Data-Driven Revenue Management

by

Joline Ann Villaranda Uichanco

Submitted to the School of Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Computation for Design and Optimization

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2007

Author .....
School of Engineering
September 17, 2007

Certified by .....
Retsef Levi
Robert N. Noyce Career Development Professor, Assistant Professor of Management
Thesis Supervisor

Certified by .....
Georgia Perakis
J. Spencer Standish Career Development Associate Professor of Operations Research
Thesis Supervisor

Accepted by .....
Jaime Peraire
Co-director, Computation for Design and Optimization Program

Data-Driven Revenue Management

by

Joline Ann Villaranda Uichanco

Submitted to the School of Engineering
on September 17, 2007, in partial fulfillment of the
requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract

In this thesis, we consider the classical newsvendor model and various important extensions. We do not assume that the demand distribution is known, rather the only information available is a set of independent samples drawn from the demand distribution. In particular, the variants of the model we consider are: the classical profit-maximization newsvendor model, the risk-averse newsvendor model and the price-setting newsvendor model. If the explicit demand distribution is known, then the exact solutions to these models can be found either analytically or numerically via simulation methods. However, in most real-life settings, the demand distribution is not available, and usually there is only historical demand data from past periods. Thus, data-driven approaches are appealing in solving these problems.

In this thesis, we evaluate the theoretical and empirical performance of nonparametric and parametric approaches for solving the variants of the newsvendor model assuming partial information on the distribution. For the classical profit-maximization newsvendor model and the risk-averse newsvendor model we describe general nonparametric approaches that do not make any prior assumption on the true demand distribution. We extend and significantly improve previous theoretical bounds on the number of samples required to guarantee with high probability that the data-driven approach provides a near-optimal solution. By near-optimal we mean that the approximate solution performs arbitrarily close to the optimal solution that is computed with respect to the true demand distributions. For the price-setting newsvendor problem, we analyze a previously proposed simulation-based approach for a linear-additive demand model, and again derive bounds on the number of samples required to ensure that the simulation-based approach provides a near-optimal solution. We also perform computational experiments to analyze the empirical performance of these data-driven approaches.

Thesis Supervisor: Retsef Levi
Title: Robert N. Noyce Career Development Professor, Assistant Professor of Management

Thesis Supervisor: Georgia Perakis
Title: J. Spencer Standish Career Development Associate Professor of Operations Research

Acknowledgments

I would like to thank first and foremost my supervisors, Prof. Retsef Levi and Prof. Georgia Perakis. Without them pushing me to do my best, this thesis wouldn't have been possible. Our discussions have been extremely enjoyable and I learned a lot about the joys and pains of research from them. Words cannot express the gratitude I feel for the time and effort they put in each week to supervise my work. They have both been my research supervisors, my technical writing teachers and my career advisors. I am truly honored to have worked under them.

I would also like to acknowledge my fellow SMA students. We finally reached the finish line.

My thanks and appreciation go out to my mother and to my brother who are always supporting me. And to Simon for his patience over the last seven months.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis, we address several important variants of the classical newsvendor model, but under the assumption that the explicit demand distribution is not known. Rather, the only information available is a set of independent samples drawn from the true demand distribution. In particular, we consider three variants of the newsvendor model: the classical profit-maximization newsvendor model, the newsvendor model with risk preferences and the price-setting newsvendor model.

The classical profit-maximization model is a problem of matching supply and demand where the supply must be chosen before observing the demand. Demand is stochastic and the market parameters (i.e., price) are exogenous. In the classical model, the firm is risk-neutral and chooses an optimal ordering quantity that maximizes its expected profit with respect to the full demand distribution. An extension of the classical model that incorporates risk preferences of a firm is the risk-averse newsvendor model which has been considered by Bertsimas and Thiele [5]. In this variant, the goal is to maximize the expected profit over a natural set of worst-case scenarios of the demand defined by the event that the demand is less than some specified quantile. Another extension is the price-setting newsvendor model. This model apart from ordering decisions also incorporates pricing decisions. Under this model, the demand is stochastic and a function of price, and the firm decides simultaneously on the price and the supply level.

If the demand distribution is known explicitly, the exact solutions to these models

can be found either analytically or numerically via simulation methods. However, in real-life settings, the explicit demand distribution is not known. Typically, partial information on the demand is available from historical data or from simulations. Data-driven approaches can thus be used to solve the model under partial information. It is of interest to analyze the theoretical and empirical performance of these data-driven approaches.

We discuss a nonparametric approach for solving the classical profit-maximization newsvendor model that makes use of observed samples of the demand without any assumptions on the true demand distribution. Levi, Roundy and Shmoys [24] consider a cost-minimization newsvendor variant and derive bounds on the number of samples required to ensure with high probability a good quality solution to the data-driven approach. We extend the results to the profit-maximization variant, and in fact significantly improve these bounds.

Bertsimas and Thiele [5] introduce a nonparametric data-driven approach based on one-sided trimming to solve the risk-averse newsvendor problem with partial information on the demand. We provide a novel analysis regarding the number of samples required to ensure that the data-driven one-sided trimming method provides a near-optimal solution with high probability.

For the price-setting newsvendor problem, Zhan and Shen [42] propose a simulation-based approach for a particular (i.e. linear-additive) form of the demand-price relation assuming that the parameters are known. The approach uses observed samples of the random component of the demand with only mild assumptions on its true distribution. By imposing additional assumptions on the distribution, we again obtain worst-case bounds on number of samples required for the simulation-based approach to provide a good quality solution with high probability.

Finally, through computational experiments, we verify the empirical performance of these data-driven approaches.

This thesis is structured as follows. In Chapter 2, we consider the classical newsvendor problem under imperfect information. We describe and analyze a nonparametric approach to the classical newsvendor problem. We also describe the

newsvendor problem with risk preferences and establish a connection between accuracy and the sample size of the one-sided trimming approach to the risk-averse newsvendor problem. In Chapter 3, we perform numerical experiments to evaluate the performance of the data-driven approach for the classical newsvendor problem under different concrete scenarios. We also compare its performance against other approaches that solve the newsvendor problem with partial information. In Chapter 4, we describe the price-setting newsvendor problem and analyze the simulation-based proposed by Zhan and Shen [42]. Finally, in Chapter 5, we conduct numerical experiments to evaluate the performance of the simulation-based approach to the price-setting newsvendor problem under different concrete scenarios.

Chapter 2

Data-Driven Approach to the Newsvendor Problem with Exogenous Price

In this chapter, we consider the classical profit-maximizing newsvendor problem with exogenous price (i.e., price is not under the firm's control), but under the assumption that the explicit demand distribution is not known. Instead, we assume that the only information available is a set of independent samples drawn from the true demand distribution. In this chapter, we discuss and analyze a data-driven approach based on solving the Sample Average Approximation (SAA) counterpart. Specifically, we provide a theoretical bound on the number of samples required to achieve a provably near-optimal solution with high probability.

We consider a firm selling a product over a single sales period. A random demand for the product occurs during the sales period. The firm needs to decide before the sales period how many units of the commodity to produce. The actual demand occurs during the sales period and is satisfied as much as possible with the units produced. The firm incurs a cost proportional to the production quantity. The firm sells each unit of product sold for an exogenously determined selling price. Any unmet demand is assumed to be lost. The objective of the firm is to maximize its expected profit.

The classical newsvendor problem is an important foundation for many problems,

especially in revenue management. It can be applied to industries where the firm has no control over the price, but can influence its profit by deciding on a quantity of the product to order. Most industries selling perishable goods fit this criterion since the selling period is too short for the price to be adjusted. The model is also a building block for revenue management problems in the service industry such as airlines and hotels where there is allocation of limited resources to optimize profit.

Under full knowledge of the demand distribution, the classical newsvendor problem has a well-defined solution that balances the expected cost of understocking and the expected cost of overstocking. Specifically, the optimal ordering quantity is a well-specified quantile of the demand distribution that can be computed if one knows the cumulative distribution function of the demand (Porteus [31]). In practice however, the true demand distribution is usually not known. Instead, only partial and imperfect information is available on the demand in the form of demand parameters (e.g., mean, variance, support) or historical demand data.

The newsvendor model and other inventory control and revenue management models with imperfect demand information have been addressed by many researchers (e.g. Savage [34]; Scarf [36]; Gallego and Moon [15]; Bertsimas and Thiele [6]; Levi, Roundy and Shmoys [24]; Perakis and Roels [28]). When the demand distribution is unknown, one may use either a parametric approach or a nonparametric approach. A parametric approach assumes that the true distribution belongs to a parametric family of distributions, but the specific values of the parameters are unknown. On the other hand, a nonparametric approach requires no assumptions regarding the parametric form of the demand distribution.

A popular parametric approach pioneered by Scarf [35] to the newsvendor problem uses a Bayesian procedure to update the belief regarding the uncertainty of the parameter based on observations that are collected over time. More recently, Liyanage and Shanthikumar [25] introduced operational statistics as an approach which, unlike the Bayesian approach, does not assume any prior knowledge on the parameter values. Under this approach, optimization and estimation are done simultaneously.

A robust optimization approach is yet another way to address imperfect informa-

tion on the demand distribution in supply chain models. In this approach, instead of fitting the data to a unique distribution, we allow a family of distributions to which we assume the true demand distribution belongs. A traditional paradigm, called the maximin approach consists of maximizing the worst-case profit over the set of allowed distributions. Scarf [36] and Gallego and Moon [15] derive the maximin order policy over family of distributions having the same mean and variance. The difficulty with this approach is that it can lead to decisions under pessimistic scenarios about the unknown demand. The robust minimax approach (Savage [34]; Perakis and Roels [29]) is an alternative approach which somewhat relaxes this inherent conservativeness by introducing an "uncertainty budget" within which the worst-case scenario is selected. The constraints are either ellipsoidal (see Ben-Tal and Nemirovski [2]) or polyhedral (Bertsimas and Sim [8]; Bertsimas and Thiele [6]). Other recent works on the regret robust approach to other revenue management models include Ball and Queyranne [1], Eren and Maglaras [13] and Perakis and Roels [28]. This approach minimizes the maximum opportunity cost for not making the optimal decision. The robust approach works well if the only information available are market demand parameters such as the mean or variance. Note that these approaches no longer consider the original objective of maximizing the expected profit. Moreover, the resulting solution may be very conservative.

Several nonparametric approaches have been applied to inventory models and the newsvendor problem with partial demand information. The concave adaptive value estimation (CAVE) approach (e.g. Godfrey and Powell [17]) successively approximates the objective cost function with a sequence of piecewise linear functions. The infinitesimal perturbation approach (IPA) is a sampling based stochastic gradient estimation technique that has been used to solve stochastic supply chain models (see Glasserman and Ho [16]). The bootstrap method (Bookbinder and Lordahl [9]) is a nonparametric approach that estimates the newsvendor quantile of the demand distribution. More recently, Huh and Rusmevichientong [20],[19] develop an adaptive algorithm for capacity allocation problems and inventory planning problems with censored demand data.

Another nonparametric approach that utilizes realizations of the demand is the data-driven approach. This approach has the advantage that realizations of the demand are easily obtained from demand data of past periods. From a practical standpoint, the data-driven approach is a simple and natural means to provide an accurate estimate of the optimal order quantity. One such data-driven approach that solves stochastic optimization problems by using only samples of the random variable is the Sample Average Approximation (SAA) approach. In the SAA approach the original objective function, which is the expectation of some random variable, is replaced with the average of the random samples. The SAA approach has been considered for two-stage stochastic optimization models by Kleywegt, Shapiro and Homem-De-Mello [23]. They show that the optimal value of the SAA approach converges to the optimal value of the original problem with probability 1 as the sample size goes to infinity. They also derive bounds on the sample size that ensures with a confidence probability that the difference between the objective value of the SAA solution and the optimal objective value is a certain value. The bounds they derive however, depend on the variance and other properties of the objective function which might be difficult to know if the demand distribution is not known.

Levi, Roundy and Shmoys [24] apply the SAA approach to the cost-minimizing newsvendor problem. Under this variant of the newsvendor model, the firm chooses an order quantity to minimize its expected cost under the presence of an understocking cost b and an overstocking cost h. In this model, it is optimal to order the $\frac{b}{b+h}$ quantile of the demand. Levi, Roundy and Shmoys [24] derive a bound on the sample size that ensures a good quality solution to the SAA approach with high probability. In particular, they show that if the sample size is greater than $\frac{9}{2\alpha^2\omega^2}\log\left(\frac{2}{\delta}\right)$, where $\omega = \frac{\min(b,h)}{b+h}$, then the solution of the SAA counterpart is at most $1+\alpha$ times the optimal solution under full knowledge of the distribution with probability at least $1-\delta$. In contrast to [23], the bound they derive is easy to compute and is free of any assumptions on the demand distribution.

However, applying the bound in [24] for the SAA counterpart of the profit-maximizing variant is not straightforward. In fact, it is usually the case that approx-

imation results for minimization problems do not directly translate to approximation results for the equivalent maximization problem. In this chapter, we introduce a new notion of regret, which we define as the expected cost of the uncertainty in the demand. In particular, it is the difference in the expected profit under the scenario where the firm knows the demand beforehand and the profit under the scenario where the demand is uncertain. By introducing this concept, we managed to leverage the bound in [24] and apply it to the profit-maximization newsvendor problem. Moreover, we managed to improve on the bound in [24] and obtain a significantly stronger theoretical bound that is inversely proportional to ω instead of ω^2. For example, in the cost-minimization newsvendor model, if the service level $\frac{b}{b+h}$ is high, as is typically the case in practice, then we significantly improve (in terms of order of magnitude) the bound on the required sample size that ensures that the SAA approach provides an accurate solution.

Finally, we extend the data-driven approach to newsvendor models that incorporate risk. Bertsimas and Thiele [5] consider a variant of the classical newsvendor problem that incorporates risk preferences of the firm. In the traditional model, the optimal order quantity is determined by maximizing the expected profit with respect to the whole demand distribution. In contrast, they assume that the firm maximizes the expected profit over a natural set of worst-case scenarios of the demand defined by the event that the demand is less than some specified risk parameter. Bertsimas and Thiele [5] propose a data-driven approach to approximate the solution of the risk-averse newsvendor problem. By considering their approach as a variant of the SAA counterpart of the risk-neutral problem with adjusted parameters, we provide a novel analysis regarding the number of samples required to guarantee with a specified confidence level that the regret of the solution to their approach is a near-optimal approximation of the regret under full knowledge of the distribution.

The remainder of this chapter is structured as follows. In Section 2.1, we consider the problem under full knowledge of the demand distribution and review optimality conditions for the solution to this problem. In Section 2.2, we assume imperfect information about the demand in the form of samples from the true demand distribution.

We introduce the SAA counterpart of the problem that provides us an approximate solution to the newsvendor problem. In Section 2.3, we will address the question of finding a sample size that ensures that the SAA counterpart is a "good" approximation of the original newsvendor problem. Finally, in Section 2.4, we find a bound on the sample size for the data-driven approach by Bertsimas and Thiele [5] that ensures an accurate solution to the problem with risk preferences.

2.1 The Profit-Maximizing Newsvendor Model

In this section, we consider the classical maximization newsvendor problem. The classical newsvendor model deals with determining an optimal order level to maximize profit under uncertain demand. Under this model, it is assumed that the explicit demand distribution is known. We provide a mathematical formulation of the profit-maximizing newsvendor problem with exogenous price. The optimal solution is characterized by a balance between the expected cost of understocking and expected cost of overstocking. Porteus [31] and Khouja [22] provide excellent summaries of this problem.

2.1.1 Problem Formulation

First we define the following notation, which will be useful throughout the chapter:

c	unit cost,
p	unit selling price,
q	order or production quantity,
D	random demand with mean μ,
$F(\cdot)$	cumulative distribution of demand.

We define the monotonically non-increasing, left-continuous function

$$\bar{F}(d) = \Pr(D \geq d), \tag{2.1}$$

where notice that we depart from the traditional notation.

A random demand D for a single commodity occurs in a single period. At the beginning of the period, before the demand is observed, the firm decides to order q units of the commodity to satisfy the random demand D. The cost to order a unit of the commodity is c. During the selling period, the actual demand d (the realization of D) is observed. The firm sells the minimum of the demand and the number of units ordered, at a unit selling price p. The profit of the firm is given by

$$\pi(d, q) = p \ \min(d, q) - cq.$$

Since the actual demand is not known when ordering decisions are made, a sensible objective for the firm is to maximize the expected profit. Thus, the problem is equivalent to solving

$$\max_{q \geq 0} \ g(q), \tag{2.2}$$

where

$$g(q) = E[\pi(D, q)] = p \ E[\min(D, q)] - cq$$

is the expected profit of the firm for a given order quantity, q. Unless stated otherwise, the expected value is taken with respect to the true demand distribution. We refer to problem (2.2) as the maximization variant of the newsvendor problem.

A standard assumption in the classical newsvendor problem is that the unit price exceeds the unit cost (i.e., $p > c$). Otherwise, there is no incentive for the firm to order any units of the commodity. Note that, without loss of generality, the salvage cost for each unit of excess inventory is assumed to be equal to zero.

We can think of the cost of uncertainty in the newsvendor model as the difference between the expected profits under the scenario where the demand is known beforehand and when the demand is uncertain. If the expectation of the demand infinite, then any ordering cost will incur negative uncertainty cost. Thus, the following assumption is needed for the newsvendor problem to be well defined.

Assumption 2.1.1 *The distribution for the random demand D satisfies $E[|D|] < \infty$.*

2.1.2 Optimal Solution

In this subsection, we characterize the optimal order quantity of the maximization variant of the newsvendor problem (2.2). Recall that a profit-maximizing firm seeks to find an order quantity that maximizes its expected profit function. That is, the firm wishes to maximize $g(q)$ with respect to q. Note that $g(q)$ is a concave function. Thus, in order to find the order quantity that maximizes $g(q)$, we can utilize standard results in convex optimization. The following properties of a concave function are discussed thoroughly in Bertsekas [7].

Definition 2.1.1 (Bertsekas [7, p. 731]) *Let $f : \Re^n \to \Re$ be a concave function. A vector $v \in \Re^n$ is a* subgradient *of f at a point x_0 if for any $x \in \Re^n$, $f(x) \leq f(x_0) + v^T(x - x_0)$. We denote the set of all subgradients of f at x_0 as the subdifferential, $\partial f(x_0)$.*

Theorem 2.1.1 (Bertsekas [7, pp. 731–732]) *Let $f : \Re^n \to \Re$ be a concave function. Then x_0 is a global maximizer of f if and only if $\mathbf{0} \in \partial f(x_0)$.*

If $f : \Re \to \Re$ is a concave function, one may show that the subdifferential at x_0 is a nonempty closed interval $[f^r(x_0), f^l(x_0)]$, where $f^l(x_0)$ and $f^r(x_0)$ are the left-sided and right-sided derivatives at x_0.

Since $g(q)$ is a concave function, then we can use the optimality condition for a concave function stated in Theorem 2.1.1. Note that the left-sided and right-sided derivatives of $g(q)$ are given by

$$\begin{aligned} g^l(q) &= p\,(1 - \Pr(D < q)) - c, \\ g^r(q) &= p\,(1 - F(q)) - c. \end{aligned}$$

If the cumulative distribution function of D is continuous, note that $g^l(q) = g^r(q) = g'(q)$.

For simplicity, let us define

$$\lambda = 1 - \frac{c}{p}. \tag{2.3}$$

which is referred to in the literature (see Porteus [31]) as the critical fractile. The critical fractile λ balances the cost of understocking (each unit of lost sale is worth $(p-c)$) and the cost of overstocking (the cost of each unit of unused supply is c).

Define

$$q^* = \inf\left\{q \middle| F(q) \geq \lambda\right\}, \tag{2.4}$$

where λ is given by (2.3).

Since $g^l(q^*) \geq 0$ and $g^r(q^*) \leq 0$, then we know that $0 \in \partial g(q^*)$. By Theorem 2.1.1, q^* is the production quantity that optimizes $g(q)$. Thus, the optimal policy is to order up to the λ quantile of demand.

If the cumulative distribution function of the demand is known, then we can explicitly solve for the optimal order quantity, q^*. However, in most real-life scenarios, the true demand distribution is not available. Usually, the information that is available comes from historical demand data. In the following section, we will introduce a data-driven approach based on solving the Sample Average Approximation (SAA) counterpart.

2.2 Sample Average Approximation

One common approach for solving stochastic optimization problems is to solve the Sample Average Approximation (SAA) counterpart. In the original problem, the objective function is the expectation of some random function taken with respect to the true underlying probability distribution. In the SAA counterpart, the objective function is the average value over finitely many independent samples that are drawn from the probability distribution. Shapiro [38] provides an excellent overview of the SAA approach to solving stochastic optimization problems.

Under full knowledge of the demand distribution, the optimal ordering quantity q^* is the λ quantile of the demand distribution. Suppose the demand distribution is unknown and the only information available is a set of independent samples of the demand. We can estimate the expected profit by a function that depends on realizations of the demand. Specifically, for a set of N independent samples of the

demand D, denoted by $d^1, \ldots, d^N$. We approximate the original objective function by

$$\max_{q \geq 0} \hat{g}_N(q) = p \, \frac{1}{N} \sum_{k=1}^{N} \min\left(q, d^k\right) - cq. \tag{2.5}$$

The SAA counterpart can be thought of as a modified newsvendor problem defined with respect to the induced empirical distribution where each of the N samples of the demand occurs with the same probability $\frac{1}{N}$.

Let $\hat{Q}_N$ denote the optimal solution to the SAA counterpart with N samples. Note that $\hat{Q}_N$ is a random variable that is dependent on the specific N samples of D. Since the SAA counterpart is a modified newsvendor problem defined on the empirical distribution, the λ sample quantile

$$\hat{q}_N = \inf\left\{q \middle| \hat{F}_N(q) \geq \lambda\right\} \tag{2.6}$$

is a realization of $\hat{Q}_N$, where $\hat{F}_N(q)$ is the cumulative distribution function of the empirical distribution. The empirical distribution has the following distribution functions:

$$\hat{F}_N(q) = \frac{1}{N} \sum_{k=1}^{N} \mathbf{1}(d^k \leq q), \tag{2.7}$$

$$\hat{\bar{F}}_N(q) = \frac{1}{N} \sum_{k=1}^{N} \mathbf{1}(d^k \geq q). \tag{2.8}$$

where $\mathbf{1}(\cdot)$ is the respective indicator function which is equal to 1 when the argument is true.

2.3 Relative Error of the Regret of the Optimal Solution to the SAA Counterpart

The SAA approach relies on samples of the demand to approximate the expected profit function. A natural question to ask is how many samples are required to

ensure with high probability a good quality solution of the SAA counterpart?

Levi, Roundy and Shmoys [24] address this question for the cost-minimizing newsvendor problem. We define the cost-minimization variant as follows. Consider the case when a firm wishes to choose an order quantity that minimizes its expected cost. Let q be the order quantity. If the number of units ordered exceeds the actual demand, then a per-unit holding cost, h, is incurred for each unit of excess inventory. On the other hand, if the actual demand exceeds the number of units ordered, then a per-unit lost-sales penalty, b, is incurred for each unsatisfied demand. The goal in this minimization variant is to choose an order quantity that minimizes the expected cost, given these costs. Thus, the problem is to solve

$$\min_{q \geq 0} \; C(q) = E\left[h(q-D)^+ + b(D-q)^+\right] \tag{2.9}$$

where $x^+ = \max(x, 0)$. We refer to problem (2.9) as the minimization variant of the newsvendor problem. Suppose q_m^* is the optimal order quantity of the cost-minimizing newsvendor problem. Then it can easily be shown that q_m^* is the $\frac{b}{b+h}$ quantile of the demand distribution [24].

Levi, Roundy and Shmoys [24] apply the SAA approach to solve the cost-minimizing newsvendor problem. The optimal order quantity of the SAA counterpart $\hat{Q}_N^m$ is the $\frac{b}{b+h}$ quantile of the implied empirical demand distribution. They establish how the accuracy of the SAA counterpart relates to the sample size. In particular, they obtain a bound on the number of samples $N = N(b, h, \alpha, \delta)$ required to guarantee that, with probability of at least $1 - \delta$, the optimal solution to the SAA counterpart defined on N samples has an expected cost $C(\hat{Q}_N^m)$ of at most $(1+\alpha)C(q_m^*)$. The following theorem holds.

Theorem 2.3.1 (Levi et al. [24, Theorem 2.2]) *Suppose Assumption 2.1.1 holds. Consider the cost-minimizing newsvendor problem specified by a per-unit holding cost $h > 0$ and a per-unit backlogging penalty $b > 0$. Let $0 < \alpha \leq 1$ be a specified accuracy level and $1 - \delta$ (for $0 < \delta < 1$) be a specified confidence level. Suppose that $N \geq \frac{9}{2\alpha^2\omega^2} \log\left(\frac{2}{\delta}\right)$, where $\omega = \frac{\min(b,h)}{b+h}$. Suppose the SAA counterpart is solved*

with respect to N i.i.d. samples of D. Let $\hat{Q}_N^m$ be the optimal solution to the SAA counterpart. Then, with probability of at least $1-\delta$, the expected cost of $\hat{Q}_N^m$ is at most $1+\alpha$ times the expected cost of an optimal solution q_m^* to the newsvendor problem. In other words, $C(\hat{Q}_N^m) \leq (1+\alpha)C(q_m^*)$ with probability of at least $1-\delta$.

To prove Theorem 2.3.1, Levi, Roundy and Shmoys [24] introduced the following property that relates to the accuracy of a solution as an approximation to the subgradient conditions.

Definition 2.3.1 *Let $\hat{q}_N^m$ be some realization of $\hat{Q}_N^m$ and let $\alpha' > 0$. We will say that $\hat{q}_N^m$ is α'-accurate if $F(\hat{q}_N^m) \geq \frac{b}{b+h} - \alpha'$ and $\bar{F}(\hat{q}_N^m) \geq \frac{h}{b+h} - \alpha'$.*

In what follows we discuss the steps behind the proof of Theorem 2.3.1 in [24]. In the first step it is shown that if $\hat{q}_N^m$ is α'-accurate, with $\alpha' = \frac{\alpha}{3}\frac{\min(b,h)}{b+h}$, then its expected cost is at most $(1+\alpha)C(q_m^*)$. Next, it is shown what is the number of samples $N = N(\alpha', \delta)$ required to guarantee that the $\frac{b}{b+h}$ quantile of the N samples of the random demand is α'-accurate with a confidence probability $1-\delta$. Theorem 2.3.1 follows by combining these two steps. The following two results outline these steps rigorously.

Corollary. 2.3.2 (Levi, Roundy and Shmoys [24, Corollary 2.1]) *For a given accuracy level $\alpha \in (0,1]$, if $\hat{q}_N^m$ is α'-accurate for $\alpha' = \frac{\alpha}{3}\frac{\min(b,h)}{b+h}$, then the cost of $\hat{q}_N^m$ is at most $(1+\alpha)$ times the optimal cost, i.e., $C(\hat{q}_N^m) \leq (1+\alpha)C(q_m^*)$.*

Lemma 2.3.3 (Levi, Roundy and Shmoys [24, Lemma 2.2]) *For each $\alpha' > 0$ and $0 < \delta < 1$, if the number of samples is $N \geq N(\alpha', \delta) = \frac{1}{2}\frac{1}{\alpha'^2}\log\left(\frac{2}{\delta}\right)$, then $\hat{Q}_N^m$, the $\frac{b}{b+h}$ quantile of the sample, is α'-accurate with probability of at least $1-\delta$.*

It is relatively straightforward to see that the definition 2.3.1 of α'-accuracy is related to the deviation between the CDF of the empirical and the true demand distribution, as well as the deviation between the function $\bar{F}$ of the empirical and the true demand distribution. In particular, the proof of Lemma 2.3.3 is based on the well-known Hoeffding's inequality that is used to bound these deviations.

Theorem 2.3.4 (Hoeffding's Inequality [18]). *Let $X^1, \ldots, X^N$ be i.i.d. random variables such that $X^1 \in [\alpha, \beta]$ for some $\alpha < \beta$. Then, for each $\epsilon > 0$, we have* $\Pr\left(\frac{1}{N}\sum_{i=1}^{N} X^i - E[X^1] \geq \epsilon\right) \leq e^{-2\epsilon^2 N/(\beta-\alpha)^2}$.

It must be noted that $(1+\alpha)$-accurate approximations of the objective function of a minimization problem do not necessarily translate to $(1+\alpha)$-accurate approximations of the objective function of a maximization problem. In fact, it is usually the case that approximation results for minimization problems cannot be leveraged to the equivalent maximization problems. Thus, applying Theorem 2.3.1 to the profit-maximization newsvendor problem is highly nontrivial.

It is natural to think of an approximation to the newsvendor problem being "good" if the relative error of its expected profit is small. With this criterion however, a possible complication is that the expected profit of the optimal order quantity is not guaranteed to be well above zero.

To circumvent this problem, we introduce the notion of regret. We define the regret as an objective function of a minimization problem equivalent to the profit-maximization newsvendor problem. If we can show that the approximation evaluated at the regret function has a small relative error, then we can say that it is a "good" approximation.

If the firm knows beforehand that it will face a demand of d, then it will order exactly d units to cover the demand. Ordering any more or less units will give the firm marginal costs. Therefore, under the scenario where the demand is known prior to ordering, the firm makes an expected profit of $(p-c)\mu$. The firm cannot achieve a better profit than this if the demand it faces is unknown. We can think of the difference between $(p-c)\mu$ and the expected profit if the demand is unknown as the cost of the demand uncertainty. Therefore, the profit-maximizing strategy of the firm also minimizes the cost of uncertainty of the demand.

We define the regret to be the difference between the expected profit when the demand is known and the expected profit under uncertain demand. That is, the regret function is

$$\rho(q) = (p-c)\mu - g(q). \tag{2.10}$$

Note that the regret function is minimized by the optimal solution to the newsvendor problem, q^*.

In what follows, we reduce the profit-maximizing newsvendor model into the cost-minimization variant of the newsvendor problem using some suitable transformation. By doing so, we can apply the results derived in Levi, Roundy and Shmoys [24] for the SAA counterpart to the profit-maximizing newsvendor.

Consider the profit-maximizing newsvendor problem (2.2). After rearranging terms, we can rewrite

$$\begin{aligned} p\ E\left[\min(D,q)\right] &= p\ E\left[D+\min(q-D,0)\right] \\ &= p\ E\left[D-(D-q)^+\right] \\ &= p\ \mu - p\ E\left[(D-q)^+\right] \end{aligned}$$

where $E[D]=\mu$, and

$$\begin{aligned} -c\ q - c\ E\left[(D-q)^+\right] + c\mu &= -c\ E\left[q+\max(D-q,0)-D\right] \\ &= -c\ E\left[(q-D)^+\right]. \end{aligned}$$

Thus, the expected profit of the newsvendor problem can be rewritten as

$$\begin{aligned} g(q) &= p\ E\left[\min(D,q)\right] - c\ q \\ &= p\ \left(\mu - E\left[(D-q)^+\right]\right) - c\ E\left[(q-D)^+\right] + c\ E\left[(D-q)^+\right] - c\ \mu \\ &= (p-c)\mu - E\left[c(q-D)^+ + (p-c)(D-q)^+\right]. \end{aligned}$$

This implies that the regret $\rho(q)$ can be expressed as

$$\rho(q) = (p-c)\mu - g(q) = E\left[c(q-D)^+ + (p-c)(D-q)^+\right]. \tag{2.11}$$

Note that equation (2.11) implies that, for every unit of quantity above the actual demand, a cost of c is incurred. On the other hand, for every unit of demand above the order quantity, a cost of $p-c$ is incurred. Thus, the regret represents the cost of

a mismatch between the demand and supply (see Cachon and Terwiesch [10]).

