

Robust Stochastic Lot-Sizing by Means of Histograms

Diego Klabjan

Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL.

E-mail: d-klabjan@northwestern.edu Phone: +1 847 491 0663 Fax: +1 847 491 8005

David Simchi-Levi

Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge,

MA. E-mail: dslevi@mit.edu Phone: +1 617 253 6160 Fax: +1 617 324 2546

Miao Song

Department of Industrial and Manufacturing Systems Engineering, The University of Hong Kong, Hong

Kong SAR, P.R. China. E-mail: msong@hku.hk Phone: +852 2859 7966 Fax: +852 2858 6535

Abstract

Traditional approaches in inventory control first estimate the demand distribution among a predefined family of distributions based on data fitting of historical demand observations, and then optimize the inventory control using the estimated distributions. These approaches often lead to fragile solutions whenever the preselected family of distributions was inadequate. In this work we propose a minimax robust model that integrates data fitting and inventory optimization for the single-item multi-period periodic review stochastic lot-sizing problem. In contrast with the standard assumption of given distributions, we assume that histograms are part of the input. The robust model generalizes the Bayesian model, and it can be interpreted as minimizing history dependent risk measures. We prove that the optimal inventory control policies of the robust model share the same structure as the traditional stochastic dynamic programming counterpart. In particular, we analyze the robust model based on the chi-square goodness-of-fit test. If demand samples are obtained from a known distribution, the robust model converges to the stochastic model with true distribution under generous conditions. Its effectiveness is also validated by numerical experiments.

Keywords: stochastic inventory control, robust optimization, dynamic programming, optimal policy

History: received June 2010; revisions received December 2011, May 2012; accepted June 2012

1. Introduction

The stochastic lot-sizing model has been extensively studied in the inventory literature. Most of the research has focused on models with complete information about the distribution of customer demand. However, in most real-world situations, the demand distribution is not known; only historical data is available. A common approach is to hypothesize a family of demand distributions and then to estimate the parameters specifying the distribution using the historical data. Once the probability distribution has been identified, the inventory problem is solved following this estimated

distribution. This implies that the inventory policy is determined under the assumption that the fitted distribution adequately characterizes the demand to be realized in the future.

The estimated demand distribution may not be accurate and hence the approach of fitting the distribution and optimizing the inventory decisions sequentially may not work as expected. As shown in Liyanage and Shanthikumar (2005) for the newsvendor model, such an approach may generate suboptimal solutions. Besides, in distribution fitting, one needs to assume a parametric family of a demand distribution in the first place, and this hypothesis may also go awry. For instance, we may fit the historical data to a lognormal distribution while it actually follows a uniform distribution.

The robust inventory models, without assuming a parametric family of distributions, provide an approach to address ambiguity in the demand distribution. A brief review of these robust models is provided in Section 1.1. These models adopt a minimax approach targeting to minimize the worst case expected cost maximized over the set of distributions. Without exception, the existing literature either considers a pre-specified set for demand distributions without detailed discussions about how to generate the set, e.g., Notzon (1970), or derives the set of distributions based on certain statistics of the historical data such as the sample mean and variance, e.g., Bertsimas and Thiele (2006) and See and Sim (2010). Comparing with the classical approach with separate fitting and optimization, the robust models based on historical statistics may miss important information about the demand distribution conveyed in the historical data set, e.g., the shape of the distribution, which, in the separate, two-phase approach, is usually used to determine the parametric family of the distributions.

In this paper, we merge the merits of both approaches, namely, (i) to fully utilize historical data as in the classical approach and (ii) to concurrently optimize the demand distribution and the inventory decision without assuming a distribution family as in the robust models. We analyze the single-item stochastic finite-horizon periodic review lot-sizing model, under the assumption that the demand is subject to an unknown distribution and only historical demand observations (given by histograms) are available. As all practitioners in inventory control start with histograms and then fit an underlying demand distribution, this assumption reflects the practical value of this research.

