a fractional KKT solution. In particular, there exists no value of t for which $\delta(w, t) = 0$ if $\eta(w)K_\alpha(u) > L_\alpha(u)$. Otherwise, the following value of t sets $\delta(u, t) = 0$:

$$t = \frac{\left([\vartheta_\alpha(u) - \eta(w)K_\alpha(u)]\frac{\sigma_\ell^2}{2\eta(u)\xi_\ell}\right)^2 - \sum_{i=1}^{\ell-1}\sigma_i^2}{\sigma_\ell^2}$$

   which yields a KKT solution provided that $t \in (0, 1]$.

In either case, the point (P, θ) given by

$$P = \xi_\ell(t) - K_\alpha(u)s_\ell(t)$$

$$\theta = \xi_\ell(t) - \vartheta_\alpha(u)s_\ell(t)$$

is potentially on the Pareto efficient frontier. In addition, for a fixed value of $\ell = 1, ..., m$ we obtain a possibly Pareto efficient solution with $w = H_\alpha^{-1}(\Phi^{-1}(\rho))$. Note that in that case, the feasible region of $(R_\theta(B))$ will often be a singleton so that the KKT conditions do not apply. Finally, we also consider the candidate solution $y = \mathbf{0}$. Here are the steps required to convert these points into a frontier:

1. Considering all n candidate points, sort the coordinates (P, C) in decreasing order according to the P values such that

$$(P_1, \theta_1) \geq (P_2, \theta_2) \geq \ldots \geq (P_n, \theta_n) \;\text{ where } P_1 \geq P_2 \geq \ldots P_n.$$

2. Starting with the second sorted point, set

$$\theta'_i = \begin{cases} \theta_i & \text{if } \theta_i > \theta_{i-1} \\ \theta_{i-1} & \text{if } \theta_i \leq \theta_{i-1} \end{cases}$$

3. The upper bound $\mathcal{U}(B)$ will consist of the points (P_1, θ_1) and (P_i, θ'_i) for every $i = 2, ..., n$.

Establishing an upper bound, $\mathcal{U}(B)$ by solving a relaxation of $(\theta(B))$ allows us to implement a branch and bound algorithm like those described in Chapter IV.

5.6 Computational Results

In this section we provide some examples to illustrate our results. First, consider a simple three market example in which the markets are ranked $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2}$ as shown by Figure 5.1. Note the inclusion of the $y = \{\mathbf{0}\}$ market selection indicates that the VaR values associated with this problem are all negative. Value-at-Risk is a threshold for the lowest expected profit returns so negative numbers are not unexpected.

For this problem, we solved the KT conditions specified in the previous section to identify a set of potential Pareto efficient frontiers. These are shown in Figure 5.2. The Pareto efficient candidate solution $(P, \theta) = (0, 0)$ is represented by the black line