If we introduce the transformation

$$C(q) = \rho(q), \tag{2.12}$$

$$h = c, \tag{2.13}$$

$$b = p - c, \tag{2.14}$$

then we find that minimizing the regret $\rho(q)$ of the profit-maximizing newsvendor is equivalent to the cost-minimization variant considered by Levi, Roundy and Shmoys [24]. Also note that under this transformation, the $\frac{b}{b+h}$ quantile is equal to the λ quantile. Thus, it follows that $\hat{q}_N = \hat{q}_N^m$ and $q^* = q_m^*$.

Therefore, we have the following version of Corollary 2.3.2.

Corollary. 2.3.5 *For a given accuracy level $\alpha \in (0,1]$, if $\hat{q}_N$ is α'-accurate for $\alpha' = \frac{\alpha}{3}\omega$ where $\omega = \min(\lambda, 1-\lambda)$, then the regret of $\hat{q}_N$ is at most $(1+\alpha)$ times the optimal regret, i.e., $\rho(\hat{q}_N) \leq (1+\alpha)\rho(q^*)$.*

Next we shall show that the bound in Theorem 2.3.1 can be significantly improved to depend only on $\frac{1}{\omega}$ instead of $\frac{1}{\omega^2}$. We state this result:

Theorem 2.3.6 *Suppose Assumption 2.1.1 holds. Consider the profit-maximizing newsvendor problem specified by a unit selling price $p > 0$ and a unit production cost $c > 0$ (where $p > c$). Let $0 < \alpha \leq 1$ be a specified accuracy level and $1-\delta$ (for $0 < \delta < 1$) be a specified confidence level. Suppose that $N \geq N(\omega, \alpha, \delta) \equiv \frac{2}{\alpha^2\omega}(9(1-\omega)+4\alpha)\log\left(\frac{2}{\delta}\right)$ where $\omega = \min(\frac{c}{p}, 1-\frac{c}{p})$. Suppose the SAA counterpart is solved with respect to N i.i.d. samples of D. Let $\hat{Q}_N$ be the optimal solution to the SAA counterpart. Then, with probability of at least $1-\delta$, the regret of $\hat{Q}_N$ is at most $1+\alpha$ times the regret of an optimal solution q^* to the newsvendor problem. In other words, $\rho(\hat{Q}_N) \leq (1+\alpha)\rho(q^*)$ with probability of at least $1-\delta$.*

Observe that the bound in Theorem 2.3.6 is inversely proportional to ω. We contrast this to the bound in Levi, Roundy and Shmoys [24] which is inversely proportional to ω^2. If ω is small, then the ratio between the cost of being understocked

and the cost of being overstocked is either very big or very small. In most industries, ω is typically small. Therefore, this result significantly reduces sample size (i.e. in order of magnitude) that is required to ensure that the SAA counterpart provides an accurate approximation of the optimal regret with a certain confidence level.

The idea behind the proof of Theorem 2.3.6 is that we replace and strengthen Lemma 2.3.8. As mentioned previously, the proof of Lemma 2.3.3 makes use of Hoeffding's inequality. Hoeffding's inequality provides an upper bound on the probability that a sum of random variables deviates from its mean using only information about the support of the random variables. We can think of Hoeffding's inequality providing this bound by assuming the worst possible variance with only the knowledge of the range. However, if additional variance information is available, then a tighter bound can be found by using Bernstein's inequality. Below we state Bernstein's inequality:

Theorem 2.3.7 (Bernstein's Inequality [4]). *Let $X_1, \ldots, X_N$ be i.i.d. random variables satisfying* $\Pr(X_1 - E[X_1] \leq d) = 1$. *Let $\sigma^2 = E[X_1^2] - E[X_1]^2$. Then for any $\varepsilon > 0$,*

$$\Pr\left(\frac{1}{N}\sum_{i=1}^{N} X_i - E[X_1] \geq \varepsilon\right) \leq \exp\left(\frac{-N\varepsilon^2}{2\sigma^2 + \frac{2\varepsilon d}{3}}\right).$$

Note that for a fixed q, we can define $X^i = X^i(q) = \mathbf{1}(d^i \leq q)$. Suppose we have N samples of the demand. Therefore, $\hat{F}_N(q) = \frac{1}{N}\sum_{k=1}^{N} X^k$. Note that, $E[X^1] = F(q)$, $\sigma^2 = F(q)(1 - F(q))$, and $X^1 - E[X^1] \leq 1$. Thus, taking $d = 1$ and applying Bernstein's inequality, it follows that for each q

$$\Pr\left(\hat{F}_N(q) - F(q) \geq \varepsilon\right) \leq \exp\left(\frac{-N\varepsilon^2}{2F(q)(1 - F(q)) + \frac{2\varepsilon}{3}}\right). \tag{2.15}$$

Similarly, we can define $Z^i = Z^i(q) = \mathbf{1}(d^i \geq q)$ for a fixed q. For N samples of the demand, we have $\hat{\bar{F}}_N(q) = \frac{1}{N}\sum_{k=1}^{N} Z^k$. Note that $E[Z^1] = \bar{F}(q)$, $\sigma^2 = \bar{F}(q)(1 - \bar{F}(q)$, and $Z^1 - E[Z^1] \leq 1$. Taking $d = 1$, we can apply Bernstein's inequality, to obtain

$$\Pr\left(\hat{\bar{F}}_N(q) - \bar{F}(q) \geq \varepsilon\right) \leq \exp\left(\frac{-N\varepsilon^2}{2\bar{F}(q)(1 - \bar{F}(q)) + \frac{2\varepsilon}{3}}\right). \tag{2.16}$$

By using Bernstein's inequality, we obtain a stronger version of Lemma 2.3.8:

Lemma 2.3.8 *For each $\alpha' > 0$ and $0 < \delta < 1$, if the number of samples is $N \geq N(\omega, \alpha', \delta) = \frac{2}{\alpha'}\left(\frac{1}{\alpha'}\omega(1-\omega) + \frac{4}{3}\right)\log\left(\frac{2}{\delta}\right)$, where $\omega = \min(1-\lambda, \lambda)$, then $\hat{Q}_N$ the λ quantile of the sample is α'-accurate with probability of at least $1-\delta$.*

Proof. Note that each given sample of size N of the demand, $\hat{q}_N$ is the λ quantile of the empirical distribution. Define the event

$$B = \left[F(\hat{q}_N) < \lambda - \alpha'\right].$$

Also define the quantile

$$q_1 = \inf\left\{q \middle| F(q) \geq \lambda - \alpha'\right\}$$

which is illustrated in Figure 2-1(a).

Since $F(\cdot)$ is nondecreasing, then $B = [\hat{q}_N < q_1]$. Consider a monotonically decreasing, nonnegative sequence $\{\tau^k\}$, where $\tau^k \downarrow 0$. Define the events

$$B_k \equiv \left[\hat{F}_N(q_1 - \tau^k) \geq \lambda\right] = \left[\hat{q}_N \leq q_1 - \tau^k\right].$$

Note that since $\hat{F}_N(q_1 - \tau^k) \leq \hat{F}_N(q_1 - \tau^{k+1})$, then it follows that $B_k \subseteq B_{k+1}$. Thus, if $\bar{B}$ is the limiting event of the sequence of events $\{B_k\}$, then $B_k \uparrow \bar{B}$. Thus, $\Pr(B_k) \uparrow \Pr(\bar{B})$. Note also that $B \subseteq \bar{B}$. This implies that $\Pr(B) \leq \Pr(\bar{B})$.

From the definition of q_1, note that for every k, there exists $\varepsilon_k > \alpha'$ such that $F(q_1 - \tau^k) = \lambda - \varepsilon_k < \lambda - \alpha'$. Note that

$$F(q_1 - \tau^k)\left(1 - F(q_1 - \tau^k)\right) < (\lambda - \alpha')(1 - \lambda + \varepsilon_k). \tag{2.17}$$

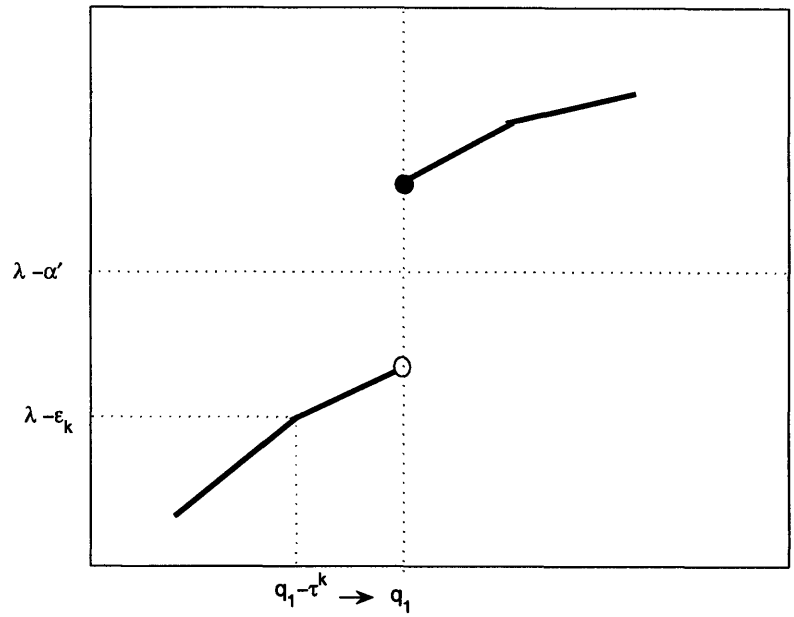

(a) $F(q) = \Pr(D \leq q)$

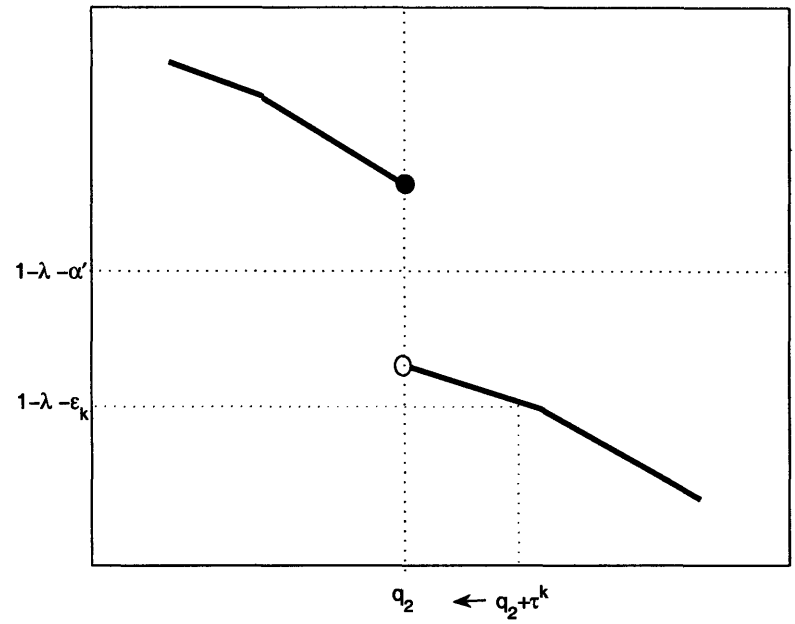

(b) $\bar{F}(q) = \Pr(D \geq q)$

Figure 2-1: Illustration of a cumulative distribution function, F, and distribution function, $\bar{F}$, and the quantiles q_1 and q_2.

Thus, we have

$$\begin{aligned}
\Pr(B_k) &= \Pr\left(\hat{F}_N(q_1 - \tau^k) \geq \lambda\right) \\
&= \Pr(\hat{F}_N(q_1 - \tau^k) - F(q_1 - \tau^k) \geq \varepsilon_k) \\
&\leq \exp\left(\frac{-N{\varepsilon_k}^2/2}{F(q_1 - \tau^k)(1 - F(q_1 - \tau^k)) + \frac{\varepsilon_k}{3}}\right)
\end{aligned}$$

where the last inequality follows from (2.15) by Bernstein's inequality.

From (2.17), it follows that

$$\begin{aligned}
\Pr(B_k) &\leq \exp\left(\frac{-N{\varepsilon_k}^2/2}{(\lambda - \alpha')(1 - \lambda + \varepsilon_k) + \frac{\varepsilon_k}{3}}\right) \\
&= \exp\left(\frac{-N\varepsilon_k/2}{\frac{1}{\varepsilon_k}(\lambda - \alpha')(1 - \lambda) + \lambda - \alpha' + \frac{1}{3}}\right)
\end{aligned}$$

Since $\varepsilon_k > \alpha'$, then

$$\begin{aligned}
\Pr(B_k) &\leq \exp\left(\frac{-N\alpha'/2}{\frac{1}{\alpha'}(\lambda - \alpha')(1 - \lambda) + \lambda - \alpha' + \frac{1}{3}}\right) \\
&= \exp\left(\frac{-N\alpha'/2}{\frac{1}{\alpha'}\lambda(1 - \lambda) - \frac{2}{3} + 2\lambda - \alpha'}\right).
\end{aligned}$$

Also, since $\omega = \min(\lambda, 1 - \lambda)$, then

$$\begin{aligned}
\Pr(B_k) &\leq \exp\left(\frac{-N\alpha'/2}{\frac{1}{\alpha'}\omega(1 - \omega) + \frac{4}{3} - 2\omega - \alpha'}\right) \\
&\leq \exp\left(\frac{-N\alpha'/2}{\frac{1}{\alpha'}\omega(1 - \omega) + \frac{4}{3}}\right).
\end{aligned}$$

Thus, by choosing

$$N \geq \frac{2}{\alpha'}\left(\frac{1}{\alpha'}\omega(1 - \omega) + \frac{4}{3}\right)\log\left(\frac{2}{\delta}\right),$$

we then have $\Pr(B_k) \leq \frac{\delta}{2}$ for all k, which implies that $\Pr(\bar{B}) \leq \frac{\delta}{2}$. Therefore, $\Pr(B) \leq \frac{\delta}{2}$.

Now define the event

$$L = \left[\bar{F}(\hat{q}_N) < 1 - \lambda - \alpha'\right].$$

Also define the quantile

$$q_2 = \sup\left\{q \middle| \bar{F}(q) \geq 1 - \lambda - \alpha'\right\}.$$

which is illustrated in Figure 2-1(b).

Since $\bar{F}(\cdot)$ is non-increasing, then $L = [\hat{q}_N > q_2]$. Define the events

$$L_k \equiv \left[\hat{\bar{F}}_N(q_2 + \tau^k) \geq 1 - \lambda\right] = \left[\hat{q}_N \geq q_2 + \tau^k\right].$$

Note that since $\hat{\bar{F}}_N(q_2 + \tau^k) \leq \hat{\bar{F}}_N(q_2 + \tau^{k+1})$, then it follows that $L_k \subseteq L_{k+1}$. Thus, if $\bar{L}$ is the limiting event of the sequence of events $\{L_k\}$, then $L_k \uparrow \bar{L}$. Thus, $\Pr(L_k) \uparrow \Pr(\bar{L})$. Note also that $L \subseteq \bar{L}$. This implies that $\Pr(L) \leq \Pr(\bar{L})$.

From the definition of q_2, note that for every k, there exists $\varepsilon_k > \alpha'$ such that $\bar{F}(q_2 + \tau^k) = 1 - \lambda - \varepsilon_k < 1 - \lambda - \alpha'$. Note that

$$\bar{F}(q_2 + \tau^k)\left(1 - \bar{F}(q_2 + \tau^k)\right) < (1 - \lambda - \alpha')(\lambda + \varepsilon_k). \tag{2.18}$$

Thus, we have

$$\begin{aligned}
\Pr(L_k) &= \Pr\left(\hat{\bar{F}}_N(q_2 + \tau^k) \geq 1 - \lambda\right) \\
&= \Pr(\hat{\bar{F}}_N(q_2 + \tau^k) - \bar{F}(q_2 + \tau^k) \geq \varepsilon_k) \\
&\leq \exp\left(\frac{-N{\varepsilon_k}^2/2}{\bar{F}(q_2 + \tau^k)(1 - \bar{F}(q_2 + \tau^k)) + \frac{\varepsilon_k}{3}}\right)
\end{aligned}$$

where the last inequality follows from (2.16) by Bernstein's inequality.

From (2.18), it follows that

$$\begin{aligned}\Pr(L_k) &\leq \exp\left(\frac{-N{\varepsilon_k}^2/2}{(1-\lambda-\alpha')(\lambda+\varepsilon_k)+\frac{\varepsilon_k}{3}}\right)\\ &= \exp\left(\frac{-N\varepsilon_k/2}{\frac{1}{\varepsilon_k}\lambda(1-\lambda-\alpha')+1-\lambda-\alpha'+\frac{1}{3}}\right)\end{aligned}$$

Thus, based on similar arguments we can conclude that choosing

$$N \geq \frac{2}{\alpha'}\left(\frac{1}{\alpha'}\omega(1-\omega)+\frac{4}{3}\right)\log\left(\frac{2}{\delta}\right),$$

implies that $\Pr(L_k) \leq \frac{\delta}{2}$ for all k. It follows that $\Pr(L) \leq \Pr(\bar{L}) \leq \frac{\delta}{2}$.

Note that $[\hat{Q}_N$ is not α'-accurate$] = B \cup L$. Thus, $\Pr(\hat{Q}_N$ is not α'-accurate$) \leq \Pr(B) + \Pr(L) \leq \delta$. Thus, for

$$N \geq N(\omega,\alpha,\delta) = \frac{2}{\alpha'}\left(\frac{1}{\alpha'}\omega(1-\omega)+\frac{4}{3}\right)\log\left(\frac{2}{\delta}\right),$$

we have $\Pr(\hat{Q}_N$ is α'-accurate$) \geq 1-\delta$. ∎

We can now proceed to prove Theorem 2.3.6, which makes use of Bernstein's inequality.

Proof of Theorem 2.3.6. Suppose $\alpha' = \frac{\alpha\omega}{3}$. Then, from Lemma 2.3.8, for

$$\begin{aligned}N &\geq \frac{2}{\alpha'}\left(\frac{1}{\alpha'}\omega(1-\omega)+\frac{4}{3}\right)\log\left(\frac{2}{\delta}\right)\\ &= \frac{6}{\alpha\omega}\left(\frac{3}{\alpha}(1-\omega)+\frac{4}{3}\right)\log\left(\frac{2}{\delta}\right)\\ &= \frac{2}{\alpha^2\omega}\left(9(1-\omega)+4\alpha\right)\log\left(\frac{2}{\delta}\right)\end{aligned}$$

it follows that $\hat{Q}_N$ is α'-accurate with probability of at least $1-\delta$. Since if $\hat{q}_N$ is α'-accurate, it follows by Corollary 2.3.5 that $\rho(\hat{q}_N) \leq (1+\alpha)\rho(q^*)$. This concludes the proof of the theorem. ∎

Since we have previously shown that the regret $\rho(q)$ of the profit-maximizing

newsvendor and the objective function of the cost-minimizing newsvendor are equivalent we can extend Theorem 2.3.6 to the cost-minimization variant.

Theorem 2.3.9 *Suppose Assumption 2.1.1 holds. Consider the cost-minimizing newsvendor problem specified by a per-unit holding cost $h > 0$ and a per-unit backlogging penalty $b > 0$. Let $0 < \alpha \leq 1$ be a specified accuracy level and $1-\delta$ (for $0 < \delta < 1$) be a specified confidence level. Suppose that $N \geq N(\omega, \alpha, \delta) \equiv \frac{2}{\alpha^2\omega}(9(1-\omega) + 4\alpha)\log\left(\frac{2}{\delta}\right)$, where $\omega = \frac{\min(b,h)}{b+h}$. Suppose the SAA counterpart is solved with respect to N i.i.d. samples of D. Let $\hat{Q}_N^m$ be the optimal solution to the SAA counterpart. Then, with probability of at least $1-\delta$, the expected cost of $\hat{Q}_N^m$ is at most $1+\alpha$ times the expected cost of an optimal solution q_m^* to the newsvendor problem. In other words, $C(\hat{Q}_N^m) \leq (1+\alpha)C(q_m^*)$ with probability of at least $1-\delta$.*

We observe that the Hoeffding sample size bound is of the order of magnitude proportional to $\frac{1}{\alpha^2\omega^2}\log\left(\frac{2}{\delta}\right)$. On the other hand, the Bernstein sample size bound is of the order of magnitude proportional to $\frac{1}{\alpha^2\omega}\log\left(\frac{2}{\delta}\right)$. Therefore, for small values of ω, we expect that the Bernstein sample size bound is significantly smaller than the Hoeffding bound. Note that if ω is small, the optimal order quantity belongs to extreme quantiles of the demand distribution.

Figures 2-2 and 2-3 illustrate the difference in the bounds provided by Bernstein and Hoeffding for various levels of ω. Notice that if ω is close to zero, then there is a drastic improvement on the sample size required by the Bernstein bound. On the other hand, if ω is close to $\frac{1}{2}$, then the difference between the two bounds is negligible.

Therefore, if ω is significantly small, then we can expect a significant improvement in the order of magnitude of the sample size required by the Bernstein bound. On the other hand, for the ranges where the Hoeffding bound is better, the improvement is small.

In the following lemma, we establish conditions on the parameters that would guarantee that Bernstein's inequality provides a tighter bound on the sample size. For each accuracy level, the range of ω values for which the Bernstein or Hoeffding

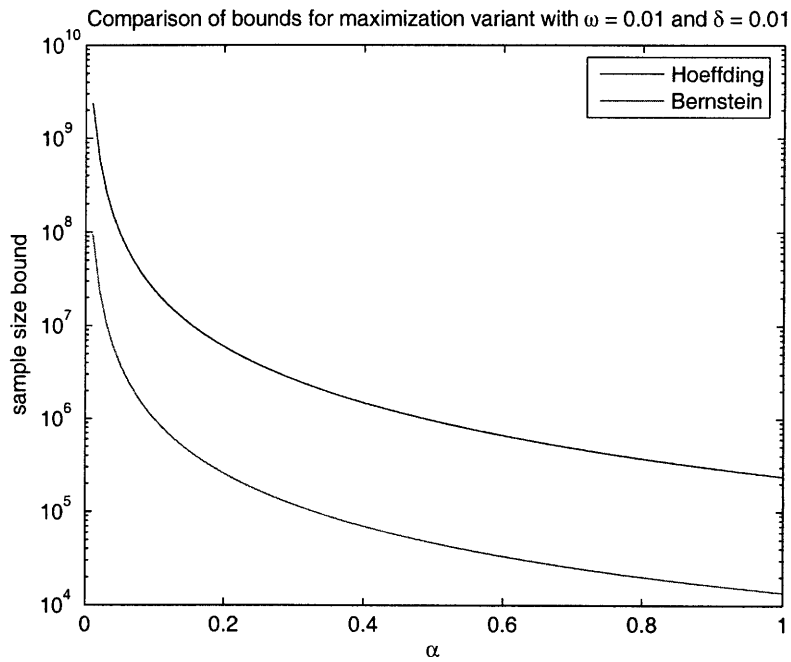

(a) $\omega = 0.01$

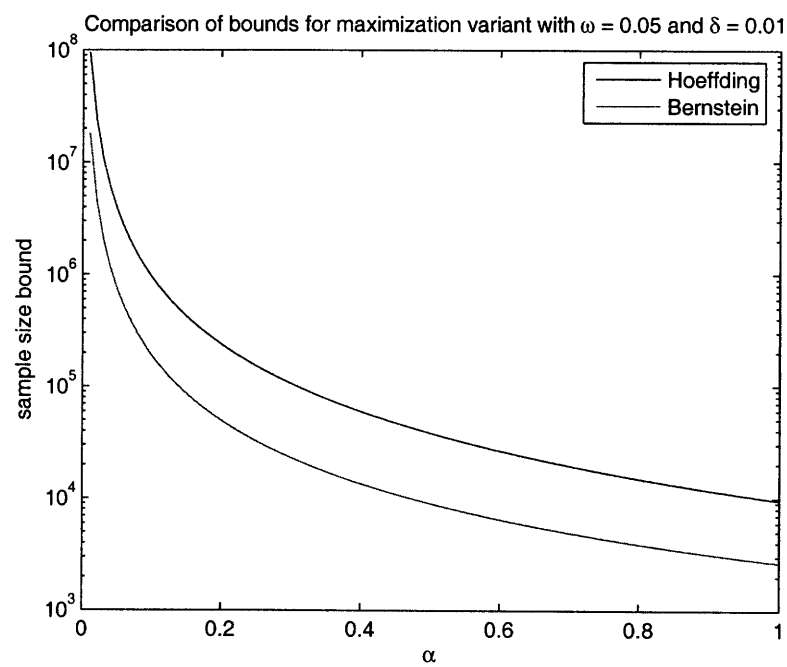

(b) $\omega = 0.05$

Figure 2-2: Comparison of bounds implied by Hoeffding, and Bernstein for the maximization variant of the newsvendor problem as a function of the accuracy rate, for $\omega = 0.01, 0.05$.

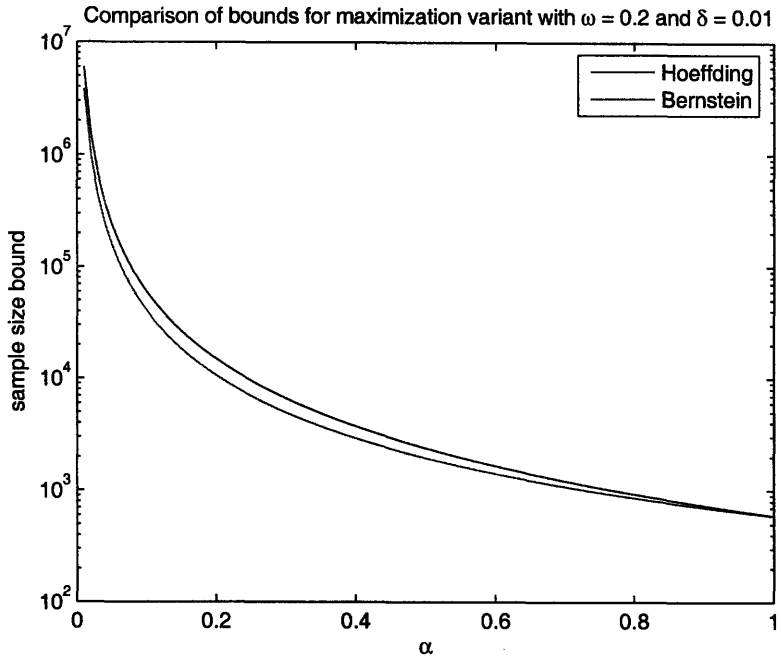

(a) $\omega = 0.2$

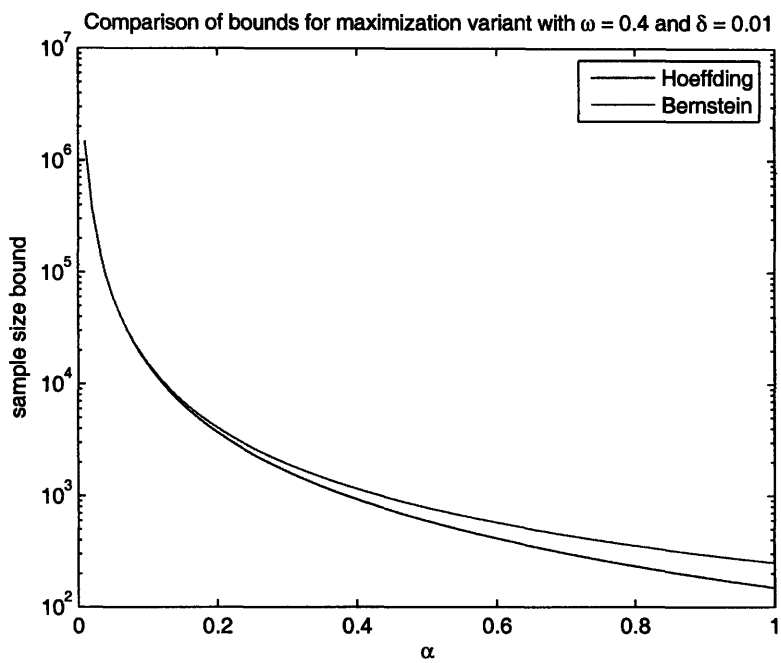

(b) $\omega = 0.4$

Figure 2-3: Comparison of bounds implied by Hoeffding, and Bernstein for the maximization variant of the newsvendor problem as a function of the accuracy rate, for $\omega = 0.2, 0.4$.

bound gives tighter bounds are plotted in Figure 2-4.

Lemma 2.3.10 *For a specified accuracy level $\alpha > 0$ and a specified confidence level $1-\delta$ (for $0 < \delta < 1$), the Bernstein bound in Theorem 2.3.9 is smaller than the Hoeffding bound in Theorem 2.3.1 if and only if $\omega < \frac{1}{18}\left(9 + 4\alpha - \sqrt{4\alpha(4\alpha+18)}\right)$.*

Proof. From Theorems 2.3.1 and 2.3.9, it follows that Bernstein's inequality gives a better bound if and only if

$$9\omega^2 - (9+4\alpha)\omega + \frac{9}{4} \geq 0. \tag{2.19}$$

For a fixed α, the left-hand side of (2.19) is a second degree polynomial of ω. The discriminant of (2.19) is given by

$$\Delta = (9+4\alpha)^2 - 9^2 = 4\alpha\,(4\alpha+18)\,, \tag{2.20}$$

which is always nonnegative. Therefore, if we define

$$\begin{aligned} \omega_{min} &= \frac{1}{18}\left(9+4\alpha-\sqrt{\Delta}\right), \\ \omega_{max} &= \frac{1}{18}\left(9+4\alpha+\sqrt{\Delta}\right), \end{aligned}$$

then the bound provided by Hoeffding's inequality in Theorem 2.3.1 is better than Bernstein in Theorem 2.3.9 only in the range of $\omega \in [\omega_{min}, \omega_{max}]$. Note however that $\Delta > 16\alpha^2$. Thus,

$$\begin{aligned} \omega_{max} &= \frac{1}{18}\left(9+4\alpha+\sqrt{\Delta}\right) \\ &> \frac{1}{18}(8\alpha+9) \geq \frac{1}{2}. \end{aligned}$$

Since $\omega = \min\left(\frac{c}{p}, 1-\frac{c}{p}\right)$, we know that $\omega \leq \frac{1}{2}$. Thus, a necessary and sufficient condition for the Bernstein bound to be better is for $\omega < \omega_{min}$. ∎

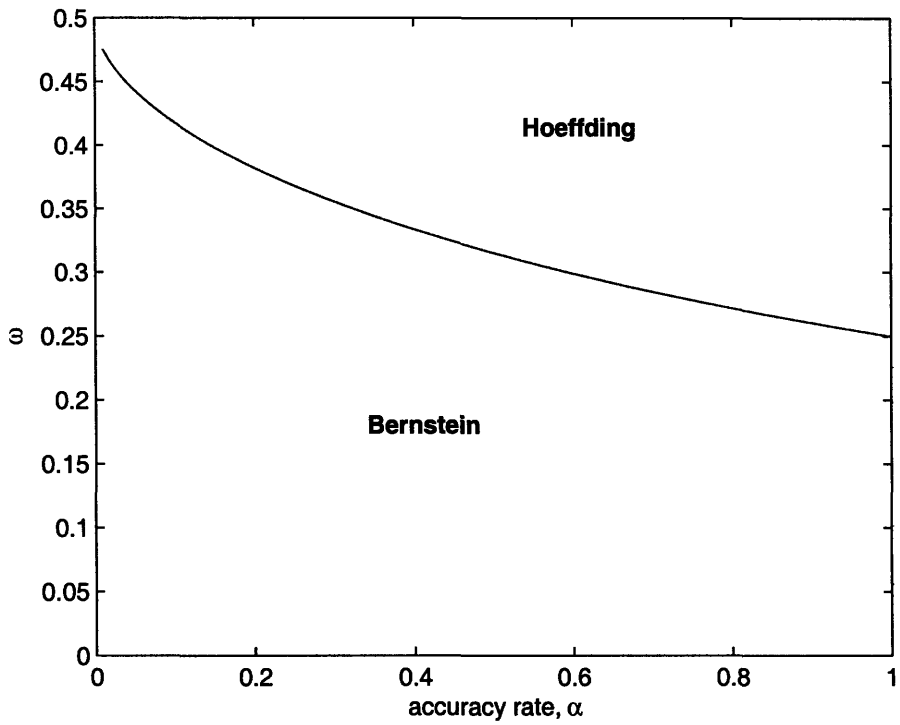

Figure 2-4: Plot of regions of $\omega = \min(1-\lambda, \lambda)$ for which the Hoeffding or the Bernstein bound is tighter, for various accuracy rates, α. The y-axis is in logarithmic scale.