By adopting the minimax robust optimization approach, rather than first estimating the demand distribution and then optimizing inventory decisions, we combine these two steps to minimize the worst case expected cost over a set of demand distributions, which is defined as all possible distributions satisfying the chi-square goodness-of-fit test. The advantage of this approach is twofold. First, the historical data is used in the same manner as in the goodness-of-fit test, thus we use all the information conveyed by the historical data that can be utilized by the goodness-of-fit test in distribution fitting. Second, we avoid the assumption about the parametric family of distributions, which is a must in distribution fitting. We show that the (s, S) policy remains optimal, discuss the behavior of the model as the number of samples increases, and demonstrate through a numerical study that this model outperforms (i) the classical approach where distribution estimation and inventory optimization are separate and (ii) a robust model where the set of distributions is defined by sample mean and variance.

Our two main contributions are as follows. First, we develop a robust minimax model that only

requires historical data, and allows correlated demand. Note that most minimax models (see, e.g., Notzon 1970 and Ahmed et al. 2007) as well as Bayesian inventory models (e.g., updating demand distributions as suggested in Iglehart (1964) in the literature can be interpreted as special cases of our framework.

Unlike the classical inventory model, which solves a single-variate optimization problem in each period, the robust model needs to identify the ordering quantity and probability distribution represented by a vector of decision variables simultaneously. Despite this complexity, the optimal policy of the robust model still shares the same structure as the corresponding policy in the classical stochastic lot-sizing model. In particular, the optimal policy is a state-dependent base-stock policy for the multi-period inventory problem without fixed procurement costs, and a state-dependent (s, S) policy if the fixed procurement cost is considered.

While the first contribution mainly serves as an extension to existing models, the second major contribution regards combining the statistical test in distribution fitting within a single inventory control model. We consider a special case of the general robust framework when the set of demand distributions is directly related to the chi-square goodness-of-fit test. Such a distribution set can be defined by a set of second-order cone constraints and hence it is tractable to compute the (s, S) levels for each period. To the best of our knowledge, this is the first endeavor to integrate the goodness-of-fit statistical test with inventory optimization and to explicitly consider the shape of the distribution in a robust framework.

We also prove that the robust model based on the chi-square test converges to the stochastic model with true demand distribution under generous conditions if samples are drawn from this distribution and they grow indefinitely. In particular, if the demand distributions are discrete, the robust model converges to the stochastic model with the true demand distribution as the number of independent samples drawn from the true distribution for each period tends to infinity. Moreover, the rate of convergence is in the order of $1/\sqrt{k}$, where k is the number of samples. Slightly weaker results are obtained for continuous distributions. These convergence results ensure the effectiveness of the robust approach when the sample size is sufficiently large.

When the sample size is relatively small, the performance of the robust model is illustrated by means of computational experiments. We argue that the robust model based on the chi-square test outperforms the traditional approach, which optimizes the inventory decisions by using fitted distributions, as well as the minimax robust model where the set of distributions is based on the set proposed by Delage and Ye (2010). We also provide insights on the performance of the robust model with different parameters and sample sizes.

In Section 2 we describe our robust model, which incorporates historical data, and present the optimality equation in a compact form. The structure of the optimal policies is characterized in Section 3. Section 4 considers a special case with robustness defined by the chi-square goodness-of-fit test. We also discuss selected convergence results for the chi-square test based models in the same section. The computational results are presented in Section 5. Finally, additional extensions are presented in Section 6. We conclude the introduction with the literature review.

1.1 Literature Review

This work is built upon two streams of literature: stochastic inventory control and robust optimization. The discrete-time stochastic inventory model has been studied since the 1950s. Scarf (1960) proposes the concept of K -convexity and proves that the (s, S) policy is optimal in the presence of a fixed ordering cost. Since then, the research in this area has flourished. We refer the reader to Zipkin (2000) for a detailed review. The concept of K -convexity has been generalized to attack various problems related to inventory control, e.g., Chen and Simchi-Levi (2004). Efficient algorithms, e.g., Guan and Miller (2008) and Halman et al. (2009), have also been proposed to solve other more general stochastic inventory problems.

Robust optimization was pioneered by Soyster (1973), which proposes robust linear programming formulations for linear programs with coefficient uncertainty. This line of research has enjoyed popularity in recent years. Some of the important works include but are not limited to Ben-Tal and Nemirovski (2000) and Bertsimas and Sim (2004) for robust linear programming, Ben-Tal and Nemirovski (1998) for robust convex optimization, and Kouvelis and Yu (1997) and Bertsimas and Sim (2003) for robust discrete optimization. More relevant to this research, Iyengar (2005) and Nilim and El Ghaoui (2005) develop a robust optimization framework for dynamic programming models, extend the Bellman recursion to the robust counterpart, and investigate its computational complexity. Delage and Ye (2010) propose a data-driven robust framework for any single-stage optimization problem, which minimizes the maximum expectation over a set of distributions defined by the sample mean and variance. They identify sufficient conditions under which the corresponding robust problem is polynomially solvable and provide probabilistic arguments for using this model by considering a confidence region for the mean and variance as a random vector.