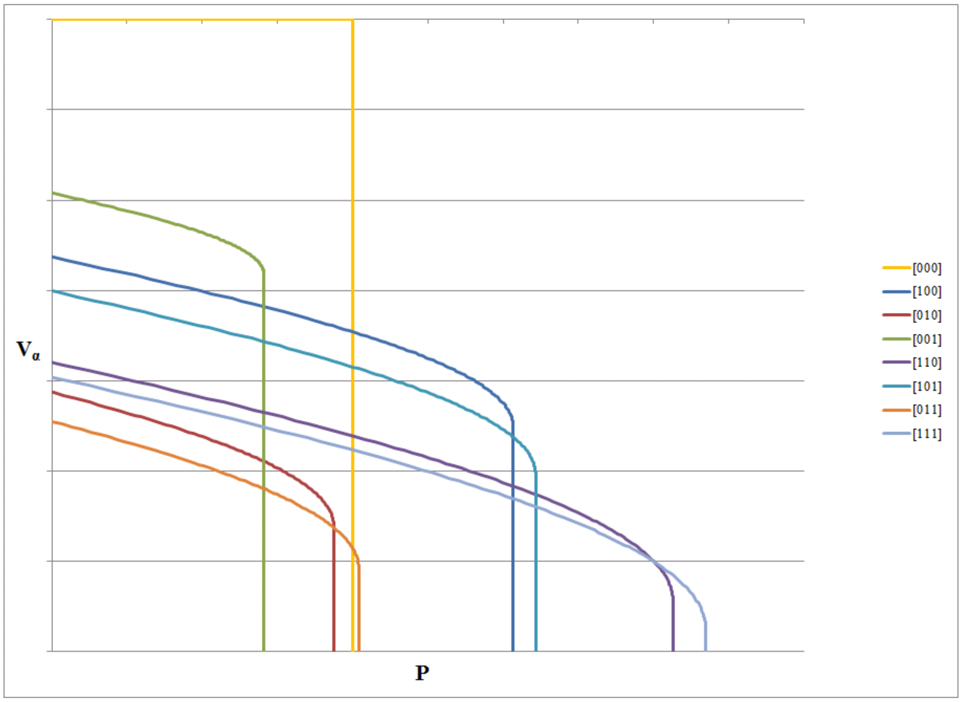

Figure 5.1: Three market example with $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2}$