2.4 The Newsvendor Model with Risk Preferences

The assumption of the newsvendor problem discussed in Section 2.1 is that the firm is risk-neutral. However, under an experimental setting, it has been shown by Schweitzer and Cachon [37] that ordering choices systematically deviate from those that maximize expected profit. The firm will often accept a smaller expected profit if it also yields a smaller risk of returns. Bertsimas and Thiele [5] model this risk-preference by introducing a single scalar risk parameter that allows the firm to adjust the trade-off between risk and return. Under a traditional newsvendor model, a risk-neutral firm chooses an ordering quantity that maximizes its expected profit with respect to all of the demand scenarios. In contrast, under the model they propose in [5], a risk-averse firm will choose to protect itself only against a specified set of worst-case scenarios of the demand. In particular, it will choose an ordering quantity that maximizes the expected profit conditioning on the demand being less than some specified risk parameter. They call this one-sided trimming of the demand. Moreover, Bertsimas and Thiele [5] propose a data-driven approach based on one-sided trimming to approximate the optimal order quantity under the newsvendor setting with risk preferences. In this section, we provide a novel analysis regarding the quality of the solution of their data-driven approach. The analysis is based on a variant of the SAA approach under the risk-neutral setting, but with appropriately adjusted parameters.

Suppose β is a risk parameter, where $\beta \in (0,1)$. Let q_β be the $1-\beta$ quantile of the random demand D. That is,

$$q_\beta = \inf\{q | F_D(q) \geq 1-\beta\}.$$

Without loss of generality, we can assume that $F_D(q_\beta) = 1-\beta$. This is because if there is a probability mass on q_β and $F_D(q_\beta) > 1-\beta$, then we can decrease β until this property is satisfied without increasing the risk.

A risk-averse firm will choose a conservative ordering policy that maximizes its expected profit conditional on the event that $D \leq q_\beta$. Define the random variable $\hat{D} = [D|D \leq q_\beta]$. To avoid confusion, all functions and expectations will be subscripted

by the random variable on they are defined. The risk-averse newsvendor problem is

$$\begin{aligned}\max_{q\geq 0} g_{\hat{D}}(q) &= pE_{\hat{D}}[\min(\hat{D},q)] - cq \\ &= pE_D[\min(D,q)|D \leq q_\beta] - cq. \end{aligned} \tag{2.21}$$

Note that the objective function is maximized at the $\lambda = 1 - \frac{c}{p}$ quantile of the random variable $\hat{D}$. In other words, if $\hat{q}^*$ is the solution to the risk-averse newsvendor problem defined by (2.21), then

$$\hat{q}^* = \inf\{q|F_{\hat{D}}(q) \geq \lambda\}.$$

Bertsimas and Thiele [5] propose a variant of the SAA approach to solve the model with risk-preferences. Let N be the total number of observations of the random demand D. Suppose the samples are ranked in an ascending order and labeled as $d^1 \leq \ldots \leq d^N$. The samples are trimmed by taking the $N_\beta = \lceil N(1-\beta)\rceil$ smallest data. The data-driven counterpart of the newsvendor problem with risk-preferences is

$$\max_{q\geq 0}\ p\frac{1}{N_\beta}\sum_{k=1}^{N_\beta}\min(d^k,q) - cq. \tag{2.22}$$

This problem is simply a classical newsvendor problem under the empirical distribution of $\hat{D}$ defined by assigning a probability of $\frac{1}{N_\beta}$ for each sample $d^1,\ldots,d^{N_\beta}$. The solution to the data-driven counterpart is the λ quantile of the empirical distribution. Thus, if we denote $\hat{q}_{N_\beta}$ to be the optimal solution of the data-driven counterpart (2.22), then

$$\hat{q}_{N_\beta} = \inf\left\{q\Big|\frac{1}{N_\beta}\sum_{k=1}^{N_\beta}\mathbf{1}(d^k \leq q) \geq \lambda\right\}. \tag{2.23}$$

Since $\hat{q}_{N_\beta} \leq d^{N_\beta}$, then $\hat{q}_{N_\beta}$ is also the $\frac{\lceil N(1-\beta)\rceil}{N}\lambda$ sample quantile of the empirical distribution of D based on all the N samples. That is, 2.23 is equivalent to

$$\hat{q}_{N_\beta} = \inf\left\{q\Big|\frac{1}{N}\sum_{k=1}^{N}\mathbf{1}(d^k \leq q) \geq \frac{\lceil N(1-\beta)\rceil}{N}\lambda\right\}. \tag{2.24}$$

Our objective in this section is to establish a connection between the sample size and the accuracy (under the regret criterion) of the solution to the data-driven counterpart. Similar to Section 2.3, we define a regret function under the random variable $\hat{D}$ as

$$\rho_{\hat{D}}(q) = (p-c)E[\hat{D}] - g_{\hat{D}}(q),$$

where $E[\hat{D}] = E[D|D \leq q_\beta]$.

Note that this regret function is minimized by the solution to the risk-averse newsvendor problem (2.21). We show that, for an appropriate choice of the sample size, we can guarantee with high probability, that the regret of the data-driven approach has a small relative error with respect to the optimal regret. We establish this relationship in the following theorem:

Theorem 2.4.1 *Consider the risk-averse, profit-maximizing newsvendor problem specified by a risk parameter β, a unit selling price $p > 0$, a unit production cost $c > 0$ (where $p > c$) and a demand distribution D with $E[|D|] < \infty$. Let $0 < \alpha \leq 1$ be a specified accuracy level and $1-\delta$ (for $0 < \delta < 1$) be a specified confidence level. Let $\omega = \min(\frac{c}{p}, 1-\frac{c}{p})$. Suppose that $N \geq N(c,p,\alpha,\delta) \equiv \max(N_1, N_2)$ where*

$$\begin{aligned} N_1 &= \frac{2}{(1-\beta)\alpha^2\omega}\left(\frac{9\lambda}{\omega}(1-(1-\beta)\lambda)+4\alpha\right)\log\left(\frac{3}{\delta}\right), \\ N_2 &= \frac{1}{(1-\lambda)^2(1-\beta)}\Bigg[\lambda(1-\lambda)+2(1-\beta)\left(\beta+\frac{1}{3}(1-\lambda)\right)\log\left(\frac{3}{\delta}\right) \\ &\qquad + \sqrt{2\log\left(\frac{3}{\delta}\right)\beta\lambda(1-\lambda)}\Bigg], \end{aligned}$$

and suppose the SAA counterpart is solved with respect to N i.i.d. samples of D. Let $\hat{Q}_{N_\beta}$ be the optimal solution to the SAA counterpart. Then, with probability of at least $1-\delta$, the regret of $\hat{Q}_{N_\beta}$ is at most $1+\alpha$ times the regret of an optimal solution $\hat{q}^*$ to the risk-averse newsvendor problem. In other words, $\rho_{\hat{D}}(\hat{Q}_{N_\beta}) \leq (1+\alpha)\rho_{\hat{D}}(q^*)$ with probability of at least $1-\delta$.

As opposed to the sample size bound for the risk-neutral newsvendor model, note

that the sample size is the maximum of two numbers, N_1 and N_2. We will show that if the sample size is at least N_1, then it is guaranteed with probability of at least $1 - \frac{2}{3}\delta$ that $\hat{Q}_{N_\beta}$ is $\alpha'(1-\beta)$-accurate with respect to D for a modified newsvendor problem. On the other hand, N_2 guarantees with probability of at least $1 - \frac{\delta}{3}$ that $\hat{Q}_N$ is at most the $1-\beta$ quantile of D.

If the trimming factor β is small, then N_1 inversely proportional to approximately ω, similar to the bound we derived without trimming (risk-neutral firm) in Theorem 2.3.6. We can also see that the bound is inversely proportional to $1-\beta$. Therefore, more samples are required if the firm is more risk-averse (i.e., the risk parameter β is larger).

We now proceed to prove Theorem 2.4.1. We know from Corollary 2.3.5 that if $\hat{q}_{N_\beta}$ is α'-accurate with respect to the distribution of $\hat{D}$ and λ, then the relative error of its regret $\rho_{\hat{D}}(\hat{q}_{N_\beta})$ is bounded. Therefore, in order to prove Theorem 2.4.1 we need to establish a relationship between the sample size N and the α'-accuracy of $\hat{q}_{N_\beta}$ with respect to $\hat{D}$ and λ.

Note that the random variable $\hat{D} = [D|D \leq q_\beta]$ has the following cumulative distribution functions

$$
\begin{aligned}
F_{\hat{D}}(q) &= \begin{cases} \frac{F_D(q)}{1-\beta}, & \text{if } q \leq q_\beta, \\ 1, & \text{if } q > q_\beta. \end{cases} \\
\bar{F}_{\hat{D}}(q) &= \begin{cases} 1 - \frac{1-\bar{F}_D(q)}{1-\beta}, & \text{if } q \leq q_\beta, \\ 0, & \text{if } q > q_\beta. \end{cases}
\end{aligned}
$$

Thus, if we can ensure that $\hat{q}_{N_\beta} \leq q_\beta$, then the conditions for α'-accuracy of $\hat{q}_{N_\beta}$ with respect to $\hat{D}$ are equivalent to

$$
\begin{aligned}
F_D(\hat{q}_{N_\beta}) &\geq (1-\beta)\lambda - (1-\beta)\alpha', \\
\bar{F}_D(\hat{q}_{N_\beta}) &\geq 1 - (1-\beta)\lambda - (1-\beta)\alpha',
\end{aligned} \tag{2.25}
$$

which we can view as conditions for $(1-\beta)\alpha'$-accuracy of $\hat{q}_{N_\beta}$ with respect to D and a modified quantile $(1-\beta)\lambda$. Thus, if we can establish that $\hat{q}_{N_\beta} \leq q_\beta$ and that

conditions (2.25) and (2.26) hold, then it follows that $\hat{q}_{N_\beta}$ is α'-accurate with respect to $\hat{D}$. We can then use Corollary 2.3.5 to prove Theorem 2.4.1.

The following two lemmas establish a sample size that guarantees that these conditions hold. We will be using Bernstein's inequality to prove the lemmas since we have shown in the previous section that it provides an improvement over the sample size implied by Hoeffding's inequality.

Lemma 2.4.2 *For each $\alpha' > 0$, $0 < \beta < 1$ and $0 < \delta_1 < 1$, if the number of samples is $N_1 \geq N(\alpha', \beta, \delta_1) = \frac{2}{(1-\beta)\alpha'}\left(\frac{1}{\alpha'}\lambda(1-(1-\beta)\lambda) + \frac{4}{3}\right)\log\left(\frac{2}{\delta_1}\right)$, then $F_D(\hat{Q}_{N_\beta}) \geq (1-\beta)\lambda - (1-\beta)\alpha'$ and $\bar{F}_D(\hat{Q}_{N_\beta}) \geq 1 - (1-\beta)\lambda - (1-\beta)\alpha'$ with probability of at least $1-\delta_1$, where $\hat{Q}_{N_\beta}$ is the $\frac{\lceil N(1-\beta)\rceil}{N}\lambda$ sample quantile.*

Proof. For ease of notation, let $\eta(N) = \frac{\lceil N(1-\beta)\rceil}{N}$. Note that $\eta(N) \geq 1-\beta$ for every N and $\eta(N) \to 1-\beta$ as $N \to \infty$. Define the event $B = [F_D(\hat{q}_{N_\beta}) < (1-\beta)\lambda - \alpha'(1-\beta)]$. Also define the quantiles of the distribution of D

$$\tilde{q}_1 = \inf\left\{q \middle| F_D(q) \geq \eta(N)\lambda - \eta(N)\alpha'\right\},$$

and

$$q_1 = \inf\left\{q \middle| F_D(q) \geq (1-\beta)\lambda - (1-\beta)\alpha'\right\}.$$

Thus, $B = [\hat{q}_{N_\beta} < q_1]$. Since $\tilde{q}_1 \geq q_1$, then $B \subseteq [\hat{q}_{N_\beta} < \tilde{q}_1]$. Consider a monotonically decreasing, nonnegative sequence $\{\tau^k\}$ where $\tau^k \downarrow 0$. Define the events

$$B_k = \left[\hat{F}_N(\tilde{q}_1 - \tau^k) \geq \eta(N)\lambda\right] = [\hat{q}_{N_\beta} \leq \tilde{q}_1 - \tau^k],$$

where $\hat{F}_N(\cdot)$ is the cumulative distribution function of the empirical distribution defined on N samples.

Since $\hat{F}_N(\tilde{q}_1 - \tau^k) \leq \hat{F}_N(\tilde{q}_1 - \tau^{k+1})$, then it follows that $B_k \subseteq B_{k+1}$. Thus, if $\bar{B}$ is the limiting event of the sequence of events $\{B_k\}$, then $B_k \uparrow \bar{B}$. Thus, $\Pr(B_k) \uparrow \Pr(\bar{B})$. Note also that $[\hat{q}_{N_\beta} < \tilde{q}_1] \subseteq \bar{B}$. This implies that $\Pr(B) \leq \Pr(\bar{B})$.

From the definition of $\tilde{q}_1$, note that for every k, there exists $\varepsilon_k > \alpha'\eta(N)$ such that $F_D(\tilde{q}_1 - \tau^k) = \eta(N)\lambda - \varepsilon_k$.

Note that

$$F_D(\tilde{q}_1 - \tau^k)(1 - F_D(\tilde{q}_1 - \tau^k)) < (\eta(N)\lambda - \eta(N)\alpha')\,(1 - \eta(N)\lambda + \varepsilon_k)\,. \tag{2.26}$$

We have by Bernstein's inequality

$$\begin{aligned}
\Pr(B_k) &= \Pr(\hat{F}_N(\tilde{q}_1 - \tau^k) \geq \eta(N)\lambda) \\
&= \Pr(\hat{F}_N(\tilde{q}_1 - \tau^k) - F_D(\tilde{q}_1 - \tau^k) \geq \varepsilon_k) \\
&\leq \exp\left(\frac{-N{\varepsilon_k}^2/2}{F_D(\tilde{q}_1 - \tau^k)(1 - F_D(\tilde{q}_1 - \tau^k)) + \frac{\varepsilon_k}{3}}\right).
\end{aligned}$$

From equation (2.26), it follows that

$$\begin{aligned}
\Pr(B_k) &\leq \exp\left(\frac{-N\varepsilon_k^2/2}{(\eta(N)\lambda - \eta(N)\alpha')\,(1 - \eta(N)\lambda + \varepsilon_k) + \frac{\varepsilon_k}{3}}\right) \\
&= \exp\left(\frac{-N\varepsilon_k/2}{\frac{1}{\varepsilon_k}\eta(N)(\lambda - \alpha')(1 - \eta(N)\lambda) + \eta(N)\lambda - \eta(N)\alpha' + \frac{1}{3}}\right).
\end{aligned}$$

Since we know that $\varepsilon_k > \eta(N)\alpha'$, then

$$\begin{aligned}
\Pr(B_k) &\leq \exp\left(\frac{-N\eta(N)\alpha'/2}{\frac{1}{\alpha'}(\lambda - \alpha')(1 - \eta(N)\lambda) + \eta(N)\lambda - \eta(N)\alpha' + \frac{1}{3}}\right) \\
&= \exp\left(\frac{-N\eta(N)\alpha'/2}{\frac{1}{\alpha'}\lambda(1 - \eta(N)\lambda) - \frac{2}{3} + 2\eta(N)\lambda - \eta(N)\alpha'}\right).
\end{aligned}$$

Also since $1 - \beta \leq \eta(N)$ and $\eta(N)\lambda \leq 1 - \min(\eta(N)\lambda, 1 - \eta(N)\lambda)$, then, it follows that

$$\begin{aligned}
\Pr(B_k) &\leq \exp\left(\frac{-N(1-\beta)\alpha'/2}{\frac{1}{\alpha'}\lambda(1 - (1-\beta)\lambda) + \frac{4}{3} - 2\min(\eta(N)\lambda, 1 - \eta(N)\lambda) - \eta(N)\alpha'}\right) \\
&\leq \exp\left(\frac{-N(1-\beta)\alpha'/2}{\frac{1}{\alpha'}\lambda(1 - (1-\beta)\lambda) + \frac{4}{3}}\right).
\end{aligned}$$

Thus, by choosing

$$N \geq \frac{2}{(1-\beta)\alpha'}\left(\frac{1}{\alpha'}\lambda(1 - (1-\beta)\lambda) + \frac{4}{3}\right)\log\left(\frac{2}{\delta_1}\right),$$

we have $\Pr(B_k) \leq \frac{\delta_1}{2}$. This implies that $\Pr(B) \leq \Pr(\bar{B}) \leq \frac{\delta_1}{2}$.

Similarly, define the event $L = [\bar{F}_D(\hat{q}_{N_\beta}) < 1 - (1-\beta)\lambda - \alpha'(1-\beta)]$. Also define the quantiles of the distribution of D

$$\tilde{q}_2 = \sup\left\{q \middle| \bar{F}_D(q) \geq 1 - \eta(N)\lambda - \eta(N)\alpha')\right\},$$

and

$$q_2 = \sup\left\{q \middle| \bar{F}_D(q) \geq 1 - (1-\beta)\lambda - (1-\beta)\alpha'\right\}.$$

Thus, $L = [\hat{q}_{N_\beta} > q_2]$. Since $\tilde{q}_2 \leq q_2$, then $L \subseteq [\hat{q}_{N_\beta} > \tilde{q}_2]$. Define the events

$$L_k = \left[\hat{\bar{F}}_N(\tilde{q}_2 + \tau^k) \geq 1 - \eta(N)\lambda\right] = [\hat{q}_{N_\beta} \geq \tilde{q}_2 + \tau^k],$$

where $\hat{\bar{F}}_N(\cdot)$ is the probability $\Pr(D \geq q)$ of the empirical distribution defined on N samples.

Since $\hat{\bar{F}}_N(\tilde{q}_2 + \tau^k) \leq \hat{\bar{F}}_N(\tilde{q}_2 + \tau^{k+1})$, then it follows that $L_k \subseteq L_{k+1}$. Thus, if $\bar{L}$ is the limiting event of the sequence of events $\{L_k\}$, then $L_k \uparrow \bar{L}$. Thus, $\Pr(L_k) \uparrow \Pr(\bar{L})$. Note also that $[\hat{q}_{N_\beta} > \tilde{q}_2] \subseteq \bar{L}$. This implies that $\Pr(L) \leq \Pr(\bar{L})$.

From the definition of $\tilde{q}_2$, note that for every k, there exists $\varepsilon_k > \alpha'\eta(N)$ such that $\bar{F}_D(\tilde{q}_2 + \tau^k) = 1 - \eta(N)\lambda - \varepsilon_k$.

Note that

$$\bar{F}_D(\tilde{q}_2 + \tau^k)(1 - \bar{F}_D(\tilde{q}_2 + \tau^k)) < (1 - \eta(N)\lambda - \eta(N)\alpha')\,(\eta(N)\lambda + \varepsilon_k)\,. \tag{2.27}$$

We have by Bernstein's inequality

$$\begin{aligned}
\Pr(L_k) &= \Pr(\hat{\bar{F}}_N(\tilde{q}_2 + \tau^k) \geq 1 - \eta(N)\lambda) \\
&= \Pr(\hat{\bar{F}}_N(\tilde{q}_2 + \tau^k) - \bar{F}_D(\tilde{q}_2 + \tau^k) \geq \varepsilon_k) \\
&\leq \exp\left(\frac{-N{\varepsilon_k}^2/2}{\bar{F}_D(\tilde{q}_2 + \tau^k)(1 - \bar{F}_D(\tilde{q}_2 + \tau^k)) + \frac{\varepsilon_k}{3}}\right).
\end{aligned}$$

From equation (2.27), it follows that

$$\begin{aligned}\Pr(L_k) &\leq \exp\left(\frac{-N\varepsilon_k^2/2}{(1-\eta(N)\lambda-\eta(N)\alpha')\left(\eta(N)\lambda+\varepsilon_k\right)+\frac{\varepsilon_k}{3}}\right)\\ &= \exp\left(\frac{-N\varepsilon_k/2}{\frac{1}{\varepsilon_k}\eta(N)\lambda(1-\eta(N)\lambda-\eta(N)\alpha')+\frac{4}{3}-\eta(N)\lambda-\eta(N)\alpha'}\right).\end{aligned}$$

Since we know that $\varepsilon_k > \eta(N)\alpha'$, then

$$\begin{aligned}\Pr(L_k) &\leq \exp\left(\frac{-N\eta(N)\alpha'/2}{\frac{1}{\alpha'}\lambda(1-\eta(N)\lambda-\eta(N)\alpha')+\frac{4}{3}-\eta(N)\lambda-\eta(N)\alpha'}\right)\\ &= \exp\left(\frac{-N\eta(N)\alpha'/2}{\frac{1}{\alpha'}\lambda(1-\eta(N)\lambda)+\frac{4}{3}-2\eta(N)\lambda-\eta(N)\alpha'}\right).\end{aligned}$$

Thus, it follows that

$$\begin{aligned}\Pr(L_k) &\leq \exp\left(\frac{-N(1-\beta)\alpha'/2}{\frac{1}{\alpha'}\lambda(1-(1-\beta)\lambda)+\frac{4}{3}-2\min(\eta(N)\lambda,1-\eta(N)\lambda)-\eta(N)\alpha'}\right)\\ &\leq \exp\left(\frac{-N(1-\beta)\alpha'/2}{\frac{1}{\alpha'}\lambda(1-(1-\beta)\lambda)+\frac{4}{3}}\right).\end{aligned}$$

Thus, by choosing

$$N \geq \frac{2}{(1-\beta)\alpha'}\left(\frac{1}{\alpha'}\lambda(1-(1-\beta)\lambda)+\frac{4}{3}\right)\log\left(\frac{2}{\delta_1}\right),$$

we have $\Pr(L_k) \leq \frac{\delta_1}{2}$. This implies that $\Pr(L) \leq \Pr(\bar{L}) \leq \frac{\delta_1}{2}$.

Therefore, for $N \geq N(\alpha,\beta,\delta_1) = \frac{2}{(1-\beta)\alpha'}\left(\frac{1}{\alpha'}\lambda(1-(1-\beta)\lambda)+\frac{4}{3}\right)\log\left(\frac{2}{\delta_1}\right)$, we have the probability of our desired event

$$\Pr(B^c \cap L^c) = 1-\Pr(B\cup L) \geq 1-\Pr(B)-\Pr(L) \geq 1-\delta_1. \quad \blacksquare$$

Lemma 2.4.3 *For each $0 < \delta_2 < 1$, if the number of samples is $N \geq N(\beta, \delta_2)$, where*

$$N_2 = N(\beta, \delta_2) = \frac{1}{(1-\lambda)^2(1-\beta)}\left[\lambda(1-\lambda) + 2(1-\beta)\left(\beta + \frac{1}{3}(1-\lambda)\right)\log\left(\frac{1}{\delta_2}\right) + \sqrt{2\log\left(\frac{1}{\delta_2}\right)\beta\lambda(1-\lambda)}\right],$$

then $\hat{Q}_N$, the $\frac{\lceil N(1-\beta)\rceil}{N}\lambda$ sample quantile, is at most q_β with probability of at least $1-\delta_2$.

Proof. Note that for every N, there exists some $\varepsilon_N \in [0,1)$ such that $\lceil N(1-\beta)\rceil = N(1-\beta)+\varepsilon_N$. Since $\hat{q}_{N_\beta}$ is the $\frac{\lceil N(1-\beta)\rceil}{N}\lambda$ quantile of the empirical distribution defined with N samples, then

$$\begin{aligned}\Pr(\hat{q}_{N_\beta} > q_\beta) &= \Pr\left(\hat{F}_N(q_\beta) < \frac{\lceil N(1-\beta)\rceil}{N}\lambda\right) \\ &= \Pr\left(\hat{F}_N(q_\beta) < \left(1-\beta+\frac{\varepsilon_N}{N}\right)\lambda\right).\end{aligned}$$

Also, since $F_D(q_\beta) \geq 1-\beta$, then

$$\Pr(\hat{q}_{N_\beta} > q_\beta) \leq \Pr\left(F_D(q_\beta) - \hat{F}_N(q_\beta) > (1-\lambda)(1-\beta) - \frac{\varepsilon_N}{N}\lambda\right).$$

We also know that $\varepsilon_N < 1$, so it follows that

$$\Pr(\hat{q}_{N_\beta} > q_\beta) \leq \Pr\left(F_D(q_\beta) - \hat{F}_N(q_\beta) > (1-\lambda)(1-\beta) - \frac{\lambda}{N}\right).$$

If $N > \frac{\lambda}{(1-\lambda)(1-\beta)}$, then we can apply Bernstein's inequality. For now, let us assume that this is true. By Bernstein's inequality, we have

$$\Pr(\hat{q}_{N_\beta} > q_\beta) \leq \exp\left(\frac{-\frac{N}{2}\left((1-\lambda)(1-\beta) - \frac{\lambda}{N}\right)^2}{\beta(1-\beta) + \frac{(1-\lambda)(1-\beta)-\frac{\lambda}{N}}{3}}\right).$$

For $\Pr(\hat{q}_{N_\beta} > q_\beta) \leq \delta_2$, we must have

$$N\left((1-\lambda)(1-\beta)-\frac{\lambda}{N}\right)^2 \geq 2\left(\beta(1-\beta)+\frac{(1-\lambda)(1-\beta)-\frac{\lambda}{N}}{3}\right)\log\left(\frac{1}{\delta_2}\right).$$

Rearranging terms, we get a second degree polynomial in N

$$\begin{aligned}(1-\lambda)^2(1-\beta)^2N^2 \quad - \quad & \left[2\lambda(1-\lambda)(1-\beta)+2(1-\beta)\log\left(\frac{1}{\delta_2}\right)\left(\beta+\frac{1}{3}(1-\lambda)\right)\right]N \\ & +\lambda^2+\frac{2}{3}\lambda\log\left(\frac{1}{\delta_2}\right) \geq 0.\end{aligned}$$

where we know if N_{max} and N_{min} are the roots of the second-degree polynomial specified by the left-hand side, then the inequality holds if $N \geq N_{max}$.

The discriminant of the quadratic inequality is

$$\begin{aligned}\Delta &= \left[2\lambda(1-\lambda)(1-\beta)+2(1-\beta)\log\left(\frac{1}{\delta_2}\right)\left(\beta+\frac{1}{3}(1-\lambda)\right)\right]^2 \\ &\quad -4(1-\lambda)^2(1-\beta)^2\left(\lambda^2+\frac{2}{3}\lambda\log\left(\frac{1}{\delta_2}\right)\right) \\ &= 4(1-\beta)^2\log\left(\frac{1}{\delta_2}\right)\left(\left(\beta+\frac{1}{3}(1-\lambda)\right)^2\log\left(\frac{1}{\delta_2}\right)+2\lambda(1-\lambda)\beta\right),\end{aligned}$$

which is always nonnegative.

To simplify the terms, note that

$$\begin{aligned}\sqrt{\Delta} &\leq \sqrt{4(1-\beta)^2\left(\beta+\frac{1}{3}(1-\lambda)\right)^2\log\left(\frac{1}{\delta_2}\right)^2}+\sqrt{8(1-\beta)^2\log\left(\frac{1}{\delta_2}\right)\beta\lambda(1-\lambda)} \\ &= 2(1-\beta)\left(\beta+\frac{1}{3}(1-\lambda)\right)\log\left(\frac{1}{\delta_2}\right)+2(1-\beta)\sqrt{2\log\left(\frac{1}{\delta_2}\right)\beta\lambda(1-\lambda)}.\end{aligned}$$

Therefore, for

$$\begin{aligned}N \quad \geq \quad & \frac{1}{(1-\lambda)^2(1-\beta)}\left[\lambda(1-\lambda)+2(1-\beta)\left(\beta+\frac{1}{3}(1-\lambda)\right)\log\left(\frac{1}{\delta_2}\right)\right. \\ & \left.+\sqrt{2\log\left(\frac{1}{\delta_2}\right)\beta\lambda(1-\lambda)}\right] \geq N_{max},\end{aligned}$$

we have $\Pr(\hat{q}_{N_\beta} \leq q_\beta) \geq 1 - \delta_2$. Note that $N > \frac{\lambda}{(1-\lambda)(1-\beta)}$, thus it is valid to use Bernstein's inequality. ∎

Now we can proceed with the proof of Theorem 2.4.1. In the following proof, we will be using Lemma 2.4.2 and Lemma 2.4.3.

Proof of Theorem 2.4.1 Define the events $S = [\hat{q}_{N_\beta} > q_\beta]$, $B = [F_D(\hat{q}_{N_\beta}) < (1 - \beta)\lambda - \alpha'(1-\beta)]$ and $L = [\bar{F}_D(\hat{q}_{N_\beta}) < 1 - (1-\beta)\lambda - \alpha'(1-\beta)]$.