In this paper, we apply the idea of robust optimization to inventory control models. This notion of robust inventory control is not new in the literature. The earliest work in minimax inventory control is attributed to Scarf (1958), where minimization of the maximum expected cost of the newsvendor model over all distributions with a given mean and variance is considered. Gallego and Moon (1993) present another proof of Scarf's result and consider various extensions of the model. The recent work by Natarajan et al. (2008) extends the result of Scarf (1958) by considering the set of distributions with a given mean, variance and semivariance information. Perakis and Roels (2008) minimize the maximum regret of the newsvendor model over a convex set of distributions with certain moments and shape.

Notzon (1970) is among the earliest works that considers a minimax multiple-period inventory model. The demand in each period is assumed to be independent and its distribution function is ambiguous but within a specified class of distribution functions. The minimax control policy minimizes the maximum expected cost. The optimality of the (s, S) policy is proved.

Bertsimas and Thiele (2006) analyze distribution-free inventory problems, in which demand in each period is assumed to be a random variable that takes values in a given range. The demand is assumed to be a random variable controlled only by two values: the lower and upper estimators. To capture the trade-off between robustness and optimality, a parameter is defined to control the budgets of uncertainty at every time period. They show that for a variety of problems, the structures

of the optimal policy remain the same as in the associated model with complete information about the distribution of customer demand. A related model from the base-stock perspective is analyzed in Bienstock and Özbay (2008).

See and Sim (2010) consider a factor-based demand model with given mean, support, and deviation measures. To obtain tractable replenishment policies, the worst case expected cost among all distributions satisfying the demand model is minimized by solving a second order cone optimization problem.

Ahmed et al. (2007) propose an inventory control model which minimizes a coherent risk measure instead of the overall cost function. They show that risk aversion treated in the form of coherence risk measures is equivalent to the minimax formulations, and it is proved that the optimal policies conserve the properties of the stochastic dynamic programming counterparts. They do not consider demand dependent evolutions.

Liyanage and Shanthikumar (2005) first provide concrete examples in a single period (news vendor) setting, which illustrate that separating distribution estimation and inventory optimization, as done in the classical approach, may lead to suboptimal solutions. They propose the use of operational statistics where it is assumed that the demand distribution function belongs to a specific (predetermined) family and estimate the (single) parameter of the family within an inventory optimization model.

In addition, selected recent papers also consider lost-sale inventory problems with censored demand data, i.e., the observed historical demand data excludes the lost-sale information as the lost sales are not observable. Huh and Rusmevichientong (2009) propose nonparametric adaptive policies to solve this problem and provide a bound for the asymptotic performance, which interestingly is the same as the convergence rate of our model under discrete distributions.

The models by Notzon (1970) and Ahmed et al. (2007) do not take the historical data into account, and they predefine the class of distribution functions. The robust optimization approaches from Bertsimas and Thiele (2006) as well as See and Sim (2010) do not use any historical data except to determine the support, expectation and deviation measures. On the other hand, Liyanage and Shanthikumar (2005) use historical data but predetermine the family of distributions. In fact, they consider only distributions characterized by a single unknown parameter. This is the only work besides the one proposed in this paper that concurrently optimizes the ordering quantity and applies techniques in distribution fitting to determine the demand distribution. Our research combines both strategies by integrating distribution fitting with robust optimization. Specifically, we consider the set of demand distributions that satisfy a certain data fitting criterion with respect to historical data and characterize an optimal policy that minimizes the maximum expected cost.

2. Formulation of Robust Stochastic Lot-Sizing

The classical multi-period inventory problem considers a finite planning horizon of T periods. We assume that all shortages are backlogged. For each period $t = 1, \dots, T$, let \tilde{D}_t be a random variable representing demand in that period. The sequence of events is as follows.