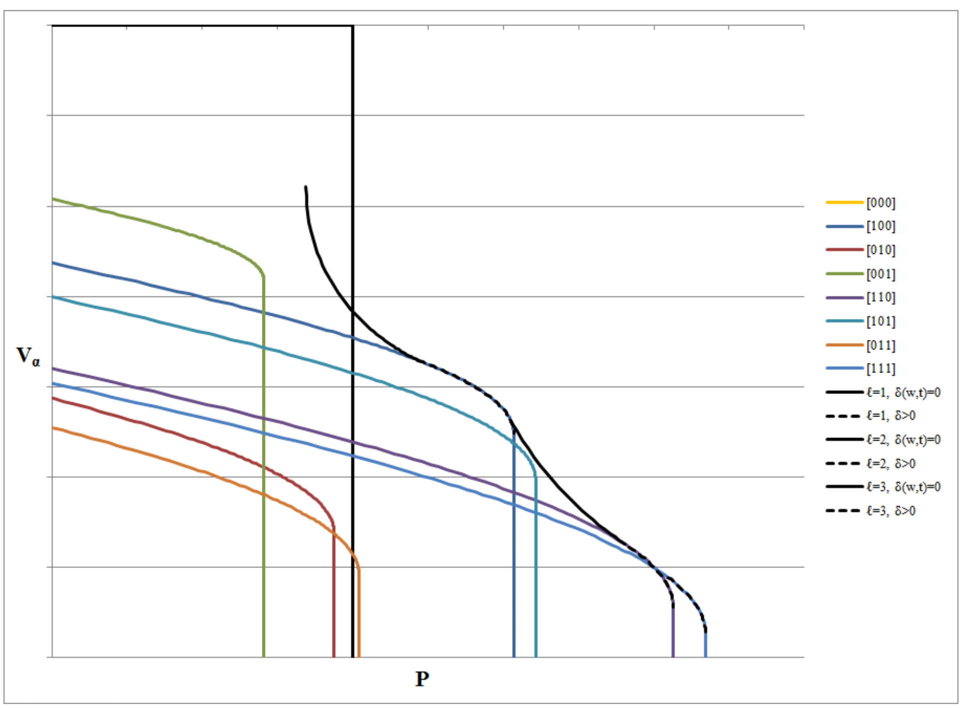

Figure 5.2: Candidate KKT solutions for an expected profit-VaR frontier

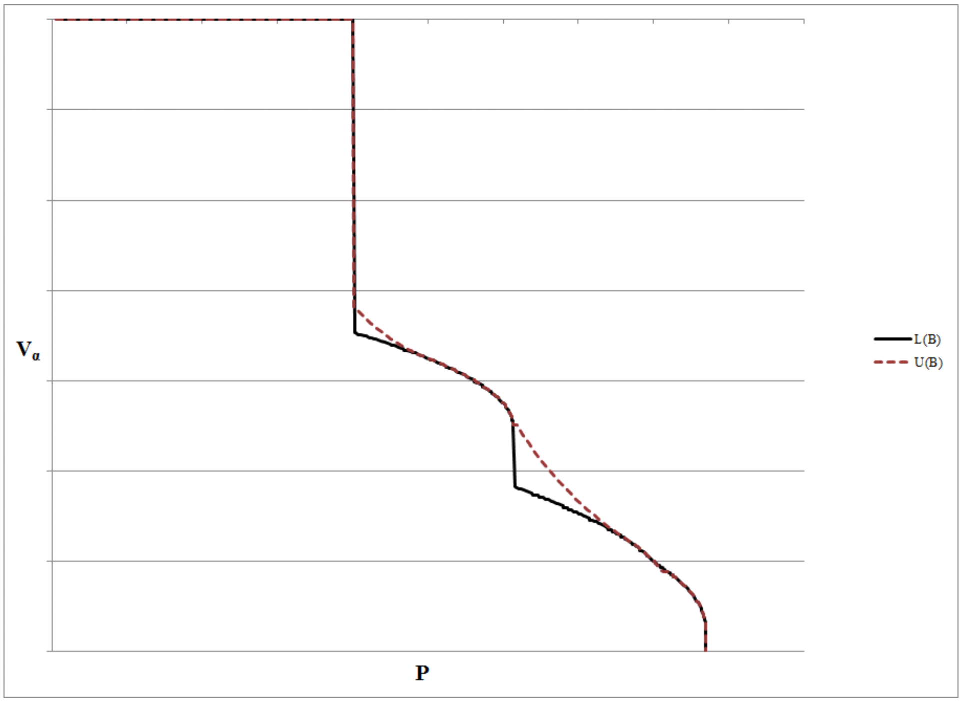

Figure 5.3: $\mathcal{L}(B)$ and $\mathcal{U}(B)$ for an expected profit-VaR frontier

on the $y = \{\mathbf{0}\}$ market selection. As always, we can compute the initial lower bound by taking the concave envelope of the sorted solutions. The initial upper and lower bound are shown by Figure 5.2. From here, we would begin a branch and bound procedure.

Utilizing a branch and bound framework to identify the mean-VaR Pareto efficient frontier is clearly useful for a decision maker solving the selective newsvendor problem who wants to evaluate market selection decisions as risk preferences change. One issue that the finance industry has run into time and time again is an exceeding reliance on using Value-at-Risk or any one risk measure. These risk measures are meant to depict a robust model and therefore it would be of use to compare the results of different risk measures.

Combining the results from this chapter with those from in Chapter IV provides a fuller picture of the risk-averse selective newsvendor. Each risk-measure has its own desirable and undesirable attributes and their use can lead to widely different results. Using both risk-measures in tandem and analyzing their similarities and differences increases the versatility of the model.

We show two simple three market examples with the markets sorted $\frac{\xi_1}{\sigma_1^2} \geq \frac{\xi_2}{\sigma_2^2} \geq \frac{\xi_3}{\sigma_3^2}$ in which the Pareto efficient frontiers vary when using a CVaR measure and a VaR measure. Figure 5.4 is the same problem described in Chapter IV and shows the eight possible market selections under consideration and the four market selections which comprise the Pareto efficient mean-CVaR frontier. Figure 5.6 shows this problem with a VaR measure. The same four market selections that comprised the original frontier remain part of the Pareto efficient mean-VaR frontier, albeit in different proportions.

Next we consider a simple example in which the sorted solutions (less $y = \{\mathbf{0}\}$) comprise the mean-CVaR Pareto efficient frontier as shown by Figure 5.6. When analyzing the mean-VaR Pareto efficient frontier you can see by Figure 5.7 that the sorted solution $y = [1\ 1\ 0]$ is no longer a part of the frontier.

The ability to efficiently identify which market selections appear on the mean-CVaR Pareto efficient frontier as well as which market selections appear on the mean-VaR Pareto efficient frontier provides a more robust depiction risk-aversion for the selective newsvendor.