Let $\alpha' = \frac{\alpha\omega}{3}$. Also let $\delta_1 = \frac{2\delta}{3}$ and $\delta_2 = \frac{\delta}{3}$. By Lemma 2.4.2, we know that for

$$\begin{aligned} N &\geq \frac{2}{(1-\beta)\alpha'}\left(\frac{1}{\alpha'}\lambda(1-(1-\beta)\lambda) + \frac{4}{3}\right)\log\left(\frac{2}{\delta_1}\right) \\ &= \frac{6}{(1-\beta)\alpha\omega}\left(\frac{3}{\alpha\omega}\lambda(1-(1-\beta)\lambda) + \frac{4}{3}\right)\log\left(\frac{3}{\delta}\right) \\ &= \frac{2}{(1-\beta)\alpha^2\omega}\left(\frac{9\lambda}{\omega}(1-(1-\beta)\lambda) + 4\alpha\right)\log\left(\frac{3}{\delta}\right) \end{aligned}$$

then $\Pr(B \cup L) \leq \delta_1 = \frac{2\delta}{3}$. Also, from Lemma 2.4.3, we know that for $N \geq N_1$, we have $\Pr(S) \leq \delta_2 = \frac{\delta}{3}$. Note that $[S^c \cap B^c \cap L^c]$ imply that $\hat{q}_{N_\beta}$ is α'-accurate with respect to $\hat{D}$. Also from Corollary 2.3.5, we know that if $\hat{q}_{N_\beta}$ is α'-accurate with respect to $\hat{D}$, then $\rho_{\hat{D}}(\hat{q}_{N_\beta}) \leq (1+\alpha)\rho(\hat{q}^*)$. Thus,

$$\begin{aligned} \Pr(\rho_{\hat{D}}(\hat{Q}_{N_\beta}) \leq (1+\alpha)\rho(\hat{q}^*)) &\geq \Pr(\hat{Q}_{N_\beta} \text{ is } \alpha'\text{-accurate}) \\ &\geq \Pr(S^c \cap B^c \cap L^c) \\ &= 1 - \Pr(S \cup B \cup L) \\ &\geq 1 - \Pr(S) - \Pr(B \cup L) \\ &\geq 1 - \delta_1 - \delta_2 = 1 - \delta. \quad ∎ \end{aligned}$$

In what follows, we will illustrate the effect of introducing risk-preferences (by the trimming factor β) to the sample size required to ensure accuracy of the regret. We compare the bound in Theorem 2.3.6 for the risk-neutral firm and the bound in Theorem 2.4.1 for the risk-averse firm. Figure 2-5 compares the two bounds for a trimming factor $\beta = 0.05$ and for critical fractiles $\lambda = 0.01, 0.99$. Note that even

though there is only a small amount of trimming ($\beta = 0.05$), if λ is very large, then the order of magnitude of the samples for the trimming case is much larger. If λ is small however, the two bounds are comparable.

Figure 2-6 plots the required sample size for the trimmed case as a function of the trimming factor β. Notice that as the trimming factor increases, more samples of the demand are required. This result follows from what we would intuitively expect since a larger trimming factor means that there are less samples from the random variable $\hat{D}$. This is compensated by taking more samples of the demand D. Also note that for $\beta = 0$, the number of samples for the trimmed and the non-trimmed cases are the same. And as the trimming factor increases, the two bounds diverge.

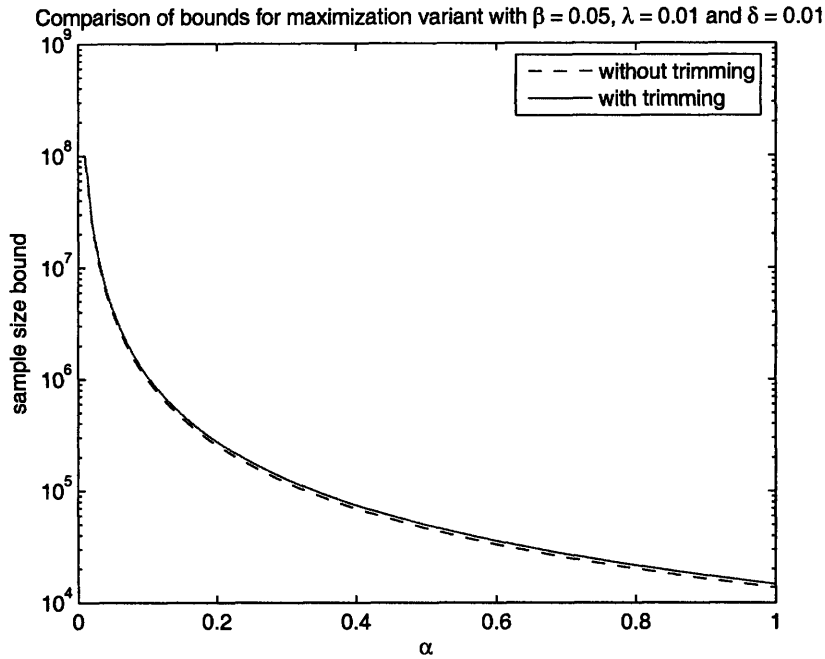

(a) $\lambda = 0.01$

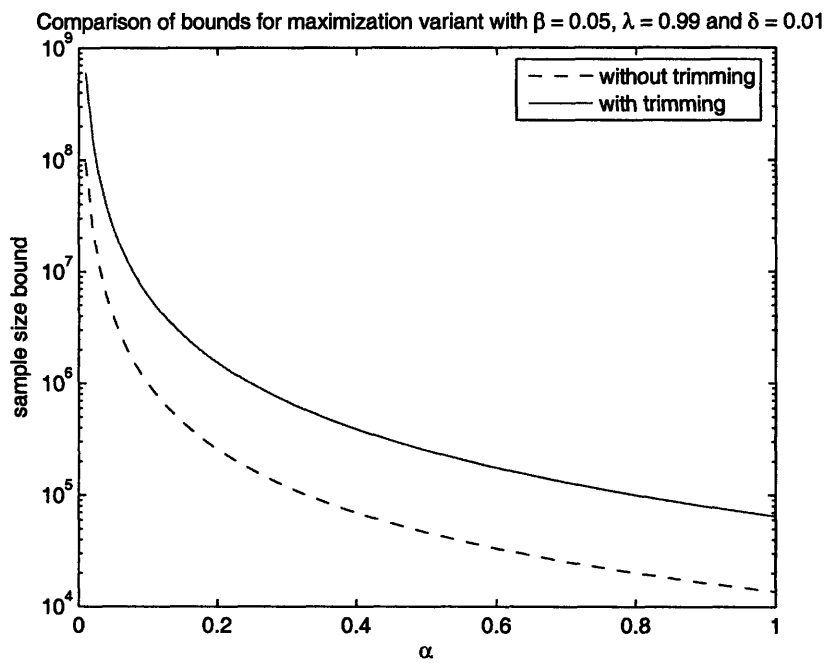

(b) $\lambda = 0.99$

Figure 2-5: Comparison of bounds for the case of trimming and without trimming (trimming factor β=0.05) as a function of the accuracy level α.

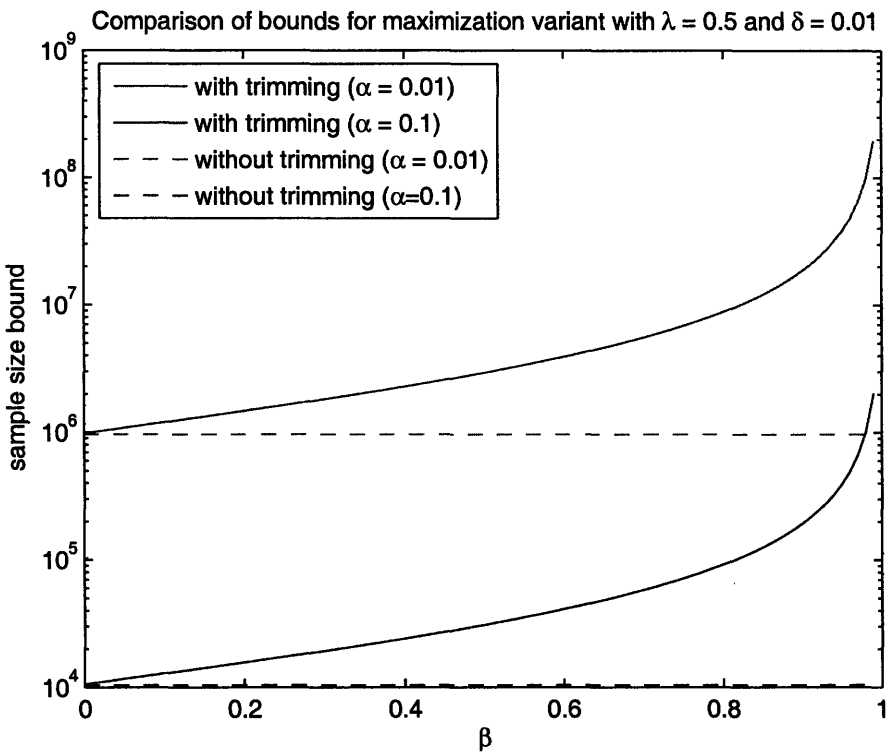

Figure 2-6: Plot of the bounds for for the required number of samples in the trimming case as a function of the trimming factor β.

Chapter 3

Numerical Results for the Data-Driven Newsvendor Problem with Exogenous Price

In the previous section, we introduced the Sample Average Approximation (SAA) method as a nonparametric approach for solving the newsvendor problem and its variants in scenarios where the explicit demand is not known. In this approach it is assumed that a set of N independent samples of the demand is available. In Chapter 2, we provide a theoretical analysis regarding the number of samples required to guarantee that the regret of the SAA solution has a bounded relative error compared to the optimal regret that is computed with respect to the true demand distribution. However, this bound is only a worst-case bound on the number of samples required. It is likely that in many cases a significantly smaller number of samples will suffice.

In this chapter, we conduct numerical experiments to evaluate the empirical performance of the SAA approach in different concrete scenarios. In particular:

- We explore how the SAA approach performs as the critical newsvendor fractile λ varies under various concrete demand distributions.
- We explore the sensitivity of the approach to different parameters of concrete distributions.

- We compare the performance of the SAA approach to other approaches for solving the newsvendor problem with partial demand information.

- Compare the theoretical and empirical confidence levels to achieve a specified accuracy of the regret of the SAA approach.

In particular, we consider the minimax regret approach and a parameter-estimation approach. The minimax regret approach minimizes the opportunity cost of not making the optimal decision over a family of distributions to which the true demand distribution is assumed to belong.

In addition, we compare the SAA approach to the classical parametric approach. In this approach, the demand is assumed to belong to a certain parametric family of distributions, and the data is used to estimate the corresponding parameters. After that, the analytical model is solved with respect to the estimated parameters.

The remainder of the chapter is structured as follows. Section 3.1 outlines experiment design as well as the performance measures of the numerical experiments. In Section 3.2, we present the results of the numerical experiments. We also make observations and present our insights from these results.

3.1 Methodology

We outline the steps of the SAA approach as we apply it in our numerical experiments.

- A set of N samples $d^1, \ldots, d^N$, is drawn from the true demand distribution.

- We denote $\hat{q}_N$ as the $\lambda = 1 - \frac{c}{p}$ sample quantile.

- The first two steps are repeated independently for a total of m times.

The optimal solution to the SAA counterpart with N samples, $\hat{Q}_N$, is a random variable that is dependent on the specific N samples of the demand D. The procedure outlined above is repeated for a total of m times to approximate a distribution for $\hat{Q}_N$ and the absolute error of the SAA quantity $|\hat{Q}_N - q^*|$, the expected profit $g(\hat{Q}_N)$

and the regret $\rho(\hat{Q}_N)$, where q^* is the optimal order quantity under full knowledge of the demand.

In the experiments, we take $N = 100$ samples for $m = 100$ runs of the procedure. We generate data from five distributions: uniform, normal, exponential, Pareto and Poisson. In the base case, the unit price is set to $p = 3$ and the cost is $c = 1$ The specifications of the distribution in the base case are as follows:

Uniform	range of $[0, 2\mu]$ with $\mu = 50$
Normal	mean $\mu = 50$ and standard deviation $\sigma = 25$
Exponential	mean $\mu = 50$
Pareto	mean $\mu = 50$ and range $[1, \infty)$
Poisson	mean $\mu = 50$

In analyzing the sensitivity to the critical fractile, the price is $p = 3$ and the unit cost takes on values $c \in 3 \times [0.05, 0.10, \ldots, 0.95]$.

The sensitivity analysis of the SAA procedure to various distribution parameters is conducted with the following specifications

Uniform	$\mu \in [12.5, 25, \ldots, 250]$
Normal	$\mu \in [12.5, 25, \ldots, 250]$
	$\sigma \in [2.5, 5, \ldots, 50]$
Exponential	$\mu \in [10, 20, \ldots, 200]$
Pareto	$k \in [1, 2, \ldots, 50]$

In the minimax regret approach, a family of distributions with mean $\mu = 50$ is assumed. The distributions used to compare with the minimax regret order quantity are: uniform, normal, exponential and Poisson, with parameters specified in the base case above.

In the parameter-estimation approach, data is generated from the following distributions: uniform, Pareto, exponential and Poisson, where the parameters are specified by the base case above. The data is then fitted to different distributions to estimate the parameters.

In comparing the theoretical and empirical confidence levels achieved by the regret of the SAA approach for specified accuracies, we consider accuracy rates $\alpha \in [0.05, 0.10, \ldots, 0.95]$ and critical fractiles $\lambda = 0.1, 0.5, 0.9$. We take $N = 100$ and repeat the SAA procedure for $m = 100$ runs. The theoretical confidence levels are computed from the Bernstein and Hoeffding bounds (Theorem 2.3.1 and Theorem 2.3.9). If the theoretical confidence level computed is negative, we simply set it to zero. Suppose $\hat{q}_N^k$ is the SAA ordering quantity for $k = 1, \ldots, m$. The empirical confidence level is computed by the following formula

$$\text{Empirical Confidence Level} = \frac{1}{m}\sum_{k=1}^{m} \mathbf{1}(\rho(\hat{q}_N^k) \leq (1+\alpha)\rho(q^*))$$

where $\mathbf{1}(\cdot)$ is the respective indicator function and q^* is the optimal ordering quantity under full knowledge of the distribution.

The performance of the SAA approach is evaluated by observing the relative error of the average expected profit and average regret. In particular, if $\hat{q}_N^k$ is the SAA ordering quantity for the kth run for $k = 1, \ldots, m$, then

$$\begin{aligned}
\text{Average Absolute Error of the Order Quantity} &= \frac{1}{m}\sum_{k=1}^{m} |\hat{q}_N^k - q^*|, \\
\text{Average Relative Error of the Expected Profit} &= \frac{g(q^*) - \frac{1}{m}\sum_{k=1}^{m} g(\hat{q}_N^k)}{g(q^*)}, \\
\text{Average Relative Error of the Regret} &= \frac{\frac{1}{m}\sum_{k=1}^{m} \rho(\hat{q}_N^k) - \rho(q^*)}{\rho(q^*)},
\end{aligned}$$

where q^* is the optimal order quantity under full knowledge of the distribution.

3.2 Results and Discussion

We present the results of our numerical experiments in this section as specified by the experimental design in Section 3.1.

α	Theoretical Confidence Level		Empirical Confidence Level			
	Hoeffding	Bernstein	Uniform	Normal	Exponential	Pareto
0.05	0	0	0.971	0.99	0.998	1
0.1	0	0	0.999	0.999	1	1
0.15	0	0	1	1	1	1
0.2	0	0	1	1	1	1
0.25	0	0	1	1	1	1
0.3	0	0	1	1	1	1
0.35	0	0	1	1	1	1
0.4	0.17778	0.12779	1	1	1	1
0.45	0.3507	0.29597	1	1	1	1
0.5	0.5013	0.44506	1	1	1	1
0.55	0.62746	0.57285	1	1	1	1
0.6	0.72933	0.67894	1	1	1	1
0.65	0.80873	0.76438	1	1	1	1
0.7	0.86854	0.83121	1	1	1	1
0.75	0.91213	0.882	1	1	1	1
0.8	0.94287	0.91952	1	1	1	1
0.85	0.96387	0.94647	1	1	1	1
0.9	0.97778	0.9653	1	1	1	1
0.95	0.98671	0.97808	1	1	1	1

Table 3.1: Theoretical and empirical confidence levels for various accuracy levels of the solution to the SAA counterpart of the profit-maximizing newsvendor with critical fractile $\lambda = 0.5$

Influence of the Critical Fractile

Figure A-1 plots the average of the error $|\hat{q}_N - q^*|$ as a function of the critical fractile. Figure A-2 plots the relative error of the expected profit and the regret achieved by the SAA counterpart as a function of the critical fractile.

Tables 3.1–3.2 summarize the theoretical and empirical confidence levels achieved by various distributions for $\lambda = 0.1, 0.5, 0.9$.

We make the following observations:

- In most of the demand distributions considered, the error $|\hat{q}_N - q^*|$ generally increases as the critical fractile increases. This is especially true of demand distributions that are unbounded from above (see Figure A-1).

- The relative error of the expected profit of the SAA solution is highly sensitive to the critical fractile (see Figure A-2(a)). In particular, the error is larger for

α	Theoretical Confidence Level		Empirical Confidence Level			
	Hoeffding	Bernstein	Uniform	Normal	Exponential	Pareto
0.05	0	0	0.946	0.924	0.906	1
0.1	0	0	0.995	0.991	0.976	1
0.15	0	0	0.997	1	0.995	1
0.2	0	0	0.999	1	0.999	1
0.25	0	0	1	1	1	1
0.3	0	0	1	1	1	1
0.35	0	0	1	1	1	1
0.4	0	0	1	1	1	1
0.45	0	0	1	1	1	1
0.5	0	0	1	1	1	1
0.55	0	0	1	1	1	1
0.6	0	0	1	1	1	1
0.65	0	0	1	1	1	1
0.7	0	0	1	1	1	1
0.75	0	0	1	1	1	1
0.8	0	0	1	1	1	1
0.85	0	0	1	1	1	1
0.9	0	0	1	1	1	1
0.95	0	0	1	1	1	1

Table 3.2: Theoretical and empirical confidence levels for various accuracy levels of the solution to the SAA counterpart of the profit-maximizing newsvendor with critical fractile $\lambda = 0.9$

smaller values of λ, which is partly influenced by the small optimal profit at small critical fractiles. In contrast, the relative error of the regret of the SAA solution is generally invariant of changes in the critical fractile (see Figure A-2(b)).

- The Pareto distribution achieves the lowest relative error in the regret (see Figure A-2(b)) for all critical fractiles, despite having infinite variance ($k \leq 2$).

- In the range of critical fractiles, the largest relative error of the expected profit is of the order 10^0. The largest relative error of the regret is of the order 10^{-1}.

- From Tables 3.1–3.2, we observe that the theoretical confidence level are much smaller than the empirical confidence levels for $N = 100$. In particular, the empirical confidence levels are very high even with a small accuracy level. This verifies that the bounds we derive in Chapter 2 are worst-case bounds (i.e., the sample size required for the regret to achieve a certain accuracy and confidence level is much smaller than what we derive).

Influence of the Distribution Parameters

Figure A-3 plots the average relative error of the expected profit and the regret achieved by the SAA counterpart as a function of the mean of a uniform distribution. The errors are compared for critical fractiles $\lambda = 0.1, 0.5, 0.9$.

Figure A-4 plots the average relative error of the expected profit and the regret achieved by the SAA counterpart as a function of the mean of a normal distribution. The errors are compared for critical fractiles $\lambda = 0.1, 0.5, 0.9$.

Figure A-5 plots the average relative error of the expected profit and the regret achieved by the SAA counterpart as a function of the standard deviation of a normal distribution. The errors are compared for critical fractiles $\lambda = 0.1, 0.5, 0.9$.

Figure A-6 plots the average relative error of the expected profit and the regret achieved by the SAA counterpart as a function of the mean of an exponential distribution. The errors are compared for critical fractiles $\lambda = 0.1, 0.5, 0.9$.

Figure A-7 plots the average relative error of the expected profit and the regret achieved by the SAA counterpart as a function of the parameter k of a Pareto distribution. The errors are compared for critical fractiles $\lambda = 0.1, 0.5, 0.9$.

We make the following observations:

- The relative error of the regret of the SAA solution is invariant of the demand distribution parameters (see Figures A-3(b), A-4(b), A-5(b), A-7(b)). In contrast, the relative error of the expected profit of the SAA solution is highly sensitive to changes in the distribution parameters (see Figures A-3(a), A-4(a), A-5(a), A-7(a)).

Comparison to other approaches

Figures A-8–A-9 compare the relative error of the regret of the minimax regret order quantity and SAA order quantity to the true optimal order quantity as a function of the critical fractile, λ.

Figures A-10–A-11 compare the relative error of the regret of the parameter-estimation order quantity and SAA order quantity to the true optimal order quantity as a function of the critical fractile, λ.

We make the following observations:

- Generally, the SAA method achieves a lower relative error in the regret than the minimax regret approach (see Figures A-9 and A-8). The only exception is for the exponential distribution (A-9(a)) which achieves an accuracy comparable and sometimes even better than that of the SAA method. This is because the exponential distribution is expected to be the most robust, given that it is entropy-maximizing over the class of nonnegative distributions with known mean (Perakis and Roels [28]).

- If the assumed distribution differs greatly from the true demand distribution, then the parameter estimation approach achieves very high relative errors (e.g., data from a uniform distribution fitted to a Pareto distribution in Figure A-10(a)). On the other hand, if the assumed distribution has similar properties

to the true demand distribution, then the parameter estimation approach can achieve low relative errors comparable to that of the SAA approach (e.g., data from a uniform distribution fitted to a normal distribution in Figure A-10(a) where the normal distribution preserves symmetry of the uniform distribution).

- From the numerical experiments, the regret of the solution to a parameter-estimation approach can grow up to 10^{15} times the optimal regret of the true distribution (see Figure A-10(a)).

- From the numerical experiments, the regret of the solution to the minimax regret approach can grow up to 100 times the optimal regret of the true distribution (see Figure A-8(a)).

From the numerical experiments, we can conclude that under the regret criterion (i.e., small relative error of the regret) the SAA approach provides a consistently good estimate of the optimal order quantity that is invariant of external demand factors (e.g. mean or variance). Moreover, it generally performs well against other approaches, namely the minimax regret approach and parameter-estimation approach.

Chapter 4

Data-Driven Approach to the Price-Setting Newsvendor Problem

In this chapter, we consider the price-setting newsvendor problem under the assumption that the explicit demand distribution is not known. We assume that the demand is composed of a deterministic component and a random component. The deterministic demand is a linear function of the price. In contrast to the nonparametric data-driven approach described in in Chapter 2, we assume that the parameters of the deterministic demand are known. On the other hand, we assume that the only information known about the random component is a set of independent samples drawn from the true distribution. We describe and analyze a simulation-based approach proposed by Zhan and Shen [42] for solving the price-setting newsvendor problem under partial information on the demand. Specifically, we provide a theoretical analysis of the number of samples required to achieve a good quality solution with high probability.

The price-setting newsvendor problem is an extension of the classical newsvendor problem that incorporates price as a decision variable. Consider a firm selling a product over a single sales period. The firm must produce units of the commodity at the beginning of the period prior to observing the demand for the product. The firm must also simultaneously set the unit selling price. The firm faces a stochastic demand for that product, which depends on the selling price. The actual demand

occurs during the sales period and is satisfied as much as possible with the units produced. The firm incurs a cost proportional to the production quantity. Any unmet demand is assumed to be lost. The objective of the firm is to then maximize its expected profit.

Studying the price-setting newsvendor problem can provide excellent insights on how operational problems (i.e., production decisions) interact with marketing issues (i.e., pricing decisions) to influence decision-making for a firm. Moreover, even though the integration of pricing and production decisions is in early stages in most manufacturing companies, it still has the potential to radically improve supply chain efficiencies [11]. Thus, the coordination between pricing and production decisions is a relevant research direction in revenue management with many practical applications.

Whitin [40] was the first to incorporate pricing decisions in the newsvendor problem. He showed how to optimize the price and order quantity decisions under the assumption that the deterministic demand is a linear function of the price. Mills [26] and Karlin and Carr [21] study the effect of the particular form of demand uncertainty to the pricing decision. Mills [26] showed that for an additive demand model, the optimal price under uncertain demand is at most equal to the optimal price under the assumption of deterministic demand (or the optimal riskless price). The optimal riskless price is the price that maximizes the expected profit under the assumption that the demand is known beforehand. Conversely, Karlin and Carr [21] find that the optimal price is at least the riskless optimal price for the multiplicative model of uncertain demand. A unified framework to reconcile this apparent contradiction has been provided by Petruzzi and Dada [30].

There have been numerous extensions to the price-setting newsvendor problem in the recent years. For instance, Dana and Petruzzi [12] consider a model where the demand depends on both the price and inventory levels. Raz and Porteus [32] remove the distributional assumption about the random demand component (e.g. additive or multiplicative) by assuming demand is a discrete random variable that depends on price. Bernstein and Federgruen [3] study a model with many newsvendors. For an extensive overview of recent developments in the area of price and production, we

refer the reader to [11].

A common assumption in most inventory models with stochastic demand is that the underlying demand distribution is known exactly. Although this results in tractable models, it does not hold in many realistic settings where the demand distribution is often unknown. Instead of full knowledge of the demand, usually there is only imperfect information (e.g. moments of the distributions or historical price-demand data). Researchers have looked at different approaches to solve the price-setting newsvendor problem under the assumption of imperfect information.

A popular approach to address imperfect information for an inventory model with simultaneous pricing and inventory decisions is through a "demand learning" approach. Monahan, Petruzzi and Zhao [27] apply a dynamic pricing to the newsvendor model and establish a practical and efficient algorithm for computing optimal prices.

Another approach proposed by Raz and Porteus [33] is by approximating the demand distribution with a finite number of fractiles and assuming that the fractile functions are piecewise linear functions of the price. They propose an algorithm that finds the optimal price-quantity pair by a method of elimination. The advantage of this approach is that it applies to more general cases of the demand than the case of additive of multiplicative uncertainty. On the other hand, the method requires enough demand samples for each price level to accurately approximate the demand fractiles.

An alternative method proposed by Zhan and Shen [42] uses the Sample Average Approximation (SAA) scheme to approximate a system of two equations that are derived from the first-order conditions of the problem. They propose a gradient search algorithm to solve the data-driven approximation of the first-order conditions. Although they assume a linear price-demand relationship with known parameters, but with no additional assumptions on the distribution of the stochastic term. However, Zhan and Shen [42] do not provide a theoretical framework for the algorithm. In particular, it is of interest to know how the accuracy of the simulation-based approach relates to the number of samples of the data.

In this chapter, we extend the results of Zhan and Shen [42] by providing a theo-

retical framework for the simulation-based approach as a method to solve the price-setting newsvendor problem. Our goal is to establish a relationship between the sample size and the accuracy of the simulation-based approach as an approximation to the price-setting newsvendor problem. We derive a bound on the sample size that ensures with high probability that the regret of the solution to the simulation-based approach has a small relative error with respect to the regret of the optimal policy under full knowledge of the demand distribution. This bound is applicable for demand distributions with a random component having bounded support.

The remainder of the chapter is structured as follows. Section 4.1 introduces the price-setting newsvendor problem where the characterization of the uncertainty is fully known. We will show that under certain assumptions on the model, the optimal policy uniquely solves certain first order conditions. Section 4.2 introduces the simulation-based approach proposed by Zhan and Shen [42]. Finally, Section 4.3 establishes the connection between the sample size and the accuracy of the simulation-based approach as an approximation to the price-setting newsvendor problem.

4.1 The Price-Setting Newsvendor Model

In this section, we consider the price-setting newsvendor model. The problem deals with determining a simultaneous pricing and ordering decision for a firm to maximize profit under uncertain demand. The demand distribution is assumed to be known. Petruzzi and Dada [30] provide a review and extensions of the price-setting newsvendor problem.

4.1.1 Problem Formulation

First let us define the following notation, which we will be using throughout the chapter:

c	unit cost
p	unit selling price
q	order quantity
$D(p, \epsilon)$	random demand
$y(p)$	deterministic demand
ϵ	random component of the demand
$F(\cdot), f(\cdot)$	cumulative and probability density distribution of ϵ

A random demand $D(p, \epsilon)$, which is a function of the price, occurs for a single commodity in a single period. At the beginning of the period, the firm decides how many units of the commodity to order to satisfy the random demand. The q units are ordered at a cost of c. The firm simultaneously decides on a unit selling price p for the product. During the period, the actual demand $D(p, e)$ (the realization of $D(p, \epsilon)$, whose distribution is affected by p) is observed. The firm sells the minimum between the demand and the number of units ordered at the selling price p. The profit of the firm is given by

$$\pi(q, p) = p \ \min(D(p, e), q) - c \ q.$$

Since the actual demand is not known when pricing and ordering decisions are made, a sensible objective for the firm is to maximize the expected profit. Thus, the problem is equivalent to

$$\max_{p,\ q} g(p, q) = E[\pi(q, p)] = p \ E[\min(D(p, \epsilon), q)] - cq. \tag{4.1}$$

A feasible pricing and ordering policy is one for which the price and the order quantity are nonnegative. However, we do not need to impose constraints on the problem because, as we will see in Proposition 4.1.2, the optimal solution is always feasible.

We will assume that the deterministic demand is a linear function of the price.

We have

$$y(p) = a - bp \tag{4.2}$$

where a and b are known positive constants. Consistent with Mills [26], we assume an additive form of the demand uncertainty. The random demand function is given by

$$D(p, \epsilon) = y(p) + \epsilon \tag{4.3}$$

where ϵ is the random component of the demand that can assume values in $[A, B]$. We assume that ϵ is a continuous random variable consistent with the models of Petruzzi and Dada [30]. A standard assumption of the additive model is that $A + a > 0$. This guarantees that the demand is positive for some valid range of the selling price. From a practical standpoint, however, it is enough to assume that $a+\mu$ is significantly greater than σ. This allows us to include unbounded distributions, such as the normal distribution in modeling the demand.

4.1.2 Optimality Conditions

We follow the analysis of Zhan and Shen [42] to derive optimality conditions for the price-setting newsvendor problem.

We can find a convenient expression of the problem consistent with Ernst [14] and Thowsen [39] by defining $z = q - y(p)$, which corresponds to the quantity ordered above the deterministic demand. If the choice of z is above the realized value of ϵ, then leftovers occur; if the choice of z is smaller than the realized value of ϵ, then shortages occur. The problem is then reduced to

$$\max_{p,\ z} \ g(z, p) \tag{4.4}$$

where

$$g(z, p) = E[\pi(z, p)] = (p - c)y(p) + p\ E[\min(\epsilon, z)] - cz.$$

Therefore, the price-setting newsvendor problem is equivalent to solving the problem defined by (4.4). For simplicity in our analysis, we will be making the following

assumptions:

Assumption 4.1.1 *The random demand component ϵ has bounded support.*

Assumption 4.1.2 *The optimal riskless price $p^0 = \frac{a+bc+\mu}{2b}$ is strictly greater than $\frac{B-A}{2b} + c$.*

Assumption 4.1.3 *The hazard rate $r(z) = \frac{f(z)}{1-F(z)}$ satisfies $r^2(z) + \frac{dr(z)}{dz} > 0$ for $z \in [A, B]$.*

Note that the optimal riskless price

$$p^0 = \frac{a + bc + \mu}{2b}$$

in Assumption 4.1.2 represents the optimal price under the assumption of deterministic demand. If the firm knows beforehand that the demand as a function of the price is $y(p) + e$ (where e is a realization of the random component ϵ of the demand), then it will order exactly $y(p) + e$ units. To maximize the expected profit, it will set a price p^0 that maximizes the expected profit under the riskless scenario $(p - c)(y(p) + \mu)$.

Note that Assumption 4.1.3 is satisfied by all nondecreasing hazard rate distributions, which include uniform, normal, logistic, chi-squared and exponential distributions (see [42]).

In the literature (see for example Whitin [40]; Zabel [41]), researchers usually reduce the original problem to an optimization problem over a single variable z. Zhan and Shen [42] on the other hand describe an approach by working directly on the first order conditions of the problem. We will be following the approach of Zhan and Shen [42] to find the optimality conditions to the problem.

The first-order conditions of the price-setting newsvendor problem are

$$\frac{\partial g}{\partial z} = -c + p(1 - F(z)) = 0, \tag{4.5}$$

$$\frac{\partial g}{\partial p} = a + bc - 2bp + E[\min(z, \epsilon)] = 0, \tag{4.6}$$

where $F(\cdot)$ is the cumulative distribution function of ϵ.

Note that, if (4.5) and (4.6) form a system of equations, then this implies the following relationship between z and p:

$$z(p) = F^{-1}\left(1 - \frac{c}{p}\right), \tag{4.7}$$

$$p(z) = p^0 - \frac{1}{2b}\int_z^B (u - z)f(u)du, \tag{4.8}$$

where p^0 is the riskless optimal price.