At the beginning of each period t , the decision maker reviews the net inventory level x_t , and places an order for q_t (possibly zero) units. The procurement cost in each period $t = 1, \dots, T - 1$ includes two components: a fixed procurement cost K if $q_t > 0$, and a unit procurement cost c_t for each unit ordered.

Assuming zero lead time, this order arrives immediately and increases the inventory level up to y_t , where $y_t = x_t + q_t$. After observing demand \tilde{D}_t , inventory holding cost is charged at a rate of h_t for any unit of excess inventory after satisfying demand \tilde{D}_t , and a unit backorder cost b_t is incurred for any unit of unsatisfied demand. The net inventory at the beginning of period $t + 1$ is reduced to $x_{t+1} = y_t - \tilde{D}_t$.

Thus, the total cost for period t given the net inventory levels before and after ordering (x_t and y_t respectively) as well as demand \tilde{D}_t in that period is

$$C_t(x_t, y_t, \tilde{D}_t) = K\mathbb{I}(y_t - x_t) + c_t(y_t - x_t) + h_t(y_t - \tilde{D}_t)^+ + b_t(y_t - \tilde{D}_t)^- \quad t = 1, \dots, T, \quad (1)$$

where $x^+ = \max(x, 0)$, $x^- = \max(-x, 0)$, $\mathbb{I}(x) = 1$ if $x > 0$ and $\mathbb{I}(x) = 0$ otherwise.

In the standard dynamic programming formulation, we consider $\tilde{V}_t(x_t)$, $t = 1, \dots, T$, which denotes the optimal expected cost over horizon $[t, T]$, given that the net inventory level at the beginning of period t is x_t and an optimal policy is adopted over horizon $[t, T]$. We assume $\tilde{V}_{T+1}(x_{T+1}) = 0$. Let $\theta \in [0, 1]$ be the discount rate. The optimality equation reads

$$\tilde{V}_t(x_t) = \min_{y_t \geq x_t} \left\{ E \left[C_t(x_t, y_t, \tilde{D}_t) \right] + \theta E \left[\tilde{V}_{t+1}(y_t - \tilde{D}_t) \right] \right\} \quad t = 1, \dots, T. \quad (2)$$

Note that the distribution of \tilde{D}_t , $t = 1, \dots, T$ is required to solve this dynamic programming formulation.

In practice, the demand distribution is not known. Rather, an inventory manager has at her disposal only historical data. Depending on the realized past demand in the planning horizon, the manager may choose different aggregations of historical data to forecast the demand distribution. For example,

- the demand data of the last n observations are considered, which is analogous to the moving average forecast, or
- the realized demand in periods 1 to $t - 1$ is accounted for when forecasting the demand in period t .

Historical observations are often aggregated to a histogram with respect to unknown distribution \tilde{D}_t . The bins are $[D_{t,i}, D_{t,i+1})$, which denotes the i th possible range of the demand in period t (all observations within a given range are indistinguishable). Let the vector $\mathbf{d}_t = [d_1, \dots, d_{t-1}]$ denote the realized demand in periods 1 to $t - 1$, where d_τ , $\tau = 1, \dots, t - 1$ corresponds to the realized demand in period τ . The number of observations falling within the i th bin is a function of the realized demand \mathbf{d}_t and is denoted by $N_{t,i}(\mathbf{d}_t)$. Finally, we define $n_t(\mathbf{d}_t) = \sum_i N_{t,i}(\mathbf{d}_t)$, which corresponds to the total number of available observations under realized demand \mathbf{d}_t . In practice, the decision maker observes only these histograms, i.e., the historical samples.

We assume that $D_{t,1} = 0$ and $D_{t,M_t+1} = +\infty$, where M_t corresponds to the number of bins in the histogram for time period t . Let $P_{t,i} = P\left(\tilde{D}_t \in [D_{t,i}, D_{t,i+1})\right)$ be the probability that demand in period t falls in the interval $[D_{t,i}, D_{t,i+1})$ under the fitted distribution. Clearly, $n_t(\mathbf{d}_t)P_{t,i}$ is the expected number of observations that fall in this interval according to the fitted distribution.