5.7 Conclusion and Future Research

In this section we studied a risk-averse selective newsvendor who utilized a Value-at-Risk risk measure. We showed that the same set of ranked solutions that contains an optimal market selection solution to the risk-neutral problem and the risk-averse

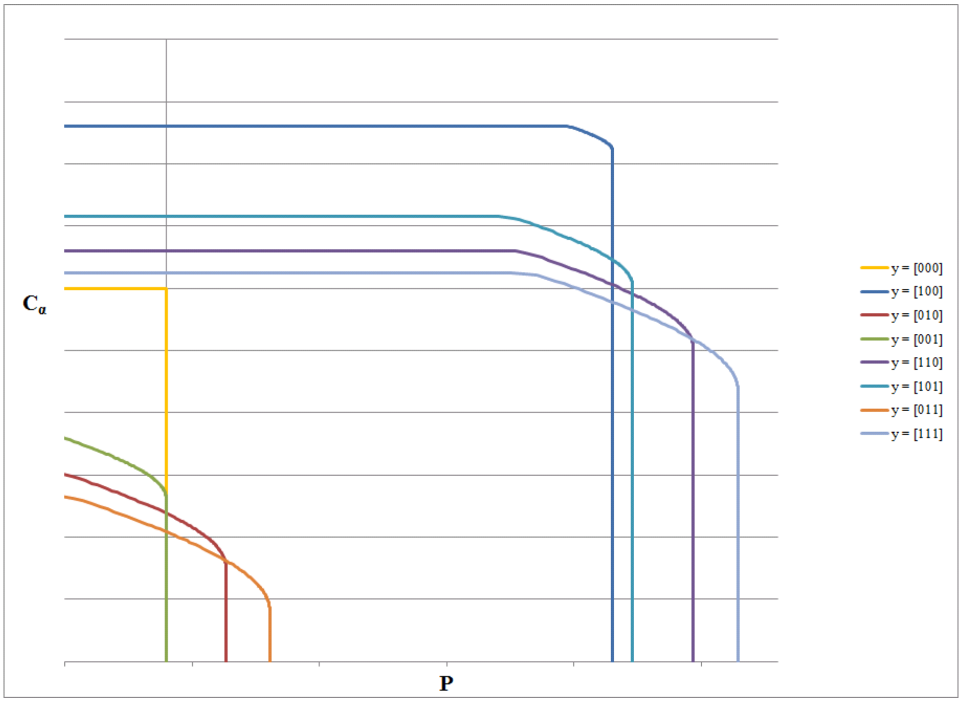

Figure 5.4: Example 1: Mean-CVaR tradeoff curves for a 3 market example

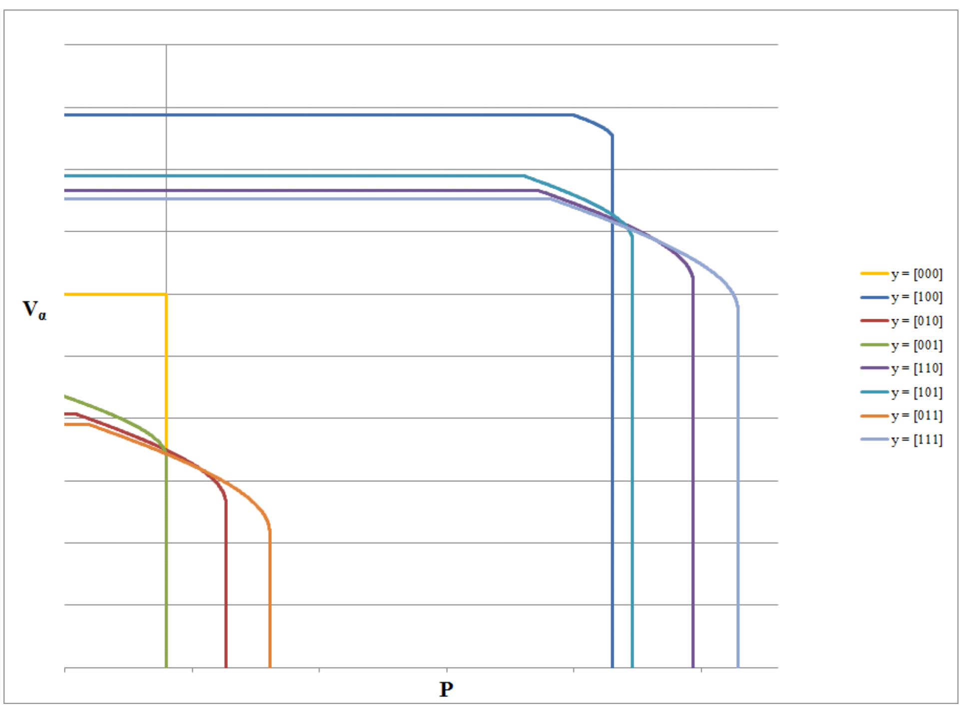

Figure 5.5: Example 1: Mean-VaR tradeoff curves for a 3 market example

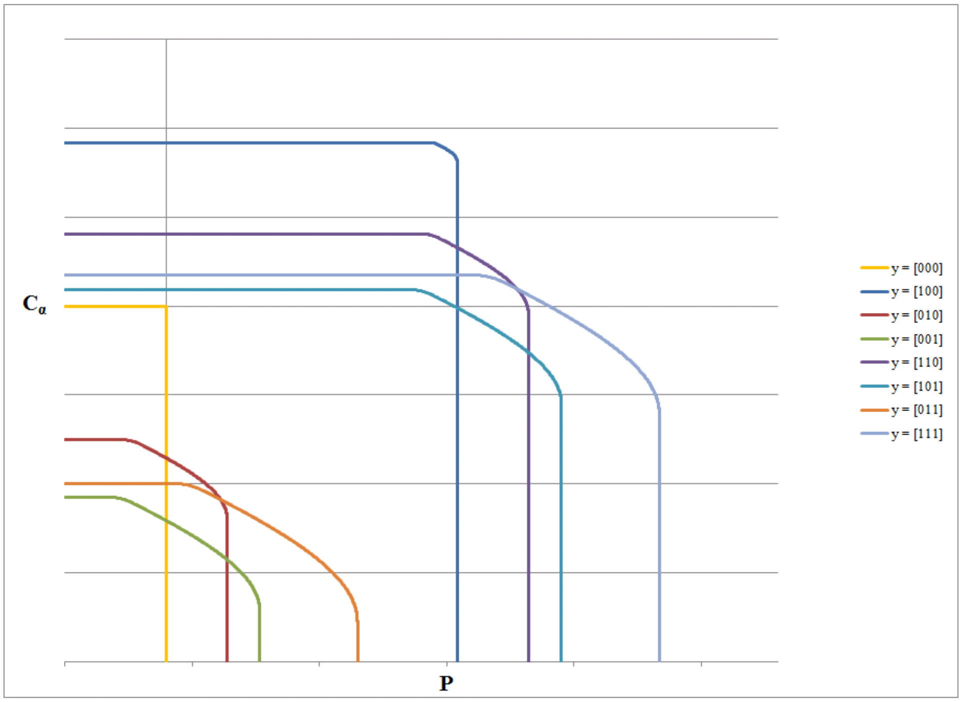

Figure 5.6: Example 2: Mean-CVaR tradeoff curves for a 3 market example

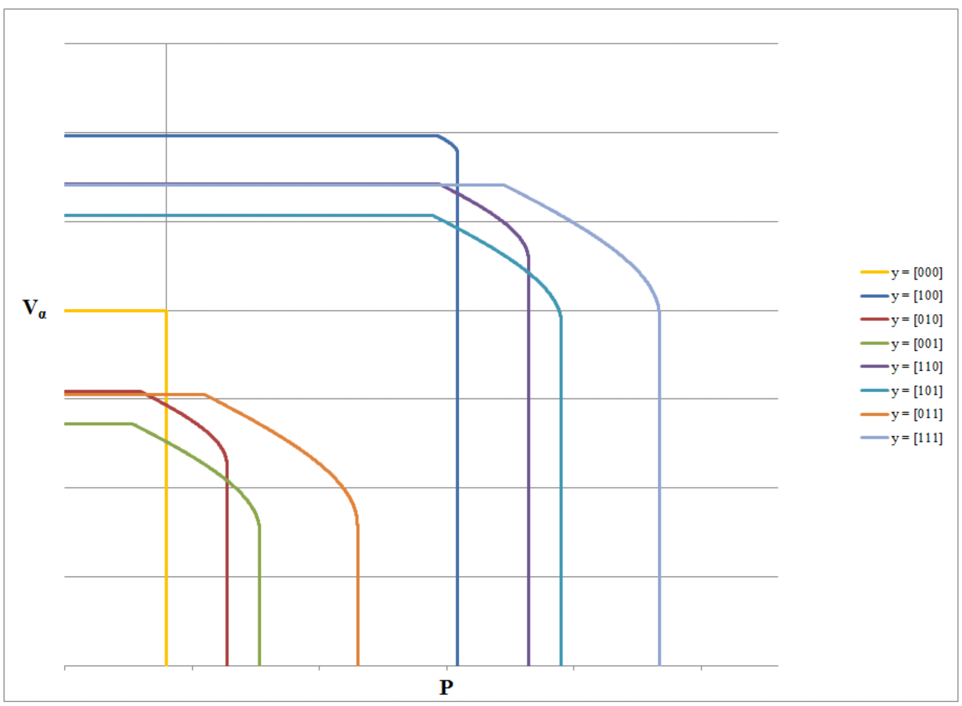

Figure 5.7: Example 2: Mean-VaR tradeoff curves for a 3 market example

problem with Conditional Value-at-Risk also contains the optimal solution for a selective newsvendor using a Value-at-Risk measure. We then analyzed the tradeoff between Value-at-Risk and expected profit. We showed that a branch and bound algorithm like those described in Chapter IV can be used to find the expected profit-VaR Pareto efficient frontier. We reviewed a method to compute the upper bound of the frontier and provided some numerical results. We also discussed the importance of considering multiple risk measures when making market selection decisions and showed two examples where this was important.

CHAPTER VI

Conclusion

6.1 Conclusion

In this dissertation, we analyzed the impact of risk-aversion in the so-called *Selective Newsvendor Problem*. The selective newsvendor incorporates a market selection component into their procurement policy. Prior to our work, Taaffe et al. [36] showed that a risk-neutral selective newsvendors optimal market selection was