Note that $p(z)$ is a concave function of z. Under Assumption 4.1.3, $z(p)$ is a concave function of p. Zhan and Shen [42] show that under this assumption, both curves intersect in at most two points. Note that any solution to this system satisfies the first order conditions. They have the following proposition:

Proposition. 4.1.1 [Zhan and Shen [42]]. *Suppose $F(\cdot)$ is a distribution function satisfying Assumption 4.1.3.*

(1) If $a - bc + A > 0$, then equations (4.7) and (4.8) have a unique solution, which is also optimal for the problem.

(2) If $a - bc + A < 0$ and $f(A) < \frac{1}{a+bc+A}$, then equations (4.7) and (4.8) have two solutions. The one with the larger p value is the optimal solution for the problem.

(3) If $a - bc + A < 0$ and $f(A) > \frac{1}{a+bc+A}$, then equations (4.7) and (4.8) have no solution.

Note that Assumption 4.1.2 implies that $a - bc + A > B - \mu$. Therefore, Assumptions 4.1.2–4.1.3 imply that Case (1) of Proposition 4.1.1 holds. In other words, the first order conditions have a unique solution which is also optimal to the problem. Thus, if we define (z^*, p^*) as the solution to the price-setting newsvendor problem, then the following relationship holds under Assumption 4.1.2:

$$z^* = F^{-1}\left(1 - \frac{c}{p^*}\right), \tag{4.9}$$

$$p^* = \frac{1}{2b}(a + bc + E[\min(z^*, \epsilon)]). \tag{4.10}$$

For simplicity of notation, we define the following quantities

$$p_{min} = p^0 - \frac{B-A}{2b}, \tag{4.11}$$

$$p_{max} = \frac{a+bc+B}{2b}. \tag{4.12}$$

The following Proposition establishes bounds on the solution to the first-order conditions.

Proposition. 4.1.2 *Under Assumption 4.1.1, if (z^*, p^*) satisfies the first-order conditions (4.5) and (4.6), then (1) $A \leq z^* \leq B$, (2) $c \leq p^* \leq p^0$, (3) $p^* \geq p_{min}$, and (4) $q^* = z^* + y(p^*) \geq 0$, where $p_{min} = p^0 - \frac{1}{2b}(B-A)$.*

Proof. Note that (1) clearly follows from condition $z^* = F^{-1}\left(\frac{c}{p^*}\right)$.

To prove (2), note that the lower bound follows since the first order conditions imply $F(z^*) = 1 - \frac{c}{p^*}$, which only makes sense when $p^* \geq c$. Also, equation (4.10) implies that $p^* = p^0 - \frac{1}{2b}\int_z^B (u-z)f(u)du$. Since $\int_z^B (u-z)f(u)du$ is always nonnegative, then this implies that $p^* \leq p^0$.

To prove (3), note that the first order conditions imply $p^* = p^0 - \frac{1}{2b}\int_{z^*}^B (u - z^*)f(u)du$. Note that $\int_{z^*}^B (u-z^*)f(u)du \geq (B-z^*)\Pr(\epsilon \geq z^*) \geq (B-z^*)$. Also from (1), we know that $z^* \geq A$. Thus, $p^* \geq p^0 - \frac{1}{2b}(B-A)$.

To prove (4), note that, from equation (4.10), we know that $z^* \geq E[\min(z^*, \epsilon)] = 2bp^* - a - bc$. Therefore, $q^* = z^* + a - bp^* \geq b(p^* - c) \geq 0$, where the last inequality follows from (2). ∎

From (4) of Proposition 4.1.2, note that the solution to the first order conditions (z^*, p^*) always provide a feasible (i.e., nonnegative) ordering quantity q^*.

4.2 Simulation-Based Procedure

One common approach to approximate the expected value of a random function is the Sample Average Approximation (SAA) method. In this method, the function is

approximated by taking the average value over finitely many independent samples that are drawn from a probability distribution.

In the previous section, we have shown that under Assumptions 4.1.2–4.1.3, the optimal solution uniquely solves the first-order conditions (4.5)–(4.6). If an explicit cumulative density function of the demand is known, then a gradient search algorithm can be used to find the optimal pricing and ordering policy. However, in most realistic settings, the firm may not know the exact distribution that the demand follows. Usually, the only information available to the firm is historical price-demand data from past periods. If we know that the deterministic demand is a linear function of the price, these data imply realizations of the random component ϵ. Therefore, we can approximate the first-order conditions by the SAA method with a function that depends on realizations of the random component ϵ. This simulation-based approach for solving the optimal pricing and ordering policy of the price-setting newsvendor problem has been proposed by Zhan and Shen in [42].

4.2.1 Problem Formulation

Suppose that the only information available to a firm is a set of N price-demand pairs $\{(d^k, p^k)\}_{k=1}^N$. We assume that the deterministic component of the demand is a linear function of price with known parameters. Suppose the demand uncertainty is additive. Therefore, we can find the corresponding realizations of the random component using the transformation $\epsilon^k = d^k - (a - bp^k)$ for $k = 1, \ldots, N$. This data allows us to obtain the data-driven approximation of the first-order conditions:

$$-c + p\left(1 - \hat{F}_N(z)\right) = 0, \tag{4.13}$$

$$a + bc - 2bp + \hat{E}_N[\min(z, \epsilon)] = 0, \tag{4.14}$$

where

$$\hat{F}_N(z) = \frac{1}{N}\sum_{k=1}^{N} \mathbf{1}(\epsilon^k \leq z), \tag{4.15}$$

$$\hat{\bar{F}}_N(z) = \frac{1}{N}\sum_{k=1}^{N} \mathbf{1}(\epsilon^k \geq z), \tag{4.16}$$

$$\hat{E}_N[\min(z,\epsilon)] = \frac{1}{N}\sum_{k=1}^{N} \min(z,\epsilon^k). \tag{4.17}$$

Let $(\hat{Z}_N, \hat{P}_N)$ denote the solution to the data-driven approximation of the first-order conditions with N samples. Note that $(\hat{Z}_N, \hat{P}_N)$ are random variables that depend on the specific N samples of ϵ. Let $(\hat{z}_N, \hat{p}_N)$ be the realization of $(\hat{Z}_N, \hat{P}_N)$ for a specific N samples of ϵ.

Proposition. 4.2.1 *If $(\hat{z}_N, \hat{p}_N)$ is the solution to the data-driven approximation to the first-order conditions, then (1) $A \leq \hat{z}_N \leq B$, (2) $c \leq \hat{p}_N \leq \hat{p}_N^0 \leq p_{max}$, and (3) $\hat{p}_N \geq \hat{p}_N^0 - \frac{1}{2b}(B-A)$, where $\hat{p}_N^0 = \frac{a+bc+\hat{\mu}_N}{2b}$ and $\hat{\mu}_N = \frac{1}{N}\sum_{k=1}^{N}\epsilon^k$*

Proof. The proof of (1), (2) and (3) is similar to the proof in Proposition 4.1.2 since $\hat{F}_N(\cdot)$ is the cumulative distribution function of the empirical distribution defined by the N samples.

4.3 Relative Error of the Regret of the Solution to the Simulation-Based Procedure

In Section 4.2, we have introduced the simulation-based method proposed by Zhan and Shen [42] for solving the price-setting newsvendor problem. A natural question to ask is how many samples are required in order to ensure with high probability that the simulation-based procedure provides a "good" solution to the price-setting newsvendor problem.

It is natural to think of an approximation to the price-setting newsvendor problem being "good" if the relative error of its expected profit is small. A potential com-

plication with this criterion is that the expected profit is not guaranteed to be well above zero. To circumvent this problem, we instead consider a minimization problem having the same solution as the price-setting newsvendor problem. If we can show that the approximation evaluated at this new objective function has a small relative error, then we can say that it is a "good" approximation of the optimal solution of the price-setting newsvendor problem.

As mentioned previously, under a scenario where the firm knows the demand beforehand, $p^0 = \frac{a+bc+\mu}{2b}$ is the price it will set to maximize the expected profit under this scenario. In particular, the expected profit the firm will receive is $(p^0 - c)(a - bp^0 + \mu)$. Under demand uncertainty, the firm cannot achieve a profit greater than this. Since we also know that $c \leq p^0 \leq p_{max} = \frac{a+bc+B}{2b}$, then a valid upper bound for the expected profit under demand uncertainty is $(p_{max} - c)(a - bc + \mu)$. Therefore, we can define the regret function

$$\rho(z,p) = (p_{max} - c)(a - bc + \mu) - g(z,p), \tag{4.18}$$

where $p_{max} = \frac{a+bc+B}{2b}$.

The price-setting newsvendor problem is equivalent to the problem of minimizing $\rho(z,p)$. Therefore, the optimal policy (z^*, p^*) also minimizes the regret function $\rho(z,p)$. In this chapter, we will say that an approximation to the price-setting newsvendor optimal policy is accurate if the relative error of its regret is small. Consider any specified accuracy level $\alpha > 0$ and confidence level $1 - \delta$. We will show that there exists a number of samples N that depend on the the accuracy and confidence level as well as the parameters of the model, such that with probability of at least $1 - \delta$, the solution to the simulation-based approach defined on N samples has a regret $\rho(\hat{Z}_N, \hat{P}_N)$ of at most $(1 + \alpha)\rho(z^*, p^*)$.

Since $(\hat{z}_N, \hat{p}_N)$ is the solution to the data-driven approximation of the first-order conditions, it is useful to define a property that $(\hat{z}_N, \hat{p}_N)$ approximately solves the first-order conditions. We state the following definition.

Definition 4.3.1 *Let $(\hat{z}_N, \hat{p}_N)$ be some realization of $(\hat{Z}_N, \hat{P}_N)$. We say that $(\hat{z}_N, \hat{p}_N)$*

is (α_1', α_2')-accurate *if* $-\alpha_1' \leq 1 - F(\hat{z}_N) - \frac{c}{\hat{p}_N} \leq \alpha_1'$ *and* $-\alpha_2' \leq a + bc + E[\min(\hat{z}_N, \epsilon)] - 2b\hat{p}_N \leq \alpha_2'$, *for some* $\alpha_1', \alpha_2' > 0$

The implication of $(\hat{z}_N, \hat{p}_N)$ satisfying the conditions for (α_1', α_2')-accuracy is that for sufficiently small α_1' and α_2', the approximation $(\hat{z}_N, \hat{p}_N)$ "almost" satisfies the first-order conditions (4.5) and (4.6).

4.3.1 Accuracy of the Regret of a First-Order Approximate Policy

We would like to show that (α_1', α_2')-accuracy of the data-driven solution relates in some way to the optimal policy of under full knowledge of the distribution. In particular, we would like to show that the regret of the data-driven policy $\rho(\hat{z}_N, \hat{p}_N)$ has a bounded relative error with respect to $\rho(z^*, p^*)$. We will show this in a two-step analysis. In what follows, we show that if $(\hat{z}_N, \hat{p}_N)$ is (α_1', α_2')-accurate and $|\hat{\mu} - \mu| \leq 2b\gamma$ for some $\gamma > 0$, then the data-driven policy under a modified one-dimensional problem (with either price of quantity fixed) has bounded relative error. Specifically:

(1) There exists $\beta_1 > 0$ such that $(1-\beta_1)\rho(z^*, \hat{p}_N) \leq \rho(z^*, p^*)$ [refer to Lemma 4.3.4];

(2) There exists $\beta_2 > 0$ such that $(1-\beta_2)\rho(\hat{z}_N, \hat{p}_N) \leq \rho(z^*, \hat{p}_N)$ [refer to Lemma 4.3.7];

The combination of these one-dimensional problems implies that the regret of the two-dimensional problem has a bounded relative error. In what follows, we prove the relationship between (α_1', α_2')-accuracy and the relative error of the regret. We do this by the two-step analysis we outline above.

First, we establish the following result:

Lemma 4.3.1 *Suppose Assumption 4.1.1 holds. Let* $(\hat{z}_N, \hat{p}_N)$ *be a* (α_1', α_2')*-accurate realization of* $(\hat{Z}_N, \hat{P}_N)$*, and* (z^*, p^*) *be the solution to the first-order conditions (4.5) and (4.6). Then* $|E[\min(z^*, \epsilon)] - E[\min(\hat{z}_N, \epsilon)]| \leq (1 + \alpha_1')(B - A)$.

Proof. Consider the case where $z^* \leq \hat{z}_N$. If $\epsilon \leq z^*$, then $\min(z^*, \epsilon) - \min(\hat{z}_N, \epsilon) = 0$. On the other hand, if $\epsilon > z^*$, then $\min(z^*, \epsilon) - \min(\hat{z}_N, \epsilon) \geq z^* - \hat{z}_N$. Thus,

$E[\min(z^*,\epsilon)] - E[\min(\hat{z}_N,\epsilon)] \geq (z^* - \hat{z}_N)(1 - F(z^*)) = -\frac{c}{p^*}|z^* - \hat{z}_N|$. Since we know that $|z^* - \hat{z}_N| \leq B - A$ and $p^* \geq c$, then it follows that $E[\min(z^*,\epsilon)] - E[\min(\hat{z}_N,\epsilon)] \geq -(B - A) \geq -(1 + \alpha_1')(B - A)$.

Now consider the case where $z^* > \hat{z}_N$. If $\epsilon \leq \hat{z}_N$, then $\min(z^*,\epsilon) - \min(\hat{z}_N,\epsilon) = 0$. On the other hand, if $\epsilon > \hat{z}_N$, then $\min(z^*,\epsilon) - \min(\hat{z}_N,\epsilon) \leq z^* - \hat{z}_N$. Thus, $E[\min(z^*,\epsilon)] - E[\min(\hat{z}_N,\epsilon)] \leq (z^* - \hat{z}_N)(1 - F(\hat{z}_N))$. Since we know that $|z^* - \hat{z}_N| \leq B - A$ and $(\hat{z}_N, \hat{p}_N)$ is (α_1', α_2')-accurate, then it follows that $E[\min(z^*,\epsilon)] - E[\min(\hat{z}_N,\epsilon)] \leq \left(\frac{c}{\hat{p}_N} + \alpha_1'\right)(B - A)$. Also, since $\hat{p}_N \geq c$, then $E[\min(z^*,\epsilon)] - E[\min(\hat{z}_N,\epsilon)] \leq (1 + \alpha_1')(B - A)$.

Thus combining the two cases, we have proved the desired result. ∎

First Step of Two-Step Analysis

In the next three lemmas, we accomplish the first part of the proof of Theorem 4.3.8. This is accomplished in the following steps:

(1) We first establish an upper bound for $g(z^*,p^*) - g(z^*,\hat{p}_N)$ [Lemma 4.3.2].

(2) Next, we establish a lower bound for $\rho(z^*,\hat{p}_N)$ [Lemma 4.3.6].

(3) By appropriately choosing β_1, we have $(1-\beta_1)\rho(z^*,\hat{p}_N) \leq \rho(z^*,p^*)$ [Lemma 4.3.4].

Lemma 4.3.2 *Suppose Assumption 4.1.1 holds. Let $(\hat{z}_N, \hat{p}_N)$ be a (α_1', α_2')-accurate realization of $(\hat{Z}_N, \hat{P}_N)$, and (z^*, p^*) be the solution to the first-order conditions (4.5) and (4.6). Then $g(z^*,p^*) - g(z^*,\hat{p}_N) \leq (\alpha_2' + (1 + \alpha_1')(B - A))\,|p^* - \hat{p}_N|$.*

Proof. From the expected profit formula, we know that

$$g(z^*,p^*) - g(z^*,\hat{p}_N) = (a + bc + E[\min(z^*,\epsilon)])(p^* - \hat{p}_N) - b(p^*)^2 + b(\hat{p}_N)^2.$$

Consider the case when $p^* \leq \hat{p}_N$. Then, from Lemma 4.3.1 it follows that

$$\begin{aligned} g(z^*,p^*) - g(z^*,\hat{p}_N) \;\leq\; & \big(a + bc + E[\min(\hat{z}_N,\epsilon)] \\ & -(1 + \alpha_1')(B - A)\big)(p^* - \hat{p}_N) - b(p^*)^2 + b(\hat{p}_N)^2. \end{aligned}$$

Since $(\hat{z}_N, \hat{p}_N)$ is assumed to be (α_1', α_2')-accurate, then

$$\begin{aligned} g(z^*, p^*) - g(z^*, \hat{p}_N) &\leq (2b\hat{p}_N - \alpha_2' - (1+\alpha_1')(B-A))(p^* - \hat{p}_N) - b(p^*)^2 + b(\hat{p}_N)^2 \\ &= -b(p^* - \hat{p}_N)^2 + (\alpha_2' + (1+\alpha_1')(B-A))\,|p^* - \hat{p}_N| \\ &\leq (\alpha_2' + (1+\alpha_1')(B-A))\,|p^* - \hat{p}_N|. \end{aligned}$$

Similarly, we can consider the case when $p^* \geq \hat{p}_N$. Then, from Lemma 4.3.1,

$$\begin{aligned} g(z^*, p^*) - g(z^*, \hat{p}_N) &\leq \big(a + bc + E[\min(\hat{z}_N, \epsilon)] \\ &\quad + (1+\alpha_1')(B-A)\big)(p^* - \hat{p}_N) - b(p^*)^2 + b(\hat{p}_N)^2. \end{aligned}$$

Since $(\hat{z}_N, \hat{p}_N)$ is assumed to be (α_1', α_2')-accurate, then

$$\begin{aligned} g(z^*, p^*) - g(z^*, \hat{p}_N) &\leq (2b\hat{p} + \alpha_2' + (1+\alpha_1')(B-A))(p^* - \hat{p}) - b(p^*)^2 + b(\hat{p}_N)^2 \\ &= -b(p^* - \hat{p}_N)^2 + (\alpha_2' + (1+\alpha_1')(B-A))\,|p^* - \hat{p}_N| \\ &\leq (\alpha_2' + (1+\alpha_1')(B-A))\,|p^* - \hat{p}_N| \end{aligned}$$

Thus, combining both cases, we prove the desired result. ∎

Lemma 4.3.3 *Suppose Assumption 4.1.1 holds. Let (z^*, p^*) be the solution to the first-order conditions (4.5) and (4.6). Then for any $(\hat{z}_N, \hat{p}_N)$, we have $\rho(z^*, \hat{p}_N) \geq 2b\,(p_{min} - c)\,|p^* - \hat{p}_N|$.*

Proof. Define $h(p) = g(z^*, p) = (a + bc + E[\min(z^*, \epsilon)])p - bp^2 - cz^* - ca$, which is the expected profit function for a fixed quantity z^*. Note that $h(p)$ is a concave quadratic function in p that is maximized at p^* (see Figure 4-1). Note that the first derivative is $h'(p) = a + bc + E[\min z^*, \epsilon] - 2bp$.

Consider the case when $\hat{p}_N \leq p^*$. Note that

$$\begin{aligned} h(c) &= cE[\min(z^*, \epsilon)] - cz^*, \\ h'(c) &= a - bc + E[\min(z^*, \epsilon)] = 2b(p^* - c) > 0. \end{aligned}$$

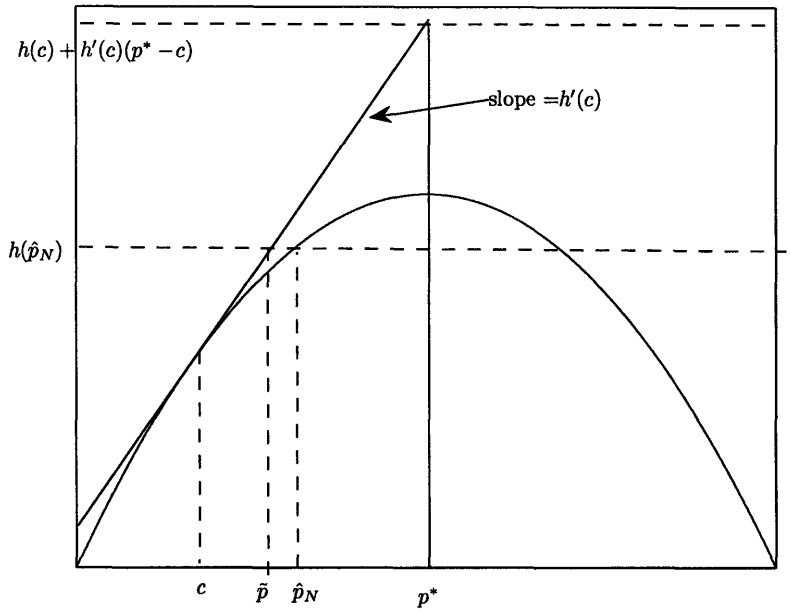

Figure 4-1: Illustration of $h(c)$ for the case when $\hat{p}_N \leq p^*$.

Since $h(p)$ is a concave function (since $h''(p) = -2b < 0$), then

$$h(p) \leq h(c) + h'(c)(p - c), \quad \forall p. \tag{4.19}$$

That is, the linear approximation at $p = c$ is an upper bound for $h(p)$. Since p^* maximizes h, then $h(\hat{p}_N) \leq h(p^*)$. Define $\tilde{p}$ as the intersection of $y = h(\hat{p}_N)$ with $y = h(c) + h'(c)(p - c)$. That is,

$$\tilde{p} = \frac{h(\hat{p}_N) - h(c)}{h'(c)} + c. \tag{4.20}$$

It is obvious from the figure that $\tilde{p} \leq \hat{p}_N \leq p^*$. From equation (4.20), we have

$$h(c) + h'(c)(p^* - c) - h(\hat{p}_N) = h'(c)(p^* - \tilde{p}). \tag{4.21}$$

Since $\tilde{p} \leq \hat{p}_N \leq p^*$ and $h'(c) > 0$, it follows that $h'(c)(p^* - \tilde{p}) \geq h'(c)(p^* - \hat{p}_N)$. Note

that

$$\begin{aligned} h(c) + h'(c)(p^* - c) &= (a + bc + E[\min(z^*, \epsilon)])c - bc^2 - cz^* - ca \\ &\quad +(a - bc + E[\min(z^*, \epsilon)])(p^* - c) \\ &= (a - bc + E[\min(z^*, \epsilon)])p^* + bc^2 - cz^* - ca. \end{aligned}$$

Since $z^* \geq E[\min(z^*, \epsilon)]$, then

$$h(c) + h'(c)(p^* - c) \leq (a - bc + E[\min(z^*, \epsilon)])(p^* - c).$$

Also since $\mu = E[\epsilon] \geq E[\min(z^*, \epsilon)]$ and $p^0 \leq p_{max}$, then

$$h(c) + h'(c)(p^* - c) \leq (a - bc + \mu)(p^0 - c) \leq (a - bc + \mu)(p_{max} - c).$$

Therefore, we have

$$\begin{aligned} \rho(z^*, \hat{p}_N) &= (p_{max} - c)(a - bc + \mu) - h(\hat{p}_N) \\ &\geq h(c) + h'(c)(p^* - c) - h(\hat{p}_N). \\ &= h'(c)(p^* - \tilde{p}) \end{aligned}$$

where the last identity follows from equation (4.21).

Since we have $\tilde{p} \leq \hat{p}_N \leq p^*$, then

$$\begin{aligned} \rho(z^*, \hat{p}_N) &\geq h'(c)(p^* - \hat{p}_N) \\ &= 2b(p^* - c)|p^* - \hat{p}_N| \\ &\geq 2b(p^0 - \frac{B - A}{2b} - c)|p^* - \hat{p}_N|. \end{aligned}$$

where the last inequality follows from the lower bound on p^* in Proposition 4.1.2.

We can make a similar argument for the case $\hat{p}_N \geq p^*$. Since $h(\cdot)$ is a concave quadratic function maximized at p^*, it is easy to verify that $h(2p^* - c) = h(c)$ and $h'(2p^* - c) = -h'(c)$. Because of the symmetry of $h(\cdot)$ around p^*, a parallel argument

can be made for this case.

Thus, combining both cases, we find that $\rho(z^*, \hat{p}_N) \geq 2b(p^0 - \frac{B-A}{2b} - c)|p^* - \hat{p}_N| = 2b(p_{min} - c)|p^* - \hat{p}_N|$. ∎

Therefore, using Lemma 4.3.2 and Lemma 4.3.3, we can accomplish the first step in proving Theorem 4.3.8. This result is stated in the following lemma.

Lemma 4.3.4 *Suppose Assumptions 4.1.1 and 4.1.2 hold. Let $(\hat{z}_N, \hat{p}_N)$ be a (α_1', α_2')-accurate realization of $(\hat{Z}_N, \hat{P}_N)$, and (z^*, p^*) be the solution to the first-order conditions (4.5) and (4.6). Define $\beta_1 = \frac{1}{2b}(p_{min} - c)^{-1}(\alpha_2' + (1 + \alpha_1')(B - A))$. Then $(1 - \beta_1)\rho(z^*, \hat{p}_N) \leq \rho(z^*, p^*)$.*

Proof. From Lemma 4.3.2, we find that

$$\begin{aligned} \rho(z^*, \hat{p}_N) - \rho(z^*, p^*) &= g(z^*, p^*) - g(z^*, \hat{p}_N) \\ &\leq (\alpha_2' + (1 + \alpha_1')(B - A))|p^* - \hat{p}_N|. \end{aligned}$$

Also, from Lemma 4.3.3, it follows that, for any $\beta_1 > 0$,

$$\beta_1 \rho(z^*, \hat{p}_N) \geq 2b\beta_1(p_{min} - c)|p^* - \hat{p}_N|.$$

Suppose we choose $\beta_1 = \frac{1}{2b}(p_{min} - c)^{-1}(\alpha_2' + (1 + \alpha_1')(B - A))$. Note that from Assumption 4.1.2, it follows that $\beta_1 > 0$. Thus our choice is valid. Rearranging terms, we have $(1 - \beta_1)\rho(z^*, \hat{p}_N) \leq \rho(z^*, p^*)$. ∎

Second Step of the Two-Step Analysis

Now, for the second step, we want to show that there exists $\beta_2 > 0$, such that $(1 - \beta_2)\rho(\hat{z}_N, \hat{p}_N) \leq \rho(z^*, \hat{p}_N)$. The way that we show this is similar to the proof in the first part. We outline the steps as follows:

(1) We first establish an upper bound for $g(z^*, \hat{p}_N) - g(\hat{z}_N, \hat{p}_N)$ [Lemma 4.3.5].

(2) Next, we establish a lower bound for $\rho(\hat{z}_N, \hat{p}_N)$ [Lemma 4.3.6].

(3) By appropriately choosing β_2, it follows that $(1-\beta_2)\rho(\hat{z}_N, \hat{p}_N) \leq \rho(z^*, \hat{p}_N)$ [Lemma 4.3.7].

Lemma 4.3.5 *Let $(\hat{z}_N, \hat{p}_N)$ be a (α_1', α_2')-accurate realization of $(\hat{Z}_N, \hat{P}_N)$, and (z^*, p^*) be the solution to the first-order conditions (4.5) and (4.6). Then $g(z^*, \hat{p}_N) - g(\hat{z}_N, \hat{p}_N) \leq \alpha_1' \hat{p}_N |z^* - \hat{z}_N|$.*

Proof. We know that $g(z,p)$ is a concave function with respect to z. Therefore, the following inequality holds

$$g(z^*, \hat{p}_N) - g(\hat{z}_N, \hat{p}_N) \leq \frac{\partial g}{\partial z}(\hat{z}_N, \hat{p}_N)(z^* - \hat{z}_N).$$

Since $(\hat{z}_N, \hat{p}_N)$ is (α_1', α_2')-accurate, therefore $\left|\frac{\partial g}{\partial z}(\hat{z}_N, \hat{p}_N)\right| \leq \hat{p}_N \alpha_1'$. Then it follows that

$$g(z^*, \hat{p}_N) - g(\hat{z}_N, \hat{p}_N) \leq \alpha_1' \hat{p}_N |z^* - \hat{z}_N|. \quad \blacksquare$$

Lemma 4.3.6 *Suppose Assumptions 4.1.1 and 4.1.2 hold. Also suppose that there exists $\gamma > 0$ such that $|\hat{p}_N^0 - p^0| \leq \gamma$ and $p_{min} - \gamma > c$. Let $(\hat{z}_N, \hat{p}_N)$ be a realization of $(\hat{Z}_N, \hat{P}_N)$, and (z^*, p^*) be the solution to the first-order conditions (4.5) and (4.6). Then we have $\rho(\hat{z}_N, \hat{p}_N) \geq \frac{1}{2} \min\left(\frac{c}{p^0+\gamma}, 1 - \frac{c}{p_{min}-\gamma}\right) \hat{p}_N |z^* - \hat{z}_N|$.*

Proof. Define $h(z) = g(z, \hat{p}_N) = (\hat{p}_N - c)(a - b\hat{p}_N) - cz - ca + \hat{p}_N E[\min(z, \epsilon)]$. Note that $h(z)$ is a concave function. Also define the function $f(z) = (\hat{p}_N - c)(a - b\hat{p}_N) - cz + \hat{p}_N \min(z, \mu)$. By Jensen's inequality, we know that $E[\min(z, \epsilon)] \leq \min(z, \mu)$. Therefore, $f(z) \geq h(z)$ for all z.

Note that

$$f(z) = \begin{cases} (\hat{p}_N - c)(a - b\hat{p}_N + z), & \text{if } z \leq \mu, \\ (\hat{p}_N - c)(a - b\hat{p}_N) + \hat{p}_N \mu - cz, & \text{if } z > \mu. \end{cases}$$

Let z_a be the intersection of $y = h(\hat{z}_N)$ with $y = (\hat{p}_N - c)(a - b\hat{p}_N + z)$, and z_b be its intersection with $y = (\hat{p}_N - c)(a - b\hat{p}_N) + \hat{p}_N \mu - cz$. Since $f(z) \geq h(z)$ and both functions are concave, then it is easy to verify that $|z^* - \hat{z}_N| \leq z_b - z_a$.

The maximum of $f(z)$ is achieved when $z = \mu$. And since $f(z) \geq h(z)$ for all z, then $f(\mu) \geq h(\hat{z}_N)$. Also note that

$$\begin{aligned} f(\mu) - h(\hat{z}_N) &= (\hat{p}_N - c)(\mu - z_a), \\ f(\mu) - h(\hat{z}_N) &= c(z_b - \mu). \end{aligned}$$

Therefore, since we know that $z_b - z_a \geq |z^* - \hat{z}_N|$, it follows that

$$f(\mu) - h(\hat{z}_N) \geq \frac{1}{2}\min(\hat{p}_N - c, c)(z_b - z_a) \geq \frac{1}{2}\min\left(\frac{c}{\hat{p}_N}, 1 - \frac{c}{\hat{p}_N}\right)\hat{p}_N|z^* - \hat{z}_N|.$$

Also, from Proposition 4.2.1, we know that $\hat{p}_N \leq \hat{p}_N^0 \leq p^0 + \gamma$ and $\hat{p}_N \geq \hat{p}_N^0 - p^0 + p_{min} \geq p_{min} - \gamma$. Therefore,

$$f(\mu) - h(\hat{z}_N) \geq \frac{1}{2}\min\left(\frac{c}{p^0 + \gamma}, 1 - \frac{c}{p_{min} - \gamma}\right)\hat{p}_N|z^* - \hat{z}_N|.$$

Since $c \leq \hat{p}_N \leq p_{max}$, then $f(\mu) = (\hat{p}_N - c)(a - b\hat{p}_N + \mu) \leq (p_{max} - c)(a - b\hat{p}_N + \mu \leq (p_{max} - c)(a - bc + \mu)$. Therefore, it follows that

$$\begin{aligned} \rho(\hat{z}_N, \hat{p}_N) &= (p_{max} - c)(a - bc + \mu) - g(\hat{z}_N, \hat{p}_N) \\ &\geq f(\mu) - h(\hat{z}_N) \\ &\geq \frac{1}{2}\min\left(\frac{c}{p^0 + \gamma}, 1 - \frac{c}{p_{min} - \gamma}\right)\hat{p}_N|z^* - \hat{z}_N|. \quad \blacksquare \end{aligned}$$

Lemma 4.3.7 *Suppose Assumptions 4.1.1 and 4.1.2 hold. Also suppose that* $|\hat{p}_N^0 - p^0| \leq \gamma$, *where* $\gamma > 0$ *such that* $p_{min} - \gamma > c$. *Let* $(\hat{z}_N, \hat{p}_N)$ *be a* (α_1', α_2')*-accurate realization of* $(\hat{Z}_N, \hat{P}_N)$, *and* (z^*, p^*) *be the optimal solution to price-setting newsvendor problem (4.4). Define* $\beta_2 = 2\alpha_1'\left[\min\left(\frac{c}{p^0+\gamma}, 1 - \frac{c}{p_{min}-\gamma}\right)\right]^{-1}$. *Then* $(1-\beta_2)\rho(\hat{z}_N, \hat{p}_N) \leq \rho(z^*, \hat{p}_N)$.