The classical approach to identify the best distribution representing the observed data is to use a goodness-of-fit test. The objective is to fit a distribution that ‘‘closely’’ follows the observed data. Under this criterion, there should be a set of distributions depending on \mathbf{d}_t , which satisfy the given goodness-of-fit test. We denote this set by $\mathcal{P}_t(\mathbf{d}_t)$. Throughout this paper, we assume that $\mathcal{P}_t(\mathbf{d}_t)$ is compact for any t and \mathbf{d}_t .

As defined in the dynamic programming field, a decision rule μ_t at time t is a function of net inventory x_t , which decides the ordering quantity at time t given x_t , i.e., $y_t = \mu_t(x_t)$. We formally state our problem in the context of a two-player game, which is also presented in Iyengar (2005). The first player chooses the decision rule μ_t at time t and pays the cost. The second player chooses a distribution of \tilde{D}_t in $\mathcal{P}_t(\mathbf{d}_t)$ after observing the order quantity, and receives a reward equal to the cost paid by the first player. Therefore, the second player may select a different distribution for different x_t and μ_t . Let $\mathbf{P}_t(x_t, \mu_t(x_t))$ denote the distribution chosen by player 2 at time t given net inventory x_t and decision rule μ_t . The set of all distributions available to player two is

$$\mathcal{Q}^{\mu_t} = \{\mathbf{P}(x_t, \mu_t(x_t)) \in \mathcal{P}_t(\mathbf{d}_t) \text{ over all } x_t, \mathbf{d}_t\}.$$

In \mathcal{Q}^{μ_t} we merely express that for each x_t, μ_t, \mathbf{d}_t , we might have a different distribution. Moreover, a policy π is defined as the decision rule to be used at every period, i.e., $\pi = (\mu_1, \dots, \mu_T)$. A policy π also yields a set of distributions \mathcal{Q}^π which can be used by the second player or adversary, where

$$\mathcal{Q}^\pi = \mathcal{Q}^{\mu_1} \times \mathcal{Q}^{\mu_2} \times \dots \times \mathcal{Q}^{\mu_T}. \quad (3)$$

As the second player will maximize her reward, given policy π , net inventory x_t , and realized demand \mathbf{d}_t , the cost paid by player one from period t to T is

$$V_t^\pi(x_t, \mathbf{d}_t) = \max_{\mathbf{Q} \in \mathcal{Q}^\pi} E^{\mathbf{Q}, \tilde{\mathbf{D}}} \left[\sum_{\tau=t}^T \theta^{\tau-t} C_\tau(x_\tau, \mu_\tau(x_\tau), \tilde{D}_\tau) + \theta^{T+1-t} V_{T+1}(x_{T+1}, \mathbf{d}_{T+1}) \right],$$

where $C_\tau(x_\tau, \mu_\tau(x_\tau), \tilde{D}_\tau)$ denotes the cost incurred in period τ in (1), and $V_{T+1}(x_{T+1}, \mathbf{d}_{T+1})$ is the terminal cost. Also note that \mathbf{Q} defines the distributions $\tilde{D}_\tau, \tau = t, \dots, T$. Unless stated otherwise, we assume that $V_{T+1}(\cdot) = 0$. We also have

$$x_{\tau+1} = \mu_\tau(x_\tau) - \tilde{D}_\tau \quad \text{and} \quad \mathbf{d}_{\tau+1} = [\mathbf{d}_\tau, \tilde{D}_\tau].$$

Since the first player will choose a policy that minimizes the cost, the optimal cost from period t to T given net inventory x_t at time t , and the realized demand \mathbf{d}_t from period 1 to $t-1$, is

$$V_t(x_t, \mathbf{d}_t) = \min_{\pi} \max_{\mathbf{Q} \in \mathcal{Q}^\pi} E^{\mathbf{Q}, \tilde{\mathbf{D}}} \left[\sum_{\tau=t}^T \theta^{\tau-t} C_\tau(x_\tau, \mu_\tau(x_\tau), \tilde{D}_\tau) + \theta^{T+1-t} V_{T+1}(x_{T+1}, \mathbf{d}_{T+1}) \right], \quad (4)$$

for $t = 1, \dots, T$. Note that the model minimizes the maximum expected cost arising from any distribution in the set $\mathcal{P}_t(\mathbf{d}_t)$ for any t , which is known as the minimax robust approach. We next state an optimality equation, which is essential to establish the optimal control policies.