one of a few candidate solutions. We initially incorporate risk-aversion by examining the selective newsvendor who wants to ensure that demand is met from inventory with for a particular probability. This straightforward extension yields the critical result that for problems of a certain form, if the optimal order quantity can be expressed by a fraction that is independent of the market selection decision, then the optimum can be found among a set of sorted solutions. We extend this central result to some, but not all, risk-averse selective newsvendor problems who use two well-known risk measures: Value-at-Risk and Conditional Value-at-Risk.

Value-at-Risk is commonly used in the finance industry as a way to measure the exposure of financial securities to market risk. Although commonly used, Value-at-Risk does not possess many suitable attributes for mathematical analysis, specifically its lack of convexity and subadditivity. Conditional Value-at-Risk was introduced as a mathematically useful alternative and a byproduct of its optimization results is the calculation of Value-at-Risk. Despite their relationship, each risk measure offers its own pros and cons and we include both to provide a more robust depiction of risk-aversion.

The first problem we considered was a selective newsvendor who maximized a weighted sum of expected profit and CVaR. For a fixed market selection, varying the weighted coefficient yields the Pareto efficient frontier. Although the single newsvendor problem had been studied before, we were able to provide an additional interpretation as well as some new insights. Extending these results to the selective newsvendor, we were able to show that the same collection of candidate solutions that yield the optimal market selection for the risk-averse newsvendor also contain the optimal market selection for this problem. This resulted in a concave upper envelope of the Pareto efficient frontier for the risk-averse *Selective Newsvendor Problem* with supported points given by market selections from the collection of sorted solutions.

Identifying the true Pareto efficient frontier for the risk-averse selective newsvendor was our next goal. The Pareto frontier was not always comprised of only the sorted market selections which required an alternative solution approach. We introduced two branch-and-bound procedures as mechanisms to identify the frontier. Our branch-and-bound methodology involves branching on collections of curves as opposed to the traditional approach in which the branching is done on individual variables. A requirement for the branch-and-bound frameworks is the existence of a bound. To that end, we identified three potential upper bounds. The second upper bound is of course the concave envelope identified by our analysis in Chapter III. The first upper bound was a continuous relaxation of the binary market selection; however, since this upper bound is at best equivalent to the one we do not focus on it. The third upper bound involved a novel approach in which we show that by relaxing the binary market selection requirement we need only look at integral sorted solutions and fractional sorted solutions. Using the KKT conditions to find these solutions and an algorithm to take the outer envelope of these potential points yields an upper bound which is commonly close to the true efficient frontier. Chapter IV also provided some graphical depictions of these results.