Proof. From Lemma 4.3.5, we find that

$$\begin{aligned} \rho(\hat{z}_N, \hat{p}_N) - \rho(z^*, \hat{p}_N) &= g(z^*, \hat{p}_N) - g(\hat{z}_N, \hat{p}_N) \\ &\leq \alpha_1'\hat{p}_N|z^* - \hat{z}_N|. \end{aligned}$$

Also, from Lemma 4.3.6, it follows that, for any $\beta_2 > 0$,

$$\beta_2 \rho(\hat{z}_N, \hat{p}_N) \geq \frac{\beta_2}{2} \min\left(\frac{c}{p^0+\gamma}, 1 - \frac{c}{p_{min}-\gamma}\right) \hat{p}_N |z^* - \hat{z}_N|.$$

Suppose we choose $\beta_2 = 2\alpha_1' \left[\min\left(\frac{c}{p^0+\gamma}, 1 - \frac{c}{p_{min}-\gamma}\right)\right]^{-1}$. Note that since $p_{min} - \gamma > c$, we have $\beta_2 > 0$. Thus, our choice is valid. Rearranging terms, we have $(1-\beta_2)\rho(\hat{z}_N, \hat{p}_N) \leq \rho(z^*, \hat{p}_N)$. ∎

Using Lemma 4.3.4 and Lemma 4.3.7, we can now prove Theorem 4.3.8 which shows that if a policy $(\hat{z}_N, \hat{p}_N)$ is (α_1', α_2')-accurate, then the relative error of its regret is bounded.

Theorem 4.3.8 *Suppose Assumptions 4.1.1–4.1.3 hold. Suppose that $|\hat{p}_N^0 - p^0| \leq \gamma$, where $\gamma > 0$ such that $p_{min} - \gamma > c$. Let $(\hat{z}_N, \hat{p}_N)$ be a (α_1', α_2')-accurate realization of $(\hat{Z}_N, \hat{P}_N)$, and (z^*, p^*) be the optimal solution to price-setting newsvendor problem. Define $\beta_1 = \frac{1}{2b}(p_{min}-c)^{-1}(\alpha_2' + (1+\alpha_1')(B-A))$, $\beta_2 = 2\alpha_1' \left[\min\left(\frac{c}{p^0+\gamma}, 1 - \frac{c}{p_{min}-\gamma}\right)\right]^{-1}$ and $\alpha = \frac{1}{(1-\beta_1)(1-\beta_2)} - 1$. If $\beta_1 < 1$ and $\beta_2 < 1$, then $\rho(\hat{z}_N, \hat{p}_N) \leq (1+\alpha)\rho(z^*, p^*)$.*

Proof. Since $\beta_1 < 1$ and $\beta_2 < 1$, then from Lemma 4.3.7, we have

$$(1-\beta_1)(1-\beta_2)\rho(\hat{z}_N, \hat{p}_N) \leq (1-\beta_1)\rho(z^*, \hat{p}_N).$$

And since from Lemma 4.3.2, we have

$$(1-\beta_1)\rho(z^*, \hat{p}_N) \leq \rho(z^*, z^*),$$

then it follows that $(1-\beta_1)(1-\beta_2)\rho(\hat{z}_N, \hat{p}_N) \leq \rho(z^*, p^*)$. Therefore, if we define $\alpha = \frac{1}{(1-\beta_1)(1-\beta_2)} - 1$, then $\rho(\hat{z}_N, \hat{p}_N) \leq (1+\alpha)\rho(z^*, p^*)$. ∎

Theorem 4.3.8 establishes the relationship between (α_1', α_2')-accuracy and bounds on the relative error of the regret for the solution to the data-driven approximation to

the first-order conditions. Note that the theorem holds if $(\hat{z}_N, \hat{p}_N)$ is (α_1', α_2')-accurate and $|p^0 - \hat{p}^0| \leq \gamma$ (or equivalently, $|\mu - \hat{\mu}^0| \leq 2b\gamma$). Intuitively, the probability that these conditions hold increases by taking more samples of the stochastic variable. The next section establishes the connection between the sample size and the conditions for which Theorem 4.3.8 hold.

4.3.2 Sample Size to Ensure First-Order Accuracy of the Solution to the Simulation-Based Procedure

In this subsection, we establish a relationship between the sample size of the data-driven approximation to the first-order conditions and the first-order accuracy of its solution. We will be using Hoeffding's inequality (see Theorem 2.3.4). The following lemma states this relationship.

Lemma 4.3.9 *Suppose Assumption 4.1.1 holds. For each $\alpha_1', \alpha_2' > 0$ and $0 < \delta_1, \delta_2 < 1$ such that $0 < \delta_1 + \delta_2 < 1$, if the number of samples is $N \geq \max(N_1, N_2)$, where*

$$\begin{aligned} N_1 &= \frac{(B-A)^2}{2\alpha_2'^2} \log\left(\frac{2}{\delta_2}\right) \\ N_2 &= \frac{1}{2\alpha_1'^2} \log\left(\frac{2(p_{max} - c)}{\delta_1}\right) \end{aligned}$$

then the solution to the data-driven approximation to the first-order conditions $(\hat{Z}_N, \hat{P}_N)$ is (α_1', α_2')-accurate with probability of at least $1 - \delta_1 - \delta_2$.

Proof. Note that each given sample of size N of the random component ϵ implies a $(\hat{z}_N, \hat{p}_N)$ that solves the system of linear equations (4.13) and (4.14). Define the event

$$B = \left[F(\hat{z}_N) < 1 - \frac{c}{\hat{p}_N} - \alpha_1'\right].$$

For an arbitrary price $\hat{p}_N^l$, we can define the event conditional on $\hat{p}_N = \hat{p}_N^l$

$$B(\hat{p}_N^l) = \left[F(\hat{z}_N) < 1 - \frac{c}{\hat{p}_N^l} - \alpha_1' \middle| \hat{p}_N = \hat{p}_N^l\right].$$

Define the quantile

$$z_1(\hat{p}_N^l) = \inf\left\{z \middle| F(z) = 1 - \frac{c}{\hat{p}_N^l} - \alpha_1'\right\}.$$

Since $F(\cdot)$ is nondecreasing, then $B(\hat{p}_N^l) = \left[\hat{z}_N < z_1\left(\hat{p}_N^l\right) \middle| \hat{p}_N = \hat{p}_N^l\right]$.

Since $(\hat{z}_N, \hat{p}_N)$ satisfies the data-driven approximation of the first-order conditions, then conditional on $\hat{p}_N = \hat{p}_N^l$, we have $\hat{F}_N(\hat{z}_N) = 1 - \frac{c}{\hat{p}_N^l}$. Thus, since $\hat{F}_N(\cdot)$ is nondecreasing, we have

$$\begin{aligned}
\Pr(B(\hat{p}_N^l)) &= \Pr\left(\hat{z} < z_1(\hat{p}_N^l) | \hat{p}_N = \hat{p}_N^l\right) \\
&\leq \Pr\left(\hat{F}_N(z_1(\hat{p}_N^l)) \geq 1 - \frac{c}{\hat{p}_N^l} \middle| \hat{p}_N = \hat{p}_N^l\right).
\end{aligned}$$

Using Bayes Theorem,

$$\begin{aligned}
\Pr(B(\hat{p}_N^l)) &\leq \frac{\Pr\left(\hat{F}_N(z_1(\hat{p}_N^l)) \geq 1 - \frac{c}{\hat{p}_N^l};\ \hat{p}_N = \hat{p}_N^l\right)}{\Pr\left(\hat{p}_N = \hat{p}_N^l\right)} \\
&\leq \frac{\Pr\left(\hat{F}_N(z_1(\hat{p}_N^l)) \geq 1 - \frac{c}{\hat{p}_N^l}\right)}{\Pr\left(\hat{p}_N = \hat{p}_N^l\right)}.
\end{aligned}$$

From the definition of $z_1(\hat{p}_N^l)$, we have

$$\begin{aligned}
\Pr(B(\hat{p}_N^l)) &\leq \frac{\Pr\left(\hat{F}_N(z_1(\hat{p}_N^l)) - F(z_1(\hat{p}_N^l)) \geq \alpha_1'\right)}{\Pr\left(\hat{p}_N = \hat{p}_N^l\right)} \\
&\leq \frac{\exp(-2N\alpha_1'^2)}{\Pr\left(\hat{p}_N = \hat{p}_N^l\right)}
\end{aligned}$$

where the last inequality follows from Hoeffding's inequality. Therefore, by choosing $N \geq \frac{1}{2\alpha_1'^2} \log\left(\frac{2(p_{max}-c)}{\delta_1}\right)$, we have $\Pr(B(\hat{p}_N^l)) \leq \frac{\delta_1}{2(p_{max}-c)} \frac{1}{\Pr\left(\hat{p}_N = \hat{p}_N^l\right)}$.

With this choice of the sample size, since $c \leq \hat{p}_N \leq p_{max}$, note that

$$\begin{aligned}
\Pr(B) &= \int_c^{p_{max}} \Pr(B(\hat{p}_N^l)) \Pr\left(\hat{p}_N = \hat{p}_N^l\right) d\hat{p}_N^l \\
&\leq \int_c^{p_{max}} \frac{\delta}{2(p_{max} - c)} d\hat{p}_N^l = \frac{\delta_1}{2}.
\end{aligned} \tag{4.22}$$

Now define the event

$$L = \left[F(\hat{z}_N) > 1 - \frac{c}{\hat{p}_N} + \alpha_1' \right].$$

We can do the same analysis as we did for event B and apply Hoeffding's inequality. Choosing $N \geq \frac{1}{2\alpha_1'^2} \log\left(\frac{2(p_{max}-c)}{\delta_1}\right)$, we have $\Pr(L) = \frac{\delta_1}{2}$.

Now define the event

$$T = \left[\hat{E}_N[\min(z,\epsilon)] - E[\min(z,\epsilon)] > \alpha_2' \right],$$

where $\hat{E}_N[\min(z,\epsilon)]$ is given by equation (4.17). Consider the random variable $X^k = \mathbf{1}(\epsilon^k \leq z)\epsilon^k + \mathbf{1}(\epsilon^k > z)z$. Note that $\hat{E}_N[\min(z,\epsilon)] = \frac{1}{N}\sum_{k=1}^{N} X^k$. Since $A \leq X^k \leq B$, and

$$\begin{aligned} E[X^k] &= \Pr(\epsilon \leq z)E[\epsilon|\epsilon \leq z] + \Pr(\epsilon > z)E[z|\epsilon > z] \\ &= \Pr(\epsilon \leq z)E[\min(z,\epsilon)|\epsilon \leq z] + \Pr(\epsilon > z)E[\min(z,\epsilon)|\epsilon > z] \\ &= E[\min(z,\epsilon)], \end{aligned}$$

then we can apply Hoeffding's inequality. Therefore,

$$\Pr(T) = \Pr\left(\hat{E}_N[\min(z,\epsilon)] - E[\min(z,\epsilon)] > \alpha_2' \right) \leq e^{-2\alpha_2'^2 N/(B-A)^2}.$$

Thus, if we choose $N \geq \frac{(B-A)^2}{2\alpha_2'^2} \log\left(\frac{2}{\delta_2}\right)$, then $\Pr(T) \leq \frac{\delta_2}{2}$.

Now define the event

$$V = \left[E[\min(z,\epsilon)] - \hat{E}_N[\min(z,\epsilon)] > \alpha_2' \right].$$

We follow an analysis similar to that for event T and apply Hoeffding's inequality, we find that for $N \geq \frac{(B-A)^2}{2\alpha_2'^2} \log\left(\frac{2}{\delta_2}\right)$, we have $\Pr(V) \leq \frac{\delta_2}{2}$.

Therefore, if we choose

$$N \geq \max\left(\frac{(B-A)^2}{2\alpha_2'^2}\log\left(\frac{2}{\delta_2}\right), \frac{1}{2\alpha_1'^2}\log\left(\frac{2}{\delta_1}\right)\right),$$

then we find that

$$\begin{aligned}\Pr(B \cup L \cup T \cup V) &\leq \Pr(B) + \Pr(L) + \Pr(T) + \Pr(V) \\ &\leq \delta_1 + \delta_2.\end{aligned}$$

Thus, it follows that $\Pr\{(\hat{Z}_N, \hat{P}_N) \text{ is } (\alpha_1', \alpha_2')\text{-accurate}\} \geq 1 - \delta_1 - \delta_2$. ∎

Before we can combine the results from Theorem 4.3.8 and Theorem 4.3.9, note that an additional assumption for Theorem 4.3.8 is that $|\hat{p}^0 - p^0| \leq \gamma$, for some $\gamma > 0$. Therefore, we first need the following lemma.

Lemma 4.3.10 *Suppose Assumption 4.1.1 holds. For each $\gamma > 0$ and $0 < \delta_3 < 1$, if the number of samples is $N \geq N_3 = \frac{(B-A)^2}{8b^2\gamma^2}\log\left(\frac{2}{\delta_3}\right)$, then $|p^0 - \hat{p}_N^0| \leq \gamma$ with probability of at least $1 - \delta_3$.*

Proof. Suppose there are N samples of ϵ, given by $\epsilon^1, \ldots, \epsilon^N$. Define

$$X^k = \frac{a + bc + \epsilon^k}{2b},$$

where $E[X^k] = \frac{a+bc+\mu}{2b} = p^0$. Note that $\hat{p}_N^0 = \frac{1}{N}\sum_{k=1}^N X^k$. Since $X^k \in [\alpha, \beta]$ where we define $\alpha = \frac{a+bc+A}{2b}$ and $\beta = \frac{a+bc+B}{2b}$, then we can apply Hoeffding's inequality to find

$$\Pr(\hat{p}_N^0 - p^0 \geq \gamma) \leq \exp\left(\frac{-2\gamma^2 N}{(\beta-\alpha)^2}\right) = \exp\left(\frac{-8b^2\gamma^2 N}{(B-A)^2}\right).$$

Thus, choosing $N \geq \frac{(B-A)^2}{8b^2\gamma^2}\log\left(\frac{2}{\delta_3}\right)$, then it follows that

$$\Pr(\hat{p}_N^0 - p^0 \geq \gamma) \leq \frac{\delta_3}{2}.$$

We can perform a similar proof for the probability that $\hat{p}_N^0 - p^0 \leq -\gamma$. Therefore, we have for $N \geq \frac{(B-A)^2}{8b^2\gamma^2} \log\left(\frac{2}{\delta_3}\right)$ that

$$\Pr(|p^0 - \hat{p}_N^0| \leq \gamma) \geq 1 - \delta_3. \quad \blacksquare$$

The main result of this chapter is the following Theorem:

Theorem 4.3.11 *Suppose Assumptions 4.1.1–4.1.3 hold. Further suppose that there exists some $\gamma > 0$ such that $p_{min} - \gamma > c$. Consider the price-setting newsvendor problem specified by a unit cost c and a demand $D(p, \epsilon) = a - bp + \epsilon$, where ϵ is the random component. Let $0 < \beta_1, \beta_2 < 1$ be specified accuracy levels such that $\alpha = \frac{1}{(1-\beta_1)(1-\beta_2)} - 1 \in (0, 1]$, and $0 < \delta < 1$ (for $0 < \delta < 1$)be a specified confidence level. Suppose for the choice of accuracy levels, the following condition holds:*

$$\alpha_2' \equiv 2b\beta_1(p_{min} - c) - (1 + \alpha_1')(B - A) > 0,$$

where

$$\alpha_1' \equiv \frac{1}{2}\beta_2 \min\left(\frac{c}{p^0 + \gamma}, 1 - \frac{c}{p_{min} - \gamma}\right).$$

Suppose the data-driven approximation to the first-order conditions is solved with respect to N i.i.d. samples of ϵ, and that $N \geq \max(N_1, N_2, N_3)$, where

$$\begin{aligned} N_1 &= \frac{(B-A)^2}{2\alpha_2'^2} \log\left(\frac{5}{\delta}\right) \\ N_2 &= \frac{1}{2\alpha_1'^2} \log\left(\frac{5(p_{max} - c)}{\delta}\right), \\ N_3 &= \frac{(B-A)^2)}{8b^2\gamma^2} \log\left(\frac{5}{\delta}\right). \end{aligned}$$

Let $(\hat{Z}_N, \hat{P}_N)$ be the solution to the data-driven approximation to the first-order conditions and let $(\hat{z}_N, \hat{p}_N)$ be denote its realization. Then, with probability at least $1 - \delta$, the regret of $(\hat{Z}_N, \hat{P}_N)$ is at most $1 + \alpha$ times the regret of an optimal solution (z^*, p^*) to the price-setting newsvendor problem. In other words, $\rho(\hat{z}_N, \hat{p}_N) \leq (1 + \alpha)\rho(z^*, p^*)$ with probability of at least $1 - \delta$.

Proof. Since Assumption 4.1.2 holds, then there exists some $\gamma > 0$ such that $p_{min} - \gamma > c$. Let $\delta_1 = \frac{2\delta}{5}$, $\delta_2 = \frac{2\delta}{5}$ and $\delta_3 = \frac{\delta}{5}$. From Lemma 4.3.10, we know that for $N \geq N_3$

$$\Pr\left(|\hat{p}_N^0 - p^0| > \gamma\right) \leq \delta_3.$$

Also from Lemma 4.3.8, we know that for $N \geq \max(N_1, N_2)$,

$$\Pr\left((\hat{Z}_N, \hat{P}_N)\text{is not}(\alpha_1', \alpha_2')\text{-accurate}\right) \leq \delta_1 + \delta_2.$$

Thus, we have for $N \geq \max(N_1, N_2, N_3)$

$$\begin{aligned}
\Pr((\hat{Z}_N, \hat{P}_N) \quad &\text{is} \quad (\alpha_1', \alpha_2')\text{-accurate and } |\hat{p}_N^0 - p^0| \leq \gamma) \\
&\geq \quad 1 - \Pr((\hat{Z}_N, \hat{P}_N) \text{ is not } (\alpha_1', \alpha_2')\text{-accurate}) \\
&\quad\quad - \Pr(|\hat{p}_N^0 - p^0| > \gamma) \geq 1 - \delta_1 - \delta_2 - \delta_3 = 1 - \delta.
\end{aligned}$$

From Theorem 4.3.8, we know that if $|\hat{p}_N^0 - p^0| \leq \gamma$ and if $(\hat{Z}_N, \hat{P}_N)$ is (α_1', α_2')-accurate, then $\rho(\hat{Z}_N, \hat{P}_N) \leq (1 + \alpha)\rho(z^*, p^*)$, where (z^*, p^*) is the solution to the first-order equations. Moreover, under Assumptions 4.1.2–4.1.3, we know that the first-order conditions (4.5) and (4.6) are uniquely solved by the optimal solution of the price-setting newsvendor. Thus, we have for $N \geq \max(N_1, N_2, N_3)$ the solution to the data-driven approximation of the first-order conditions satisfies

$$\Pr(\rho(\hat{Z}_N, \hat{P}_N) \leq (1 + \alpha)\rho(z^*, p^*)) \geq 1 - \delta,$$

where (z^*, p^*) is the optimal solution to the price-setting newsvendor problem. ∎

Theorem 4.3.11 outlines a procedure to choose a sample size that guarantees that the regret of the solution to the data-driven approximation to the first-order conditions satisfies a specified accuracy level with a specified confidence level. Compared to the bound derived for the classical newsvendor problem in Theorem 2.3.6, note that in the above theorem, accuracy depends on two parameters β_1 and β_2. The introduction of these two parameters correspond to accuracy in each of the two dimensions

(i.e., z and p).

N_1 and N_2 are the bounds needed to ensure that $(\hat{z}_N, \hat{p}_N)$ is (α_1', α_2')-accurate with high probability. On the other hand, N_3 ensures that the sample size is large enough such that the true mean μ and the sample mean $\hat{\mu}$ are γ-close with high probability. Note that N_2 is similar to the Hoeffding bound for the classical newsvendor problem, except for the term $p_{max} - c$. This term appears, because unlike in the classical model, p^* (and $\hat{p}_N$) is not fixed and belongs to a range $[c, p_{max}]$. Similarly, note that the term $(B - A)$ in N_3 appears because z^* and $\hat{z}_N$ belong in the range $[A, B]$.

Chapter 5

Numerical Results for the Data-Driven Price-Setting Newsvendor Problem

In Chapter 4, we described a simulation-based procedure proposed by Zhan and Shen [42] for solving the price-setting newsvendor problem if the explicit demand distribution is not known. Under their model, the demand is composed of a deterministic component and a random component. The deterministic demand is assumed to be a linear function of the price with known parameters. That is, the stochastic demand is $D(p, \epsilon) = a - bp + \epsilon$, where a is the zero-price demand level, b is the price-sensitivity parameter and ϵ is the random component. Under the simulation-based approach, it is assumed that a set of N independent samples of the demand is available. In Chapter 4, we prove under certain assumptions on the random component of the demand (e.g., boundedness) that the sample size is related to the accuracy of the regret achieved by the solution to the procedure with respect to the optimal regret computed with respect to the true distribution. However, the bound is only a worst-case bound on the number of samples required. Moreover, bounded accuracy may also be achieved for distributions that violate the assumptions in Chapter 4.

In this chapter, we conduct computational experiments to evaluate the performance of the simulation-based procedure under various concrete scenarios. In partic-

ular:

- We explore how the simulation-based procedure performs as the unit cost c varies.
- We explore how the simulation-based procedure performs as the zero-price demand level a varies.
- We explore how the procedure performs as the price-sensitivity parameter b varies.
- We explore the sensitivity of the procedure to different parameters of concrete distributions.

The remainder of the chapter is outlined as follows. In Section 5.1, the experimental design is described for the various computational experiments. In Section 5.2, the results of the experiments are presented and discussed.

5.1 Methodology

We outline the steps of the simulation-based procedure by Zhan and Shen [42] as we apply it in our numerical experiments.

- Obtain a set of N samples $\epsilon^1, \ldots, \epsilon^N$ from the true demand distribution (see pg. 78 on how to obtain these samples from demand-price data).
- Perform a gradient search with initial guess $(\hat{\mu}, \hat{p}^0)$, where $\hat{\mu}$ is a guess on the expected value of the random component and $\hat{p}^0$ is a guess on the optimal riskless price (see pg. 75 for definition). We outline the gradient procedure:

  1: Set tolerance level, δ_e
  2: Set $(z_{old}, p_{old}) \leftarrow (\hat{\mu}, \hat{p}^0)$
  3: **loop**
  4: Compute $dz = -c + p_{old}\left(1 - \frac{1}{N}\sum_{k=1}^{N} \mathbf{1}(\epsilon^k \leq z_{old})\right)$
  5: Compute $dp = a + bc - 2bp_{old} + \frac{1}{N}\sum_{k=1}^{N} \min(\epsilon^k, z_{old})$

```
    Choose stepsize s by Armijo rule
    Set z_new ← z_old + s dz
    Set p_new ← p_old + s dp
    if (z_old − z_new)^2 + (p_old − p_new)^2 ≤ δ_e^2 then
        Stop the loop
    else
        Set (z_old, p_old) ← (z_new, p_new)
    end if
end loop
```

- The first two steps are repeated for a total of m times.

The optimal solution to the SAA counterpart with N samples, $(\hat{Z}_N, \hat{P}_N)$, is a random variable that is dependent on the specific N samples of the random component of demand ϵ. The procedure outlined above is done for a total of m times to approximate a distribution for $(\hat{Z}_N, \hat{P}_N)$. Note also that different stepsize selection schemes can be adopted, such as Armijo rule or diminishing stepsize rule (see Bertsekas [7] for more details on stepsize selection). In these numerical examples we will be adopting an Armijo stepsize selection rule.

In the experiments, we take $N = 100$ samples for $m = 100$ runs of the procedure. We generate data from four distributions: uniform, normal, exponential, Pareto and Poisson. In the base case, the zero-price demand level $a = 200$, the price-sensitivity parameter $b = 35$ and the unit cost $c = 1$. The specifications of the distributions of the base case are as follows:

Uniform	range of $[-10, 2\mu + 10]$ with $\mu = 10$
Normal	mean $\mu = 10$ and standard deviation $\sigma = 25$
Exponential	mean $\mu = 10$
Pareto	mean $\mu = 10$ and range $[1, \infty)$

In analyzing the influence of the unit cost c, we consider various unit costs c. We take $c \in [1, 1.5, \ldots, 4]$.

In analyzing the influence of the zero-price demand level a, we consider $a \in [100, 200, \ldots, 1000]$.

In analyzing the influence of the price-sensitivity parameters b, we take $b \in [10, 15, \ldots, 50]$.

The sensitivity analysis of the simulation-based procedure to various parameters of concrete distributions is performed with the following specifications

Uniform	$\mu \in [10, 20, \ldots, 100]$
Normal	$\mu \in [10, 20, \ldots, 100]$
	$\sigma \in [10, 15, \ldots, 50]$
Exponential	$\mu \in [10, 20, \ldots, 100]$
Pareto	$k \in [2, 3, \ldots, 50]$

Since the distributions are not all bounded, we need to define a new regret function. In this chapter, the regret function we consider is $\rho(z,p) = (p^0 - c)(a - bc + \mu) - g(z,p)$, where p^0 is the riskless optimal price and $g(z,p)$ is the expected profit. The performance of the simulation is evaluated by observing the relative error of the average expected profit and average regret. In particular, if $(\hat{z}_N^k, \hat{p}_N^k)$ is the solution to the simulation-based procedure for the kth run for $k = 1, \ldots, m$, then

$$\begin{aligned}
\text{Average Relative Error of the Expected Profit} &= \frac{g(z^*,p^*) - \frac{1}{m}\sum_{k=1}^{m} g(\hat{z}_N^k, \hat{p}_N^k)}{g(z^*,p^*)}, \\
\text{Average Relative Error of the Regret} &= \frac{\frac{1}{m}\sum_{k=1}^{m} \rho(\hat{z}_N^k, \hat{p}_N^k) - \rho(z^*,p^*)}{\rho(z^*,p^*)},
\end{aligned}$$

where (z^*, p^*) is the optimal ordering and pricing policy under full knowledge of the distribution.

5.2 Results and Discussion

In this section, we present the results of our computational experiments. We also discuss these results and mention insights gained from the experiments.

Sensitivity to the Unit Cost

Figure A-12 plots the relative error in the expected profit and the regret achieved by the SAA solution as a function of the unit cost. We make the following observations:

- From Figure A-12, we can observe that the relative error in the expected profit and the regret achieved by the SAA solution is increasing against c for most of the distributions. The relative error achieved by the Pareto distribution is relatively invariant against c.

- We also observe that the relative error of the regret and the relative error of the expected profit are highly correlated as c varies.

- The largest relative error of the expected profit is of the order 10^{-1}. On the other hand, the largest relative error of the regret is of the order 10^{-2}.

Sensitivity to the Zero-Price Demand Level

Figure A-13 shows the relative error of the regret and the expected profit as a function of the zero-price demand level a. We make the following observations:

- We can observe a decreasing trend in the relative error of the expected profit achieved by the SAA counterpart as the zero-price level increases (Figure A-13(a)). A similar decreasing trend can be observed for the relative error of the regret (Figure A-13(b)). Again, the relative error of the Pareto distribution is relatively consistent under different values of a.

- The relative error of the expected profit and the regret are highly correlated as a varies.

- In both the regret and expected profit of the SAA solution, the relative error is small (less than 10^{-2}).

Sensitivity to the Price-Sensitivity Parameter

Figure A-14 plots the relative error in the expected profit and the regret achieved by the SAA solution as a function of b.

We make the following observations:

- The relative error in the expected profit and the regret of all distributions except Pareto exhibit an increasing trend as the price-sensitivity b increases. Pareto on the other hand exhibits an opposite and decreasing trend.
- The relative error of the regret and the relative error of the expected profit exhibit high correlation as b is varied.
- The relative errors of the expected profit and regret are at most of the order 10^{-3}.

Sensitivity to Parameters of the Distribution

Figure A-15 plots the relative error of the regret and the expected profit of the SAA solution as a function of the mean μ of a uniform distribution.

Figure A-16 plots the relative error of the regret and the expected profit of the SAA solution as a function of the mean μ of an exponential distribution.

Figure A-17 plots the relative error of the regret and the expected profit of the SAA solution as a function of the mean μ of a normal distribution.

Figure A-18 plots the relative error of the regret and the expected profit of the SAA solution as a function of the standard deviation σ of the normal distribution.

Figure A-19 plots the relative error of the regret and the expected profit of the SAA solution as a function of the parameter k of the Pareto distribution.

We make the following observations:

- The relative error of the regret and the expected profit are both exhibit an increasing trend as the following distribution parameters are increased: μ of the uniform distribution (Figure A-15), μ of the exponential distribution (Figure A-16), and σ of the normal distribution (Figure A-18).

- The relative error of the regret and of the expected profit exhibit a decreasing trend for increasing: μ of the normal distribution and k of the Pareto distribution.

- In all experiments where the distribution parameters are varied, the relative error of both the expected profit and regret are at most of the order 10^{-3}.

- There is a steep decreasing trend in the relative error (in regret and expected profit) of the Pareto distribution for small values of k. Moreover, for $k = 2$ (where the variance is infinite), the relative error is less than 10^{-4}. We can infer that the simulation-based procedure performs well under an infinite second moment.

- Again, the relative error in the expected profit and the relative error in the regret exhibit high correlation for all the computational experiments where demand parameters are varied.

From our discussion and observations, we can conclude that, unlike the SAA approach for the classical newsvendor problem, the performance of the simulation-based approach is highly dependent on the model parameters. Also, from the high correlation between the relative error of the regret and the expected profit, we can infer that the accuracy with respect to the expected profit is a valid criterion for a good quality solution.