Proposition 2.1. *The optimality equation of the robust model is*

$$V_t(x_t, \mathbf{d}_t) = \min_{y_t \geq x_t} \max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \left\{ \sum_i P_{t,i} \left(C_t(x_t, y_t, D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right) \right\} \quad (5)$$

for $t = 1, \dots, T$, where $\mathcal{P}_t(\mathbf{d}_t)$ is the set of distributions satisfying the goodness-of-fit condition at period t , and $C_t(x_t, y_t, D_{t,i})$ is defined by (1).

Proof. It follows from Theorem 2.1 in Iyengar (2005) when $\mathcal{P}_t(\mathbf{d}_t)$ is arbitrary. If $\mathcal{P}_t(\mathbf{d}_t)$ is convex, the proposition can also be proved by the von Neumann's minimax theorem (see, e.g., von Neumann 1928). \square

An immediate observation from Proposition 2.1 is that we minimize the worst case expected cost over a set of distributions. Therefore, our robust stochastic model may not be as conservative as the classical minimax models, where the worst case is defined by the realized demand instead of distribution, e.g., the minimax model discussed in Section 2.4 of Notzon (1970).

Note that the Bayesian inventory models assume a prior demand distribution, and the posterior distribution at time t is obtained by updating the prior distribution using \mathbf{d}_t , e.g., Iglehart (1964) updates the demand distribution belonging to the exponential and range families after observing realized demand information. Our model only requires the set of distributions $\mathcal{P}_t(\mathbf{d}_t)$ to be a function of the realized demand \mathbf{d}_t . Therefore, we can define it as a singleton updated by a Bayesian rule. In this case, the robust minmax model is reduced to a Bayesian inventory model, which indicates that the Bayesian models are special cases of our minimax model.

Proposition 2.1 also gives us an interpretation of the robust model from a risk measure perspective when set $\mathcal{P}_t(\mathbf{d}_t)$ is convex. Ahmed et al. (2007) establish the correspondence between coherent risk measures and minimax models over convex sets of distributions. From this perspective, our minimax robust model essentially minimizes a coherent risk measure with respect to the total cost. If we consider $\mathcal{P}_t(\mathbf{d}_t) \equiv \mathcal{P}_t$ for any \mathbf{d}_t and t , i.e., the set of distributions is independent of any realized demand in previous periods, then the minimax robust model (5) minimizes a coherent risk measure in any period t and it reduces to the model considered in Ahmed et al. (2007). When the set of distributions $\mathcal{P}_t(\mathbf{d}_t)$ depends on demand realization \mathbf{d}_t , model (5) also minimizes a coherent risk measure in every period t . However, this model is different from that in Ahmed et al. (2007) in the sense that the risk measure in period t is updated by the realized demand in previous periods. Intuitively, if the decision maker lost a significant amount in the previous period, he or she would tend to be more risk-averse in subsequent periods. Therefore, it is reasonable to adjust the risk measure based on the realized demand information \mathbf{d}_t .

In addition, let constant p_t denote the selling price of the product in period t . We can maximize the expected total profit from periods 1 to T by subtracting term $p_t \sum_i D_{t,i} P_{t,i}$ in the objective function of (5). All of the results, such as the optimal policy and the convergence properties, still hold for such an objective function. In addition, if we suppose that all the distributions in set $\mathcal{P}_t(\mathbf{d}_t)$ could have the same expectation $\hat{D}_t(\mathbf{d}_t)$, i.e., constraint $\sum_i D_{t,i} P_{t,i} = \hat{D}_t(\mathbf{d}_t)$ is included in the definition of $\mathcal{P}_t(\mathbf{d}_t)$, then the models that minimize cost and that maximize profit are equivalent to each other. However, as long as the demands follow certain distributions, which are not necessarily

known to the decision maker, the expected total revenue is independent of any inventory decision, i.e., the order quantity in any period t . Therefore, it is sufficient to consider the cost minimization model presented in (5).

3. Properties of Optimal Policies

In this section we study optimal policies of the general robust stochastic model (5). Notzon (1970) and Ahmed et al. (2007) show the optimality of (s, S) policy when the set of distributions in the minimax model is independent of the realized demand \mathbf{d}_t (Ahmed et al. 2007 also assume the set of distributions is convex). Here we extend the optimality of (s, S) policy to the more general model in (5).