In Chapter V, we shifted our focus to the risk-averse *Selective Newsvendor Problem* with a Value-at-Risk constraint. We examined optimal order quantities for single newsvendors using a Value-at-Risk type measure and then extended the results to the selective case. We recognized that we could find the Pareto efficient frontier for the selective newsvendor by employing a branch-and-bound method. Similarly to Chapter IV we derived a method to compute an upper bound for the branch-and-bound framework. We showed that risk-averse selective newsvendors who develop Pareto efficient frontiers using VaR and CVaR risk measures may not result in the same collection of markets which reinforces our decision to improve the robustness of our model by incorporating two risk-measures.

6.2 Future Work

As with all research projects, there are many areas for future work that were beyond the scope of this dissertation. We detail some potential projects here. One limitation of our work is the fact that, in the presence of selection decisions we require market demands to be normally distributed and, more importantly, independent. Relaxing these assumptions will require an entirely different solution approach and is the subject of future research.

Another potential extension is to extend the analysis to a more generic risk measure. We studied Value-at-Risk and Conditional Value-at-Risk but the financial mathematics research is increasing rapidly and new risk measures are consistently being defined and tested. Thus a logical next step is study alternative risk measures.

We focused solely on the market selection component of the supply chain which afforded us the ability to pool inventory and share risk. An area that is rapidly growing in the operations research literature is the supplier selection problem. There is potential to extend our problem to evaluate which suppliers are required to meet the demand in a given market so that we accurately match demand to delivery in a supply chain.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] V. Agrawal and S. Seshadri. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing & Service Operations Management*, 2(4):410–423, 2000.

[2] S. Ahmed, U. Çakmak, and A. Shapiro. Coherent risk measures in inventory problems. *European Journal of Operational Research*, 182(1):226–238, 2007.

[3] K. J. Arrow, T. Harris, and J. Marschak. Optimal inventory policy. *Econometrica*, 19(3):250–272, 1951.

[4] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 3(9):203–228, 1999.

[5] I. Bakal, J. Geunes, and H. E. Romeijn. Market selection decisions for inventory models with price-sensitive demand. *Journal of Global Optimization*, 41:633–657, 2008.

[6] K. Chahar and K. Taaffe. Risk averse demand selection with all-or-nothing orders. *Omega*, 37(5):996 – 1006, 2009.

[7] F. Y. Chen, M. Xu, and G. Zhang. A risk-averse newsvendor model under the CVaR criterion. *Operations Research*, 4(57):1040–1044, 2009.

[8] C. H. Chiu and T. M. Choi. Optimal pricing and stocking decisions for newsvendor problem with value-at-risk consideration. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 40(5):1116–1119, 2010.

[9] S. Choi and A. Ruszczyński. A risk-averse newsvendor with law invariant coherent measures of risk. *Operations Research Letters*, 36(1):77–82, 2008.

[10] F. Y. Edgeworth. The mathematical theory of banking. *Journal of the Royal Statistical Society*, 51(1):113–127, 1888.

[11] L. Eeckhoudt, C. Gollier, and H. Schlesinger. The risk-averse (and prudent) newsboy. *Management Science*, 41(5):786–794, 1995.

[12] A. A. Gaivoronski and G. C. Pflug. Value-at-risk in portfolio optimization: Properties and computational approach. *Journal of Risk*, 7(2):1–31, 2005.

[13] X. Gan, S. P. Sethi, and H. Yan. Channel coordination with a risk-neutral supplier and a downside-risk-averse retailer. *Production and Operations Management*, 14(1):80–89, 2005.

[14] J. Gotoh and Y. Takano. Newsvendor solutions via conditional value-at-risk minimization. *European Journal of Operational Research*, 179(1):80–96, 2007.

[15] B. Ismail and J. Louderback. Optimizing and satisfying in stochastic cost-volume-profit analysis. *Decision Sciences*, 10:205–217, 1979.

[16] W. Jammernegg and P. Kischka. A newsvendor model with service and loss constraints. Technical report, School of Economics and Business Administration, Friedrich-Schiller University, Jena, Germany, 2010.