Appendix A

Figures

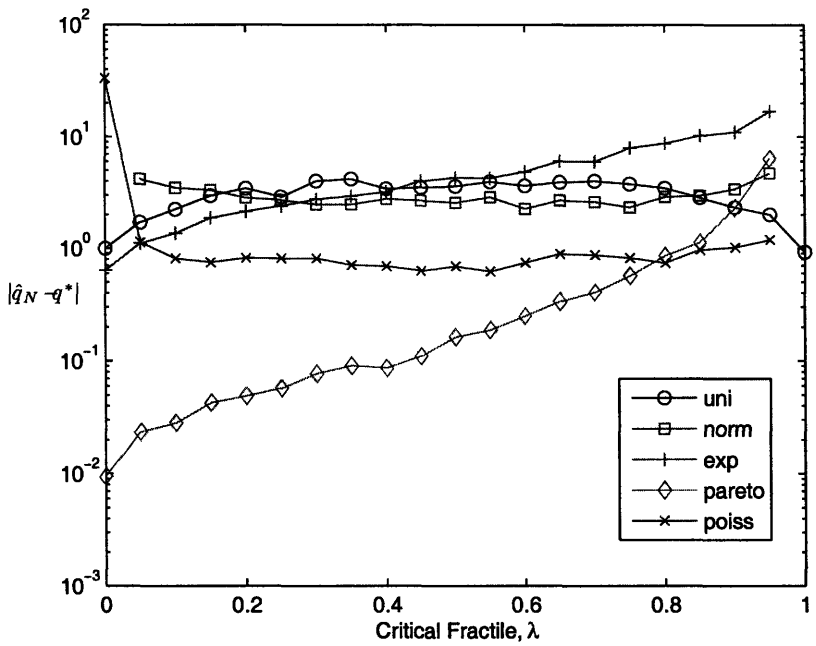

Figure A-1: The sample average of the error of the SAA order quantity as a function of the critical fractile λ.

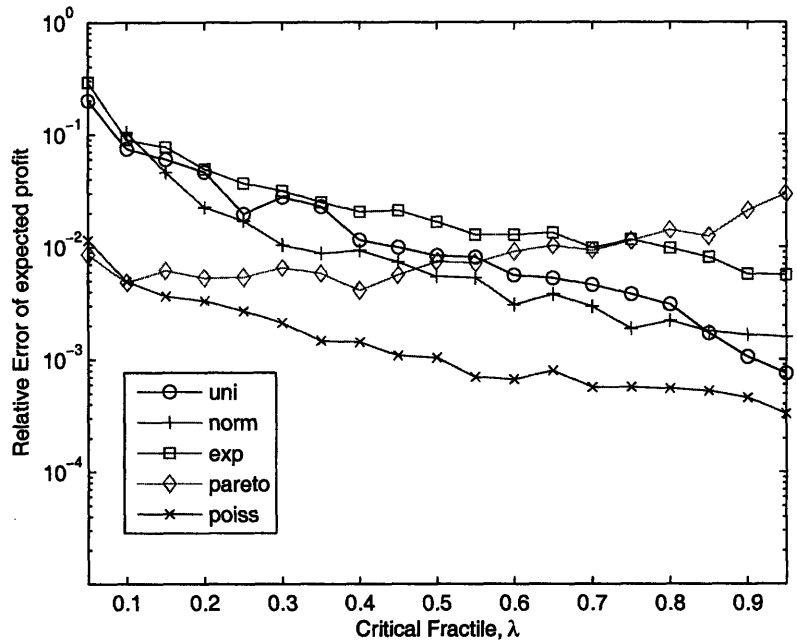

(a) Relative Error of the Average expected profit of the SAA order quantity

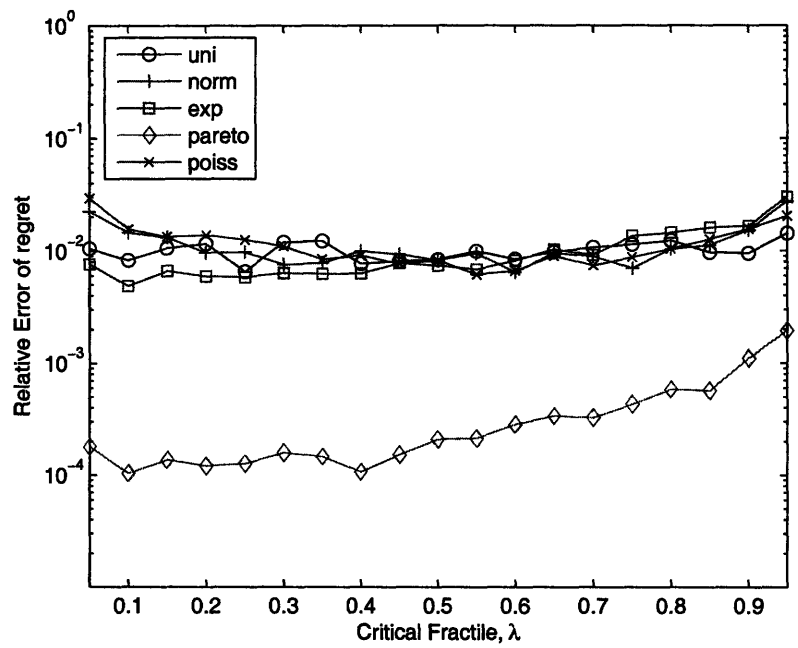

(b) Relative Error of the Average Regret of the SAA order quantity

Figure A-2: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the critical fractile λ for various distributions.

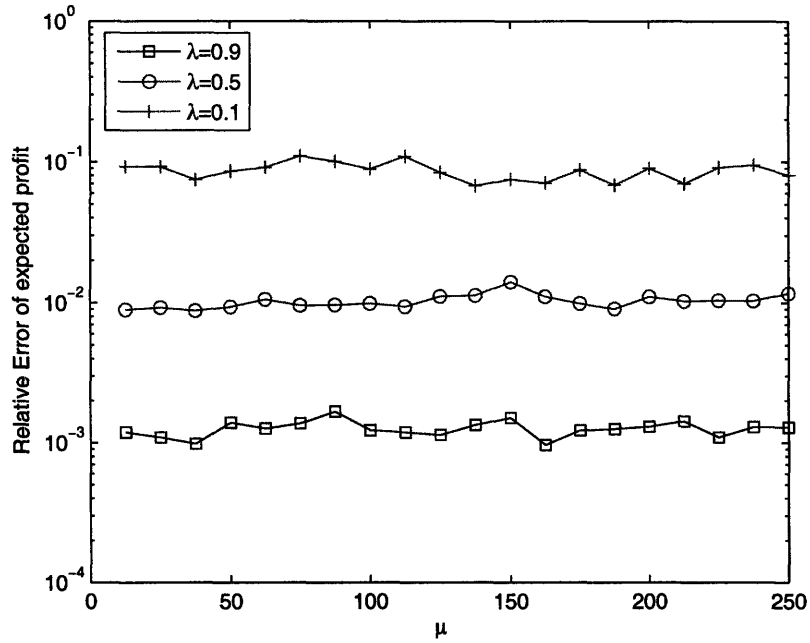

(a) Relative Error of the Average expected profit of the SAA order quantity

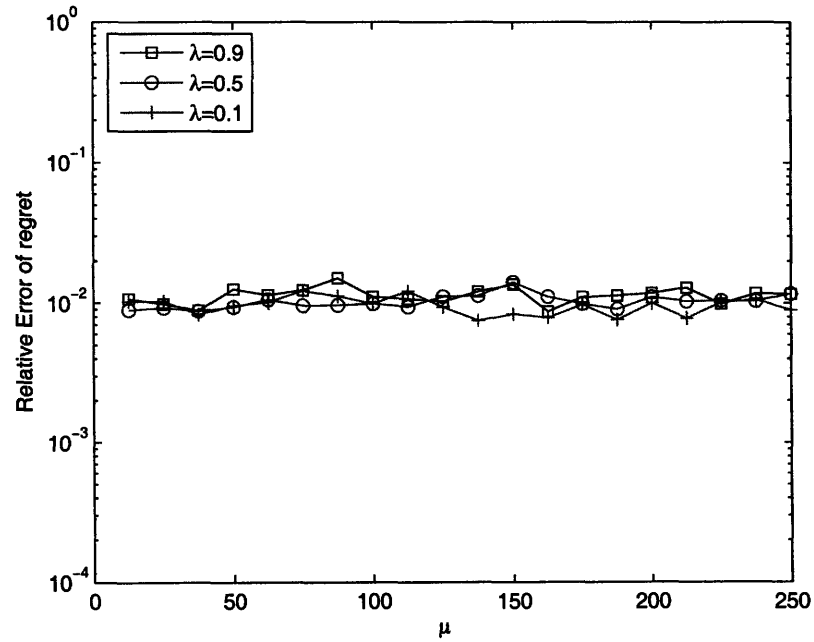

(b) Relative Error of the Average Regret of the SAA order quantity

Figure A-3: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the mean μ of the uniform distribution.

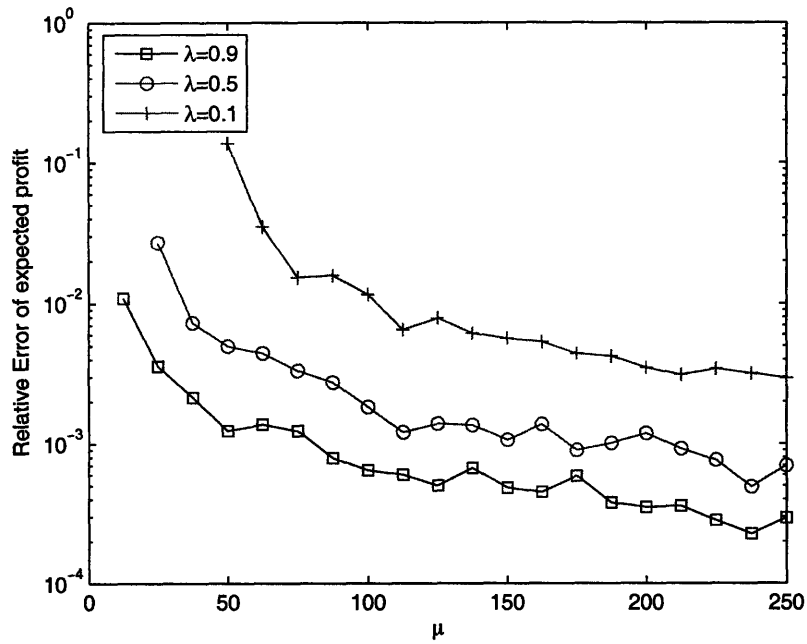

(a) Relative Error of the Average expected profit of the SAA order quantity

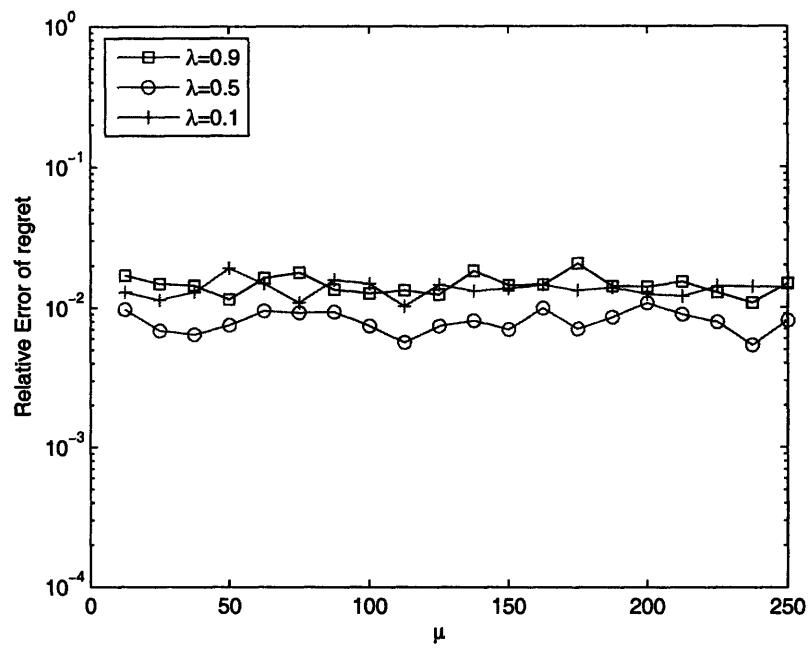

(b) Relative Error of the Average Regret of the SAA order quantity

Figure A-4: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the mean μ of the normal distribution.

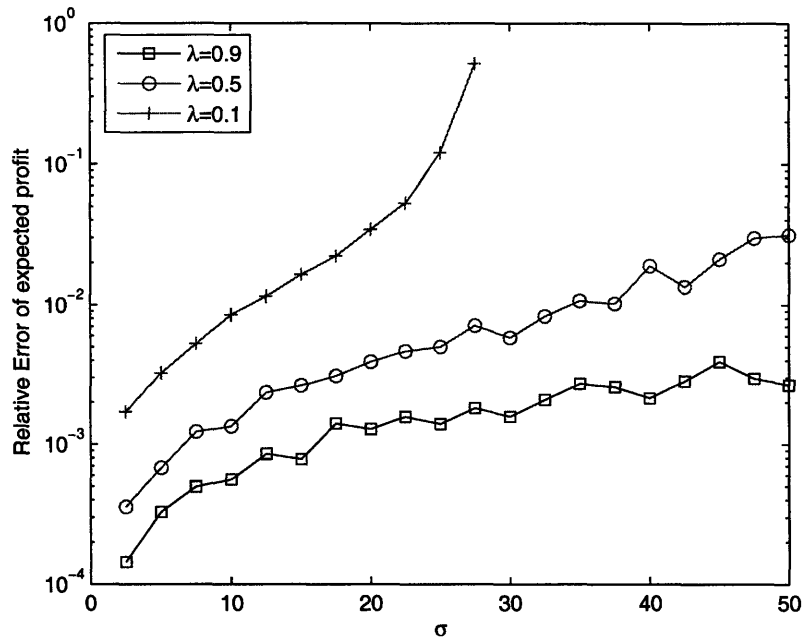

(a) Relative Error of the Average expected profit of the SAA order quantity

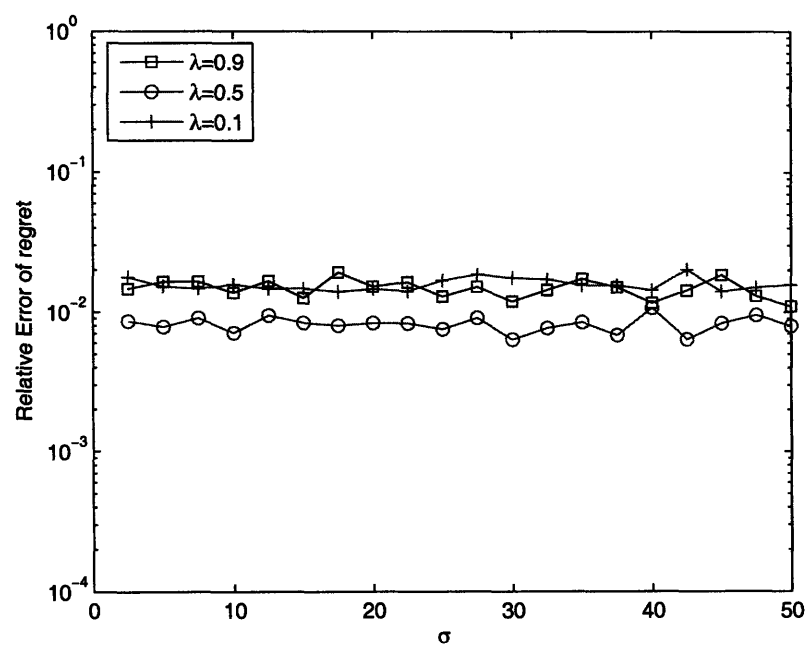

(b) Relative Error of the Average Regret of the SAA order quantity

Figure A-5: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the standard deviation σ of the normal distribution.

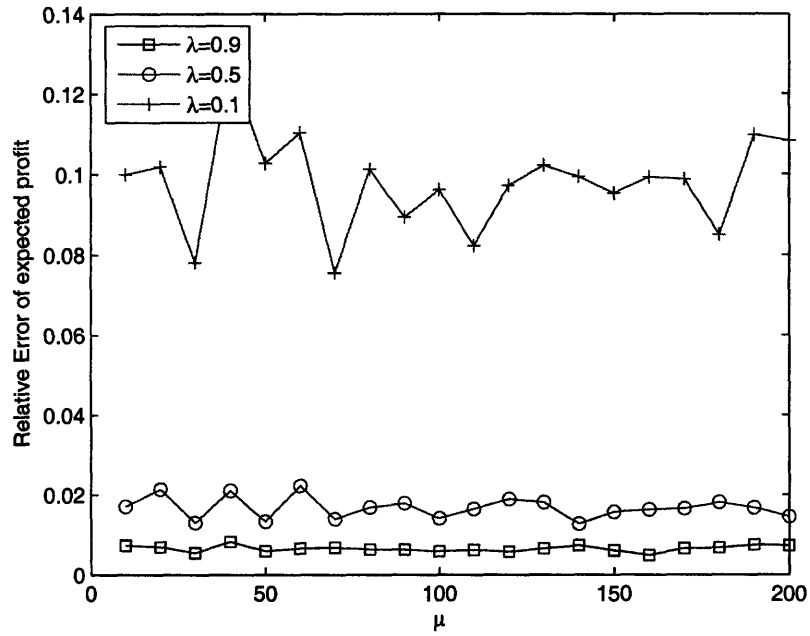

(a) Relative Error of the Average expected profit of the SAA order quantity

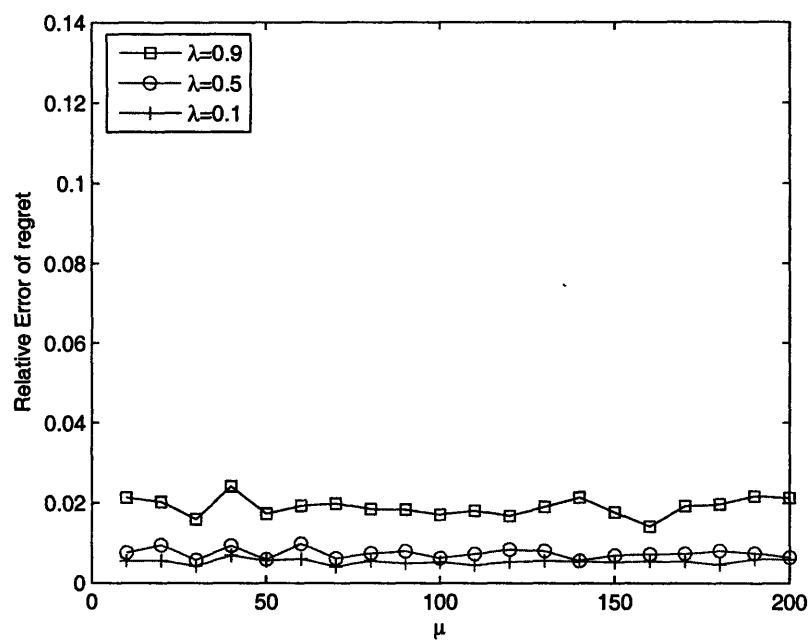

(b) Relative Error of the Average Regret of the SAA order quantity

Figure A-6: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the mean μ of the exponential distribution.

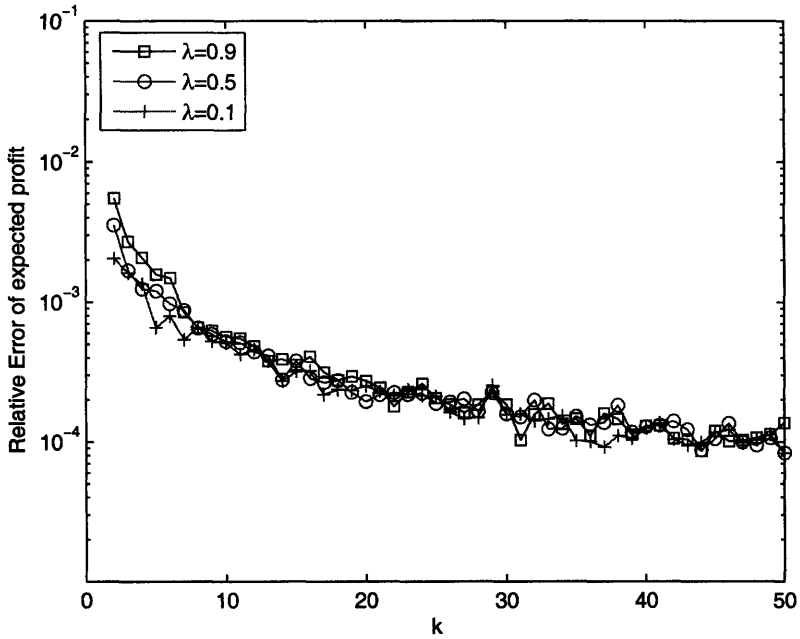

(a) Relative Error of the Average expected profit of the SAA order quantity

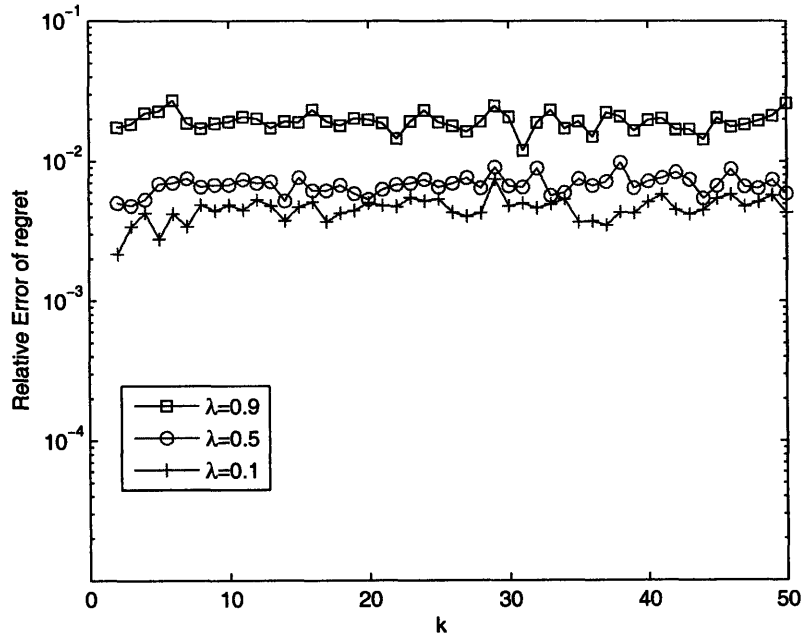

(b) Relative Error of the Average Regret of the SAA order quantity

Figure A-7: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the parameter k of the Pareto distribution.

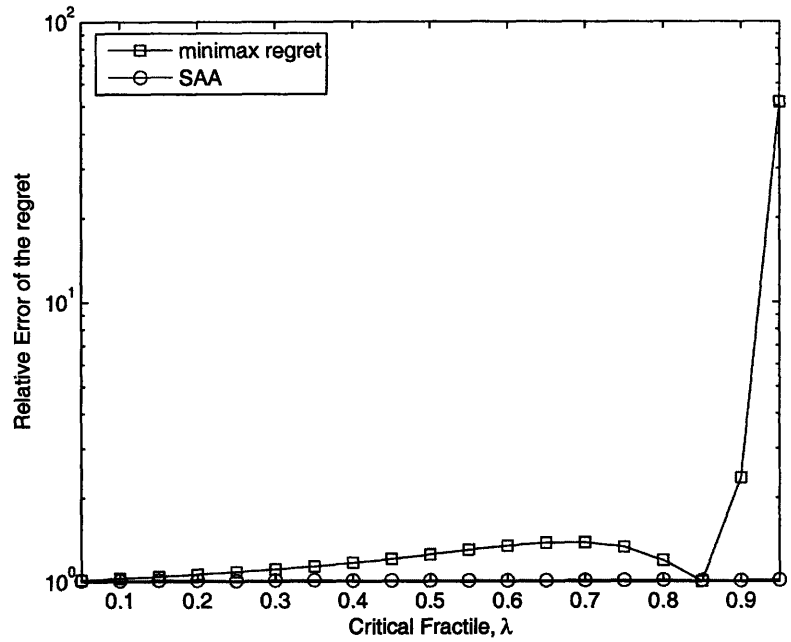

(a) Uniform

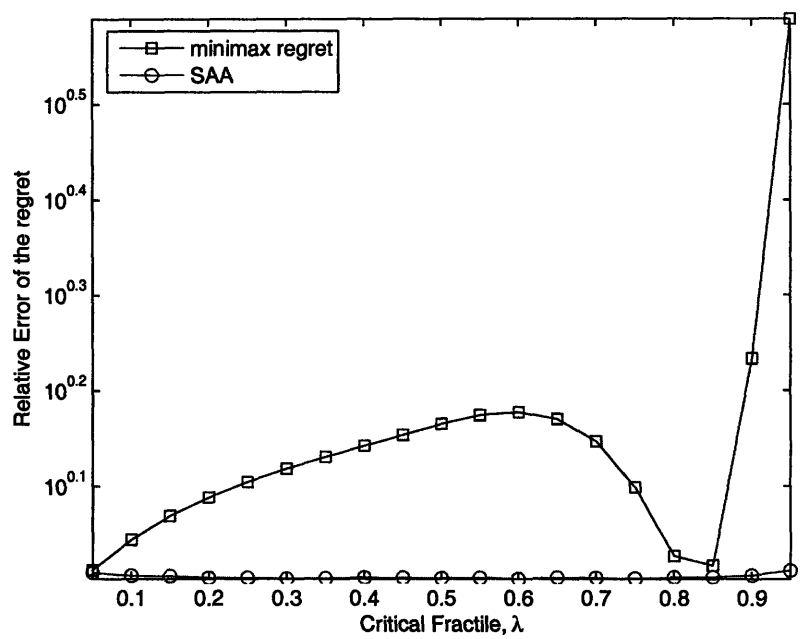

(b) Normal

Figure A-8: The relative error of the regret achieved by the minimax regret order quantity and the SAA order quantity for the uniform and normal distribution.

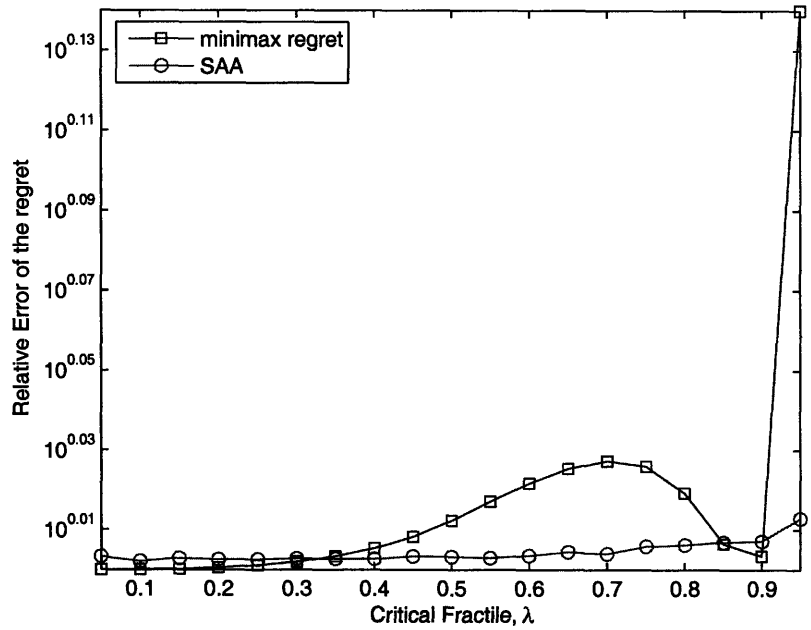

(a) Exponential

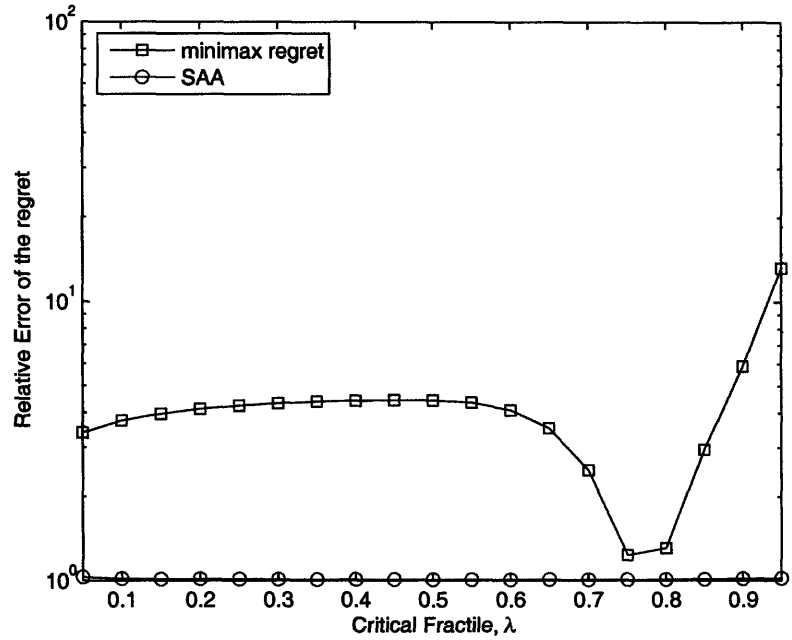

(b) Poisson

Figure A-9: The relative error of the regret achieved by the minimax regret order quantity and the SAA order quantity for the exponential and Poisson distribution.

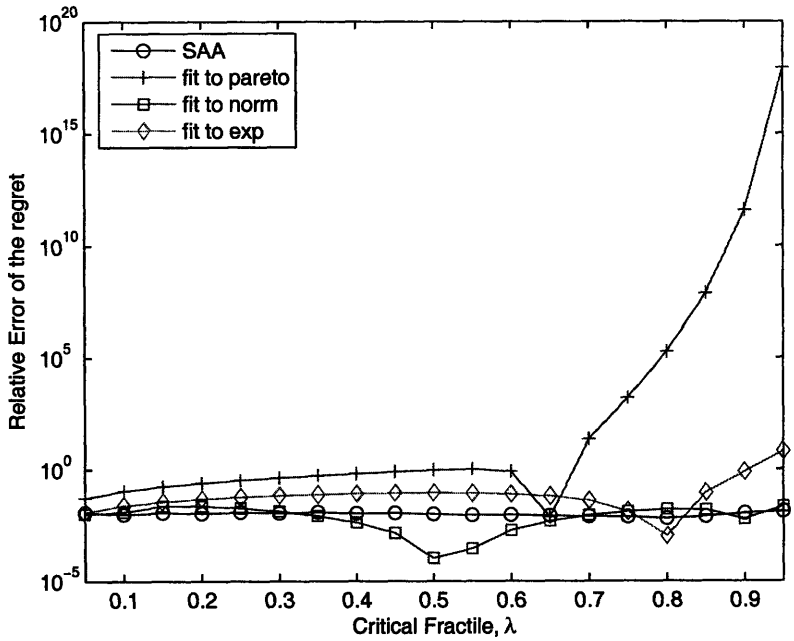

(a) Uniform

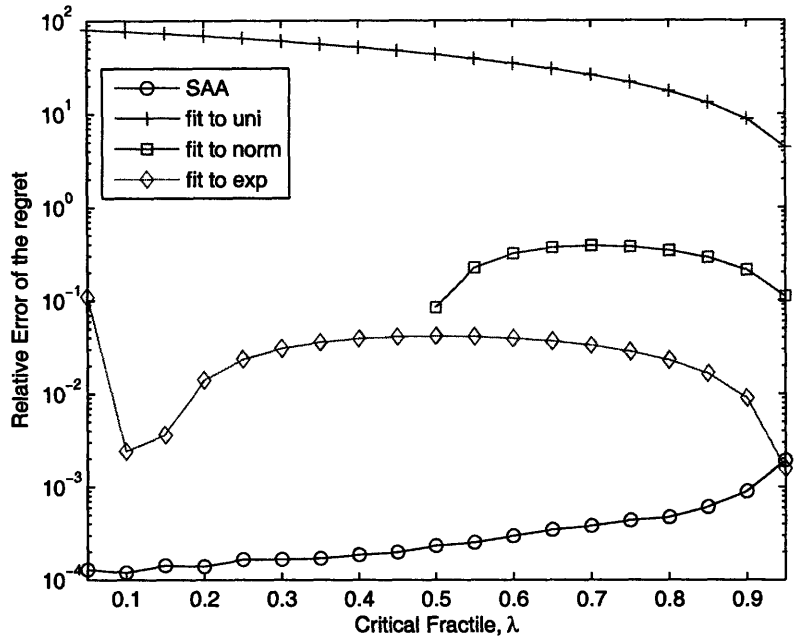

(b) Pareto

Figure A-10: The relative error of the regret achieved by the parameter-estimation order quantity and the SAA order quantity for the uniform and Pareto distribution.