We assume that the reader is familiar with standard concepts in inventory theory such as K -convexity and (s, S) policies (see, e.g., Zipkin 2000 and Porteus 2002).

Let us define

$$U_i(y, \mathbf{d}) = h_t (y_t - D_{t,i})^+ + b_t (y_t - D_{t,i})^- + \theta V_{t+1} (y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]), \quad (6)$$

which corresponds to the expected cost incurred from period t to T if the inventory level after receiving the order in period t is y_t and the demand in period t is $D_{t,i}$. Consider the function

$$f(y, \mathbf{d}) = \max_{\mathbf{P} \in \mathcal{P}(\mathbf{d})} \sum_i U_i(y, \mathbf{d}) P_i.^1$$

Since optimality of the (s, S) policy follows directly from K -convexity, first we are going to establish that the function $f(y, \mathbf{d})$ is K -convex in y .

Lemma 3.1. *If $U_i(y, \mathbf{d})$ is K -convex in y for any given \mathbf{d} , then $f(y, \mathbf{d})$ is a K -convex function in y for any given \mathbf{d} .*

Proof. Please refer to the Online Supplement. □

Lemma 3.1 shows that K -convexity is preserved under maximization over a set of distributions. Base on this property, we show the K -convexity of the cost-to-go functions.

Proposition 3.1. *If $V_{t+1}(x_{t+1}, \mathbf{d}_{t+1})$ is a K -convex function in x_{t+1} for any fixed \mathbf{d}_{t+1} , the cost-to-go function $V_t(x_t, \mathbf{d}_t)$ is a K -convex function in x_t for any fixed \mathbf{d}_t , and for any $t = 1, \dots, T$.*

Proof. The proposition is trivially true for $t = T + 1$. Suppose that the proposition holds for period $t + 1$, and consider period t .

To simplify the notation, let us define

$$f_t(y_t, \mathbf{d}_t) = c_t y_t + \max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \sum_i P_{t,i} [h_t (y_t - D_{t,i})^+ + b_t (y_t - D_{t,i})^- + \theta V_{t+1} (y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}])]. \quad (7)$$

¹Note that here we drop subscript t in order to simplify the notation.

Therefore, the optimality equation in (5) is equal to

$$V_t(x_t, \mathbf{d}_t) = -c_t x_t + \min_{y_t \geq x_t} \{K\mathbb{I}(y_t - x_t) + f_t(y_t, \mathbf{d}_t)\}.$$

According to Lemma 3.1, if $V_{t+1}(x_{t+1}, \mathbf{d}_{t+1})$ is K -convex in x_t , $f_t(y_t, \mathbf{d}_t)$ is K -convex in y_t . Let $S_t(\mathbf{d}_t)$ be a global minimizer of $f_t(y_t, \mathbf{d}_t)$ for any given \mathbf{d}_t . Moreover, let $s_t(\mathbf{d}_t)$ be the smallest element of the set $\{s_t(\mathbf{d}_t) \mid s_t(\mathbf{d}_t) \leq S_t(\mathbf{d}_t), f_t(s_t, \mathbf{d}_t) = f_t(S_t, \mathbf{d}_t) + K\}$. According to the properties of K -convex functions (see, e.g., Zipkin 2000 and Porteus 2002), we have

$$V_t(x_t, \mathbf{d}_t) = \begin{cases} K - c_t x_t + f_t(S_t(\mathbf{d}_t), \mathbf{d}_t) & \text{if } x_t \leq s_t(\mathbf{d}_t), \\ -c_t x_t + f_t(x_t, \mathbf{d}_t) & \text{otherwise.} \end{cases}$$

K -convexity of $V_t(x_t, \mathbf{d}_t)$ follows from K -convexity of $f_t(y_t, \mathbf{d}_t)$. \square

From the structure of $V_{t+1}(\cdot)$, we can derive an optimal policy.

Theorem 3.1. *A state dependent (s, S) policy is optimal for the robust stochastic model. More precisely, for any t and \mathbf{d}_t , there exists $S_t(\mathbf{d}_t)$ and $s_t(\mathbf{d}_t)$ such that $S_t(\mathbf{d}_t) - x_t$ units are ordered in period t if $x_t \leq s_t(\mathbf{d}_t)$ and no order is placed otherwise.*

Proof. The structure of the policy follows directly from the proof of Proposition 3.1 and general theory of K -convexity (see, e.g., Zipkin 2000 and Porteus 2002). \square

If there is no fixed cost, then $V_t(x_t, \mathbf{d}_t)$ is convex in x_t for any t . Therefore, a state dependent base-stock policy is optimal, and the base-stock level given the realized demand \mathbf{d}_t is $S_t(\mathbf{d}_t)$.

A drawback from the practical point of view is the fact that s_t and S_t depend on \mathbf{d}_t . We next characterize a special case where different values of \mathbf{d}_t correspond to the same (s, S) levels. Suppose that \mathbf{d}_t and \mathbf{d}'_t denote two different demand realizations from period 1 to $t-1$. Let us assume that if demand realizations in periods 1 to $t-1$ are \mathbf{d}_t or \mathbf{d}'_t , then the same demand realization in period t to T generates the same histogram in any period t, \dots, T . Then vectors \mathbf{d}_t and \mathbf{d}'_t correspond to the same (s, S) levels. To formalize this property, let $s_t(\mathbf{d}_t)$ and $S_t(\mathbf{d}_t)$ (respectively $s_t(\mathbf{d}'_t)$ and $S_t(\mathbf{d}'_t)$) denote the (s, S) levels corresponding to history \mathbf{d}_t (respectively \mathbf{d}'_t). For any $\tau \geq t$, let the vector $[\mathbf{d}_t, \tilde{d}_t, \tilde{d}_{t+1}, \dots, \tilde{d}_{\tau-1}]$ denote the realized demand up to period $\tau-1$ where the demands from periods 1 to $t-1$ are aggregated in vector \mathbf{d}_t , and the realized demand in periods t to $\tau-1$ is labeled by $\tilde{d}_t, \tilde{d}_{t+1}, \dots, \tilde{d}_{\tau-1}$, respectively.

Proposition 3.2. *Let $V_{T+1}(x_{T+1}, \mathbf{d}_{T+1}) = V_{T+1}(x_{T+1}, \mathbf{d}'_{T+1})$ for any x_{T+1} , \mathbf{d}_{T+1} , \mathbf{d}'_{T+1} , and consider any $\tau = t, \dots, T$. Suppose that realizations \mathbf{d}_t and \mathbf{d}'_t give the same number of samples in interval $[D_{\tau,i}, D_{\tau,i+1})$ for any i as long as the realized demand in periods t to $\tau-1$ is the same, i.e.,*

$$N_{\tau,i}([\mathbf{d}_t, \tilde{d}_t, \tilde{d}_{t+1}, \dots, \tilde{d}_{\tau-1}]) = N_{\tau,i}([\mathbf{d}'_t, \tilde{d}_t, \tilde{d}_{t+1}, \dots, \tilde{d}_{\tau-1}])$$

for any i and any realization $[\tilde{d}_t, \tilde{d}_{t+1}, \dots, \tilde{d}_{\tau-1}]$ of $[\tilde{D}_t, \tilde{D}_{t+1}, \dots, \tilde{D}_{\tau-1}]$. Then we have $s_t(\mathbf{d}_t) = s_t(\mathbf{d}'_t)$, $S_t(\mathbf{d}_t) = S_t(\mathbf{d}'_t)$, and $V_t(x_t, \mathbf{d}_t) = V_t(x_t, \mathbf{d}'_t)$ for any x_t .

Proof. Consider period T . According to the assumption stated, $N_{T,i}(\mathbf{d}_T) = N_{T,i}(\mathbf{d}'_T)$ for any i , and hence we have $n_T(\mathbf{d}_T) = n_T(\mathbf{d}'_T)$ and $\mathcal{P}_T(\mathbf{d}_T) = \mathcal{P}_T(\mathbf{d}'_T)$. By assumption on $V_{T+1}(\cdot)$, we obtain $s_T(\mathbf{d}_T) = s_T(\mathbf{d}'_T)$ and $S_T(\mathbf{d}_T) = S_T(\mathbf{d}'_T)$ from Theorem 4.1. Moreover, the result $V_T(x_T, \mathbf{d}_T) = V_T(x_T, \mathbf{d}'_T)$ follows from (5).

Suppose that the proposition is true for any period $\tau > t$. Hence, $V_{t+1}(x_{t+1}, [\mathbf