[17] I. Kabak and A. Schiff. Inventory models and management objectives. *Sloan Management Review*, 19(2):53 – 59, 1978.

[18] M. Khouja. The single-period (news-vendor) problem: literature review and suggestions for future research. *Omega*, 27(5):537 – 553, 1999.

[19] A. H. Land and A. G. Doig. An automatic method of solving discrete programming problems. *Econometrica*, 28(3):497–520, 1960.

[20] H. S. Lau. The newsboy problem under alternative optimization objectives. *The Journal of the Operational Research Society*, 31(6):525–535, 1980.

[21] E. L. Lawler and D. E. Wood. Branch-and-bound methods: A survey. *Operations Research*, 14(4):699–719, 1966.

[22] J. D. C. Little, K. G. Murty, D. W. Sweeney, and C. Karel. An algorithm for the traveling salesman problem. *Operations Research*, 11(6), 1963.

[23] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

[24] R. T. Marler and J. S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26:369–395, 2004.

[25] A. Özler, B. Tan, and F. Karaesmen. Multi-product newsvendor problem with value-at-risk considerations. *International Journal of Production Economics*, 117(2):244 – 255, 2009.

[26] G. C. Pflug. *Probabilistic Constrained Optimization: Methodology and Applications*, chapter Some Remarks on the Value-at-Risk and the Conditional Value-at-Risk. Kluwer Academic Publishers, 2000.

[27] G. C. Pflug. A value-of-information approach to measuring risk in multi-period economic activity. *Journal of Banking and Finance*, 30(2):695–715, 2006. Risk Management and Optimization in Finance.

[28] Y. Qin, R. Wang, A. J. Vakharia, Y. Chen, and M. M. H. Seref. The newsvendor problem: review and directions for future research. *European Journal of Operational Research*, 213(2):361–374, 2011.

[29] RiskMetrics™. Technical report, J.P.Morgan, 1996.

[30] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41, 2000.

[31] W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 1976.

[32] M. E. Schweitzer and G. P. Cachon. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, 46(3):pp. 404–420, 2000.

[33] Z. M. A. Strinka and H. E. Romeijn. Approximation algorithms for selective newsvendor problems. *Working paper*, 2012.

[34] Z. M. A. Strinka, H. E. Romeijn, and J. Wu. Exact and heuristic methods for a class of selective newsvendor problems with normally distributed demands. *Omega*, 2012.

[35] K. Taaffe, K. Chahar, and D. Tirumalasetty. Risk averse selective newsvendor problems. *International Journal of Operational Research*, 3(6):681–703, 2008.

[36] K. Taaffe, J. Geunes, and H. E. Romeijn. Target market selection and marketing effort under uncertainty: The selective newsvendor. *European Journal of Operational Research*, 189(3):987–1003, 2007.

[37] K. Taaffe, H. E. Romeijn, and D. Tirumalasetty. A selective newsvendor approach to order management. *Naval Research Logistics*, 55(8):769–784, 2008.

[38] C. Y. Tsao. Portfolio selection based on the mean-var efficient frontier. *Quantitative Finance*, 10(8):931–945, 2010.

[39] M. Xu. A price-setting newsvendor model under cvar decision criterion with emergency procurement. *Journal of Systems Science and Systems Engineering*, 19(1):85–104, 2010.

[40] M. Xu and F. Y. Chen. Tradeoff between expected reward and Conditional Value-at-Risk criterion in newsvendor models. In *Proceedings of the IEEE International Conference on Industrial Engineering and Engineering Management*, pages 1553–1557, 2007.

[41] M. Xu and J. Li. Comparative analysis of optimal strategies with two purchase modes under different risk-averse criterions. *Wuhan University Journal of Natural Sciences*, 4(14):287–292, 2009.

[42] M. Xu and J. Li. Optimal decisions when balancing expected profit and conditional value-at-risk in newsvendor models. *Journal of Systems Science and Complexity*, 23:1054–1070, 2009.

[43] Q. Zhou and F. Zhao. The loss-averse newsvendor problem with a service requirement. In *Proceedings of the 2010 International Conference on E-Business and E-Government*, ICEE '10, pages 3459–3462, Washington, DC, USA, 2010. IEEE Computer Society.