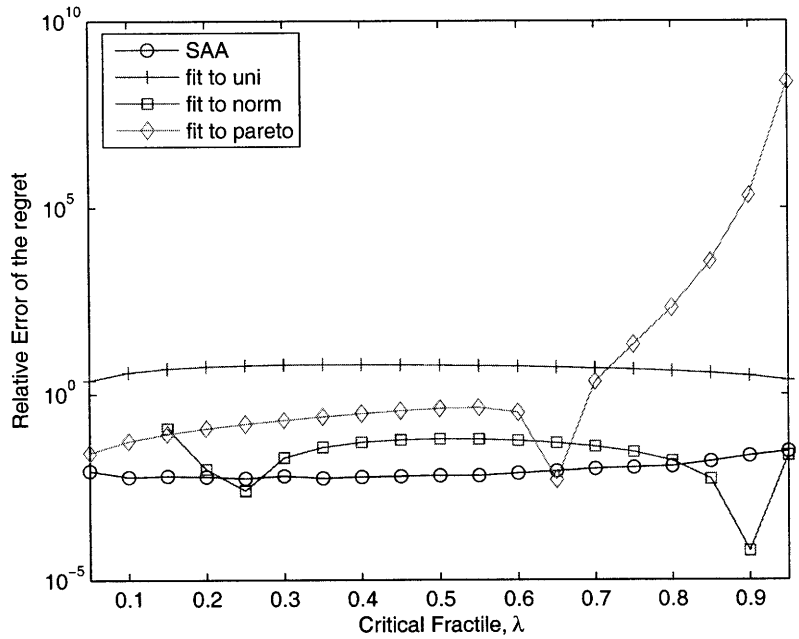

(a) Exponential

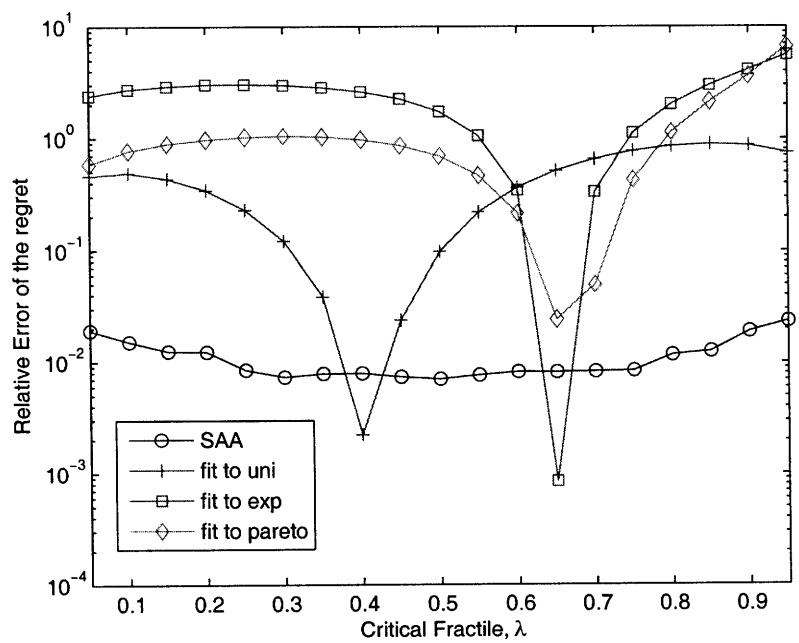

(b) Poisson

Figure A-11: The relative error of the regret achieved by the parameter-estimation order quantity and the SAA order quantity for the exponential and Poisson distribution.

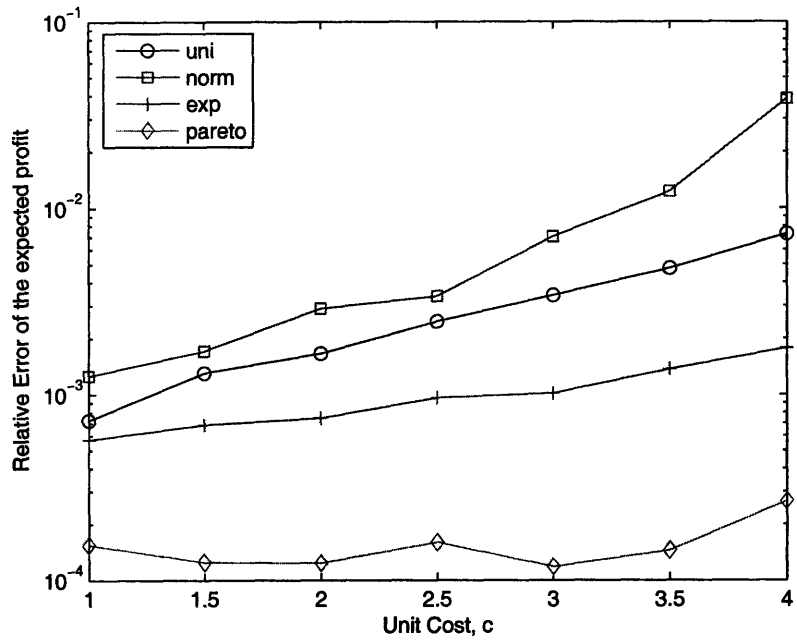

(a) Relative error of the average expected profit

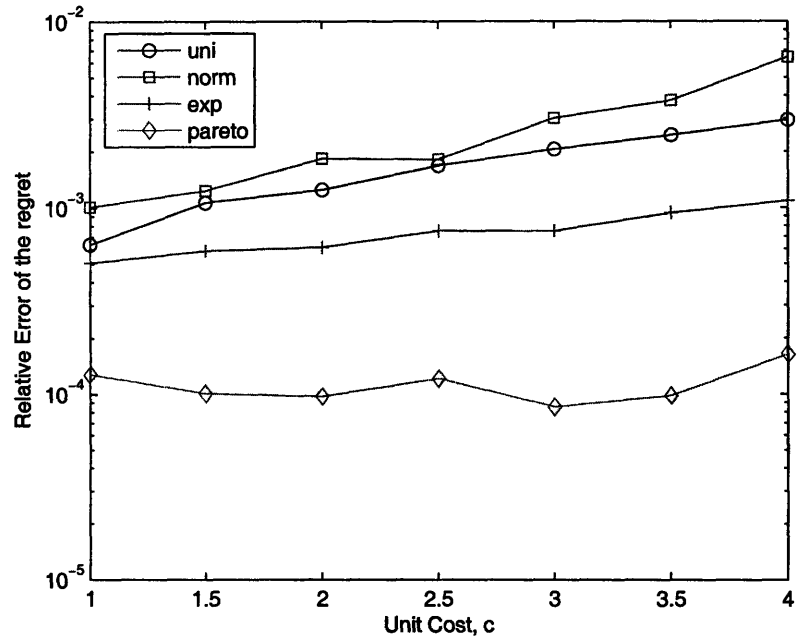

(b) Relative Error of the average regret

Figure A-12: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the unit cost c.

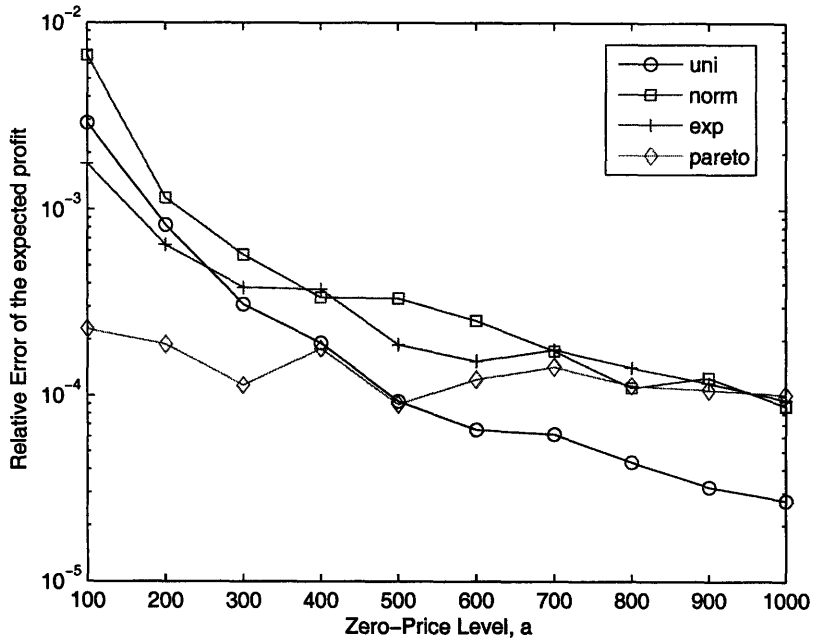

(a) Relative error of the average expected profit

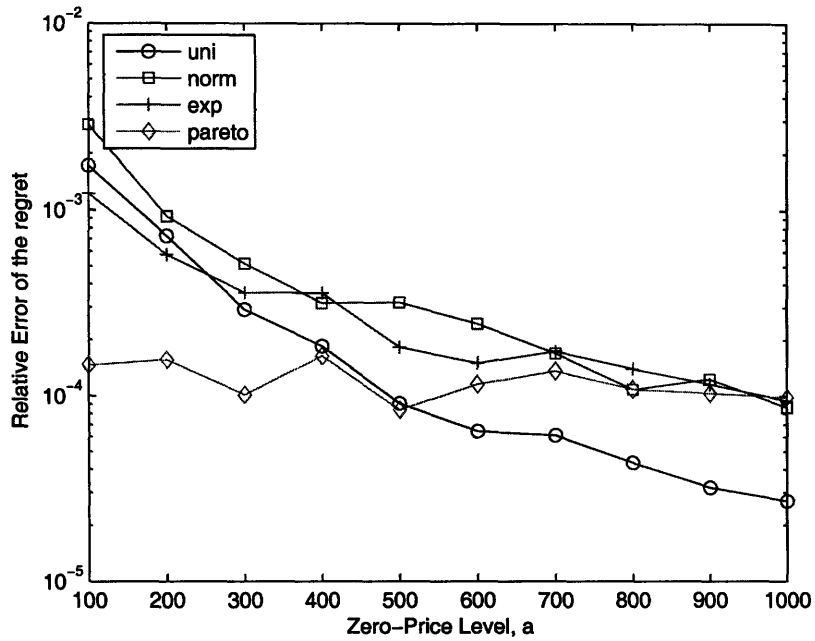

(b) Relative Error of the average regret

Figure A-13: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the zero-price demand a.

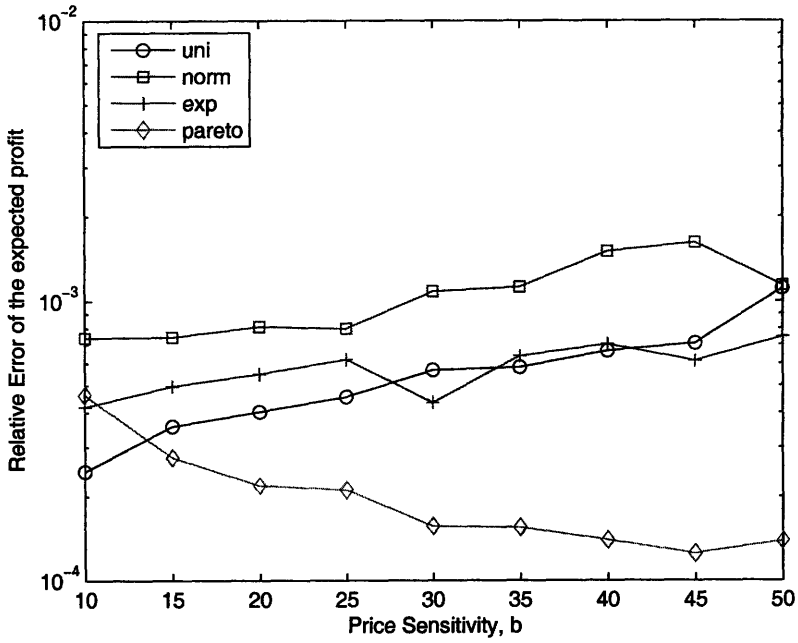

(a) Relative error of the average expected profit

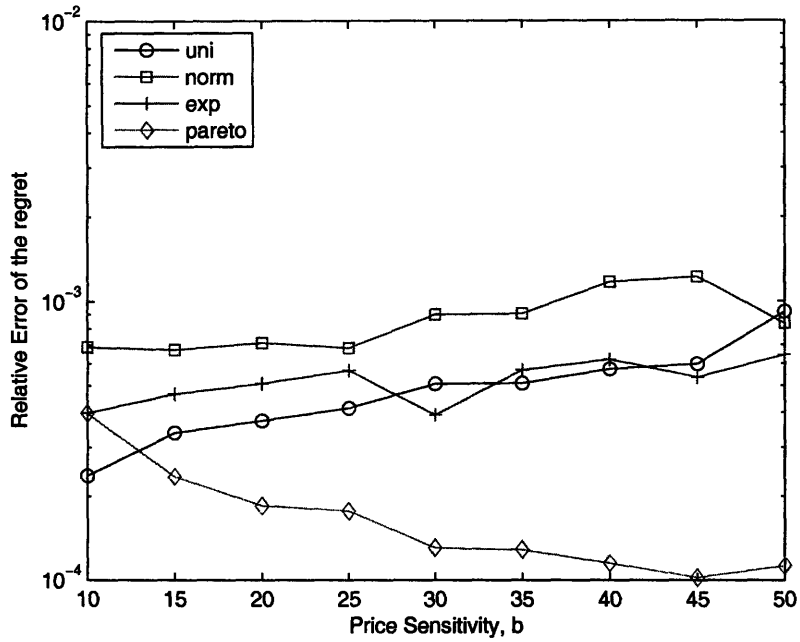

(b) Relative Error of the average regret

Figure A-14: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the price-sensitivity parameter b.

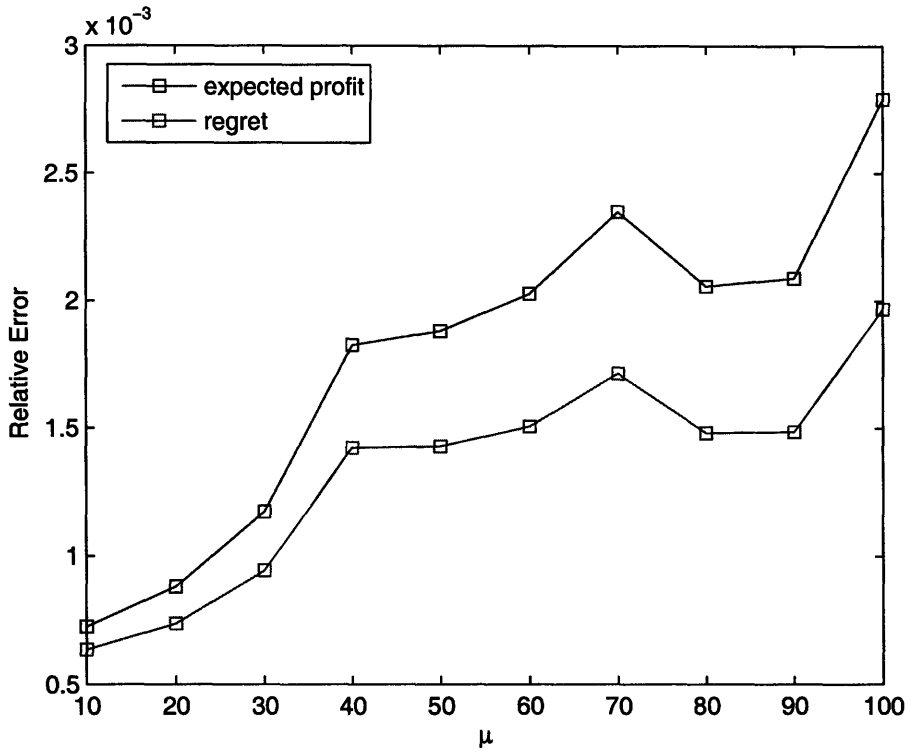

Figure A-15: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the mean μ of the uniform distribution.

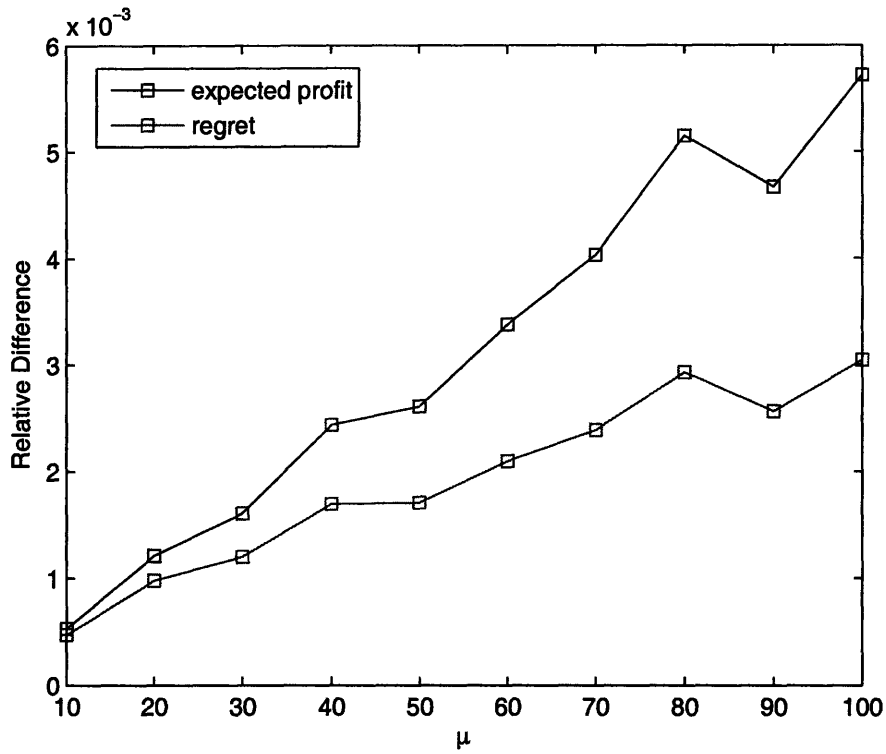

Figure A-16: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the mean μ of the exponential distribution.

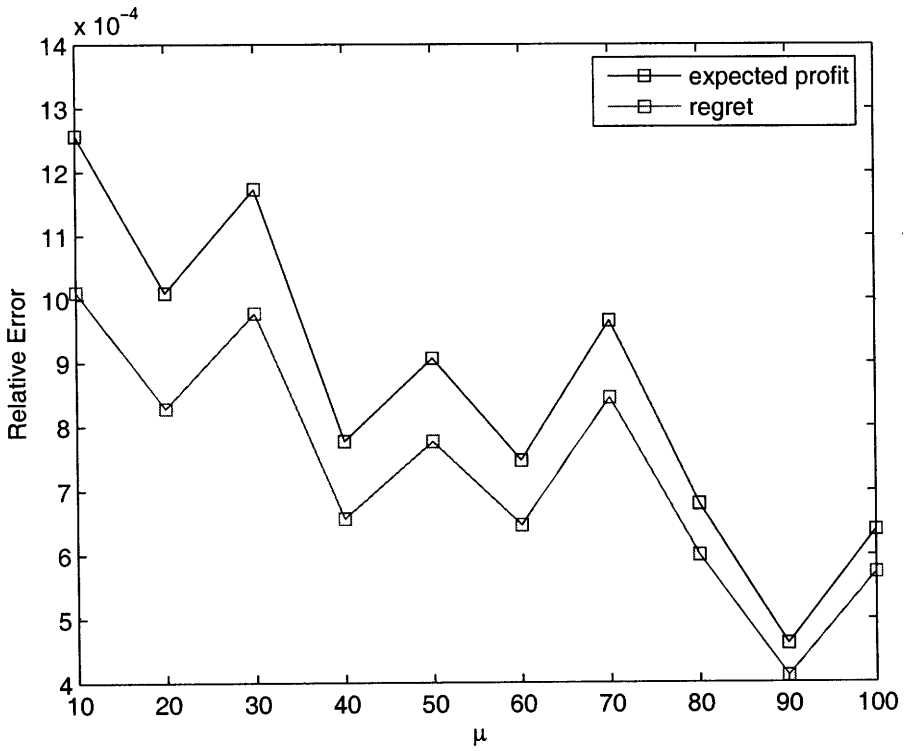

Figure A-17: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the mean μ of the normal distribution.

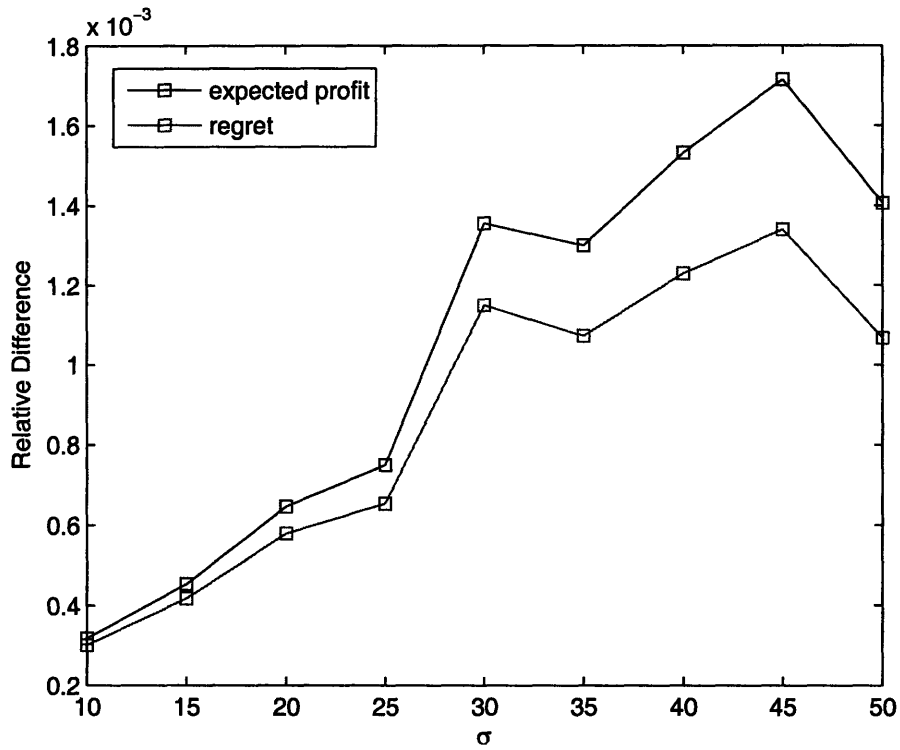

Figure A-18: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the standard deviation σ of the normal distribution.

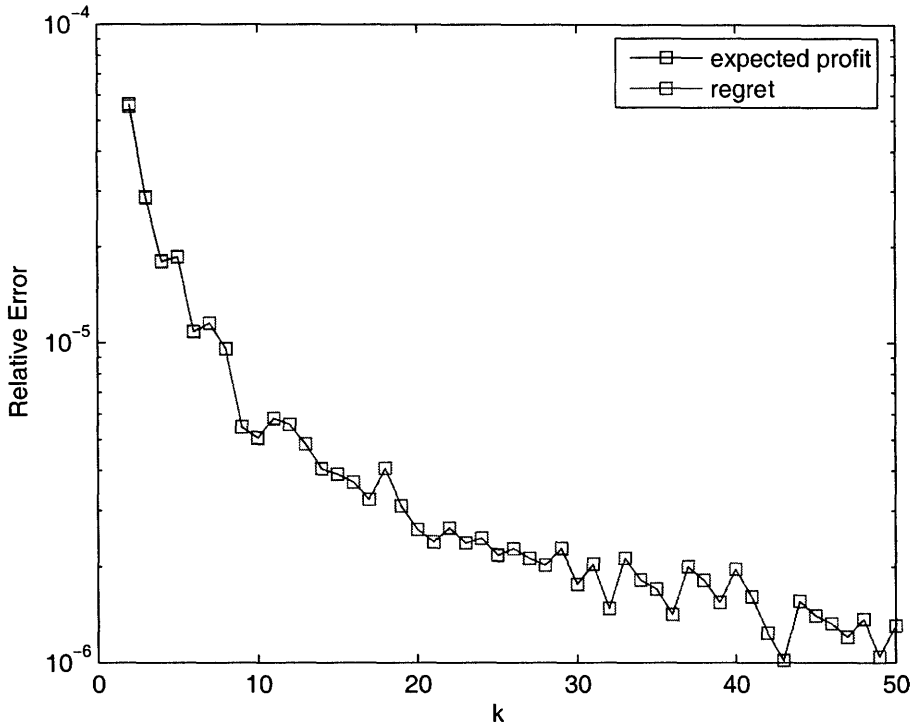

Figure A-19: The accuracy level achieved by the sample average expected profit and the sample average regret of the SAA order quantity as a function of the parameter k of the Pareto distribution.

Bibliography

[1] M. Ball and M. Queyranne. Toward robust revenue management: competitive analysis of online booking. Working paper, available at SSRN: http://ssrn.com/abstract=896547, 2006.

[2] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[3] F. Bernstein and F. Federgruen. Decentralized supply chains with competing retailers under demand uncertainty. *Management Science*, 51(1):18–29, 2005.

[4] S. N. Bernstein. *Theory of Probability.* Moscow, 1927.

[5] D. Bertimas and A. Thiele. A data-driven approach to newsvendor problems. 2005. Working Paper.

[6] D. Bertimas and A. Thiele. A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168, 2006.

[7] D. P. Bertsekas. *Nonlinear Programming.* Athena Scientific, Belmont, MA, second edition, 1999.

[8] D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2003.

[9] J. H. Bookbinder and A. E. Lordahl. Estimation of inventory reorder level using the bootstrap statistical procedure. *IEE Transactions*, 21:302–312, 1989.

[10] G. Cachon and C. Terwiesch. *Matching Supply with Demand: An Introduction to Operations Management.* Irwin - McGraw Hill, 2005.

[11] L. Chan, M. Shen, D. Simchi-Levi, and J. Swann. Coordination of pricing and inventory decisions: A survey and classification. In D. Simchi-Levi, D. Wu, and M. Shen, editors, *Handbook of Quantitative Supply Chain Analysis: Modelling in the E-Business Era.* Kluwer Academic Publishers, 2004.

[12] J. D. Dana and N. C. Petruzzi. Note: The newsvendor model with endogenous demand. *Management Science*, 47(11):1488–1497, 2001.

[13] S. Eren and C. Maglaras. Revenue management heuristics under limited market information: A maximum entropy approach. Presented at The 6th Annual INFORMS Revenue Management Conference, Jun 5-6, Columbia University, 2006.

[14] R. L. Ernst. A linear inventory model of a monopolistic firm. Ph.D. Dissertation, Department of Economics, University of California, Berkeley, CA, 1970.

[15] G. Gallego and I. Moon. The distribution free newsboy problem: Review and extensions. *Journal of the Operational Research Society*, 44(8):825–834, 1993.

[16] P. Glasserman and Y. C. Ho. *Gradient Estimation via perturbation analysis.* Kluwer Academic Publishers, 1991.

[17] G. A. Godfrey and W. B. Powell. An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science*, 47:1101–1112, 2001.

[18] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

[19] W. T. Huh and P. Rusmevichientong. Adaptive capacity allocation with censored demand data: Application of concave umbrella functions. Working paper, 2006.

[20] W. T. Huh and P. Rusmevichientong. A non-parametric asymptotic analysis of inventory planning with censored demand. Working paper, 2006.

[21] S. Karlin and C. R. Carr. Prices and optimal inventory policy. In K. J. Arrow, S. Karlin, and H. Scarf, editors, *Studies in Applied Probability and Management Science*, pages 159–172. Stanford University Press, Stanford, CA, 1962.

[22] M. Khouja. The single-period (news-vendor) problem: literature review and suggestions for future research. *Omega, the International Journal of Management Science*, 27:537–553, 1999.

[23] A. J. Kleywegt, A. Shapiro, and T. Homem-De-Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12:479–502, 2001.

[24] R. Levi, R. Roundy, and D. B. Shmoys. Provably near-optimal sample-based policies for stochastic inventory control models. Cornell Working Paper, 2005.

[25] L. H. Liyanage and J. G. Shanthikumar. A practical inventory control policy using operational statistics. *Operations Research Letters*, 33:341–348, 2005.

[26] E. S. Mills. Uncertainty and price theory. *Quarterly Journal of Economics*, 73:116–130, 1959.

[27] G. E. Monahan, N. C. Petruzzi, and W. Zhao. The dynamic pricing problem from a newsvendor's perspective. *Manufacturing and Service Operations Management*, 6:73–91, 2004.

[28] G. Perakis and G. Roels. Regret in the newsvendor model with partial information. To appear in Operations Research, 2006.

[29] G. Perakis and G. Roels. Robust controls for network revenue management. Working Paper, 2006.

[30] N. C. Petruzzi and M. Dada. Pricing and the newsvendor problem: A review with extensions. *Operations Research*, 47(2):183–194, 1999.

[31] E. L. Porteus. Stochastic inventory theory. In D. P. Heyman and M. J. Sobel, editors, *Handbooks in OR & MS*, volume 2, pages 605–652. Elsevier, North Holland, 1990.

[32] G. Raz and E. L. Porteus. A discrete service levels perspective to the newsvendor model with simultaneous pricing. Stanford University Working Paper, 2001.

[33] G. Raz and E. L. Porteus. A fractiles perspective to the joint price/quantity newsvendor model. *Management Science*, 52(11):1764–1777, 2006.

[34] L. J. Savage. The theory of statistical decisions. *Journal of the American Statistical Association*, 46(253):55–67, 1951.

[35] H. Scarf. Bayes solution to the statistical inventory problem. *Annals of Mathematical Statistics*, 30(2):490–508, 1959.

[36] H. E. Scarf. A min-max solution to an inventory problem. In K. J. Arrow and S. Karlin and H. E. Scarf, editor, *Studies in Mathematical Theory of Inventory and Production*, pages 201–209. Stanford Univerity Press, Stanford, CA, 1958.

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

[38] A. Shapiro. Monte Carlo sampling methods. In A. Ruszczyński and A. Shapiro, editor, *Stochastic Programming*, volume 10 of *Handbooks in operations research and management science*, chapter 6. Elsevier, 2003.

[39] G. T. Thowsen. A dynamic, nonstationary inventory problem for a price/quantity setting firm. *Naval Research Logistics Quarterly*, 22:461–476, 1975.

[40] T. M. Whitin. Inventory control and price theory. *Management Science*, 2:61–68, 1955.

[41] E. Zabel. Monopoly and uncertainty. *Review of Economic Studies*, 37:205–219, 1970.

[42] R. L. Zhan and Z. J. M. Shen. Newsvendor problem with pricing: Properties, algorithms, and simulation. In *Winter Simulation Conference*, pages 1743–1748, 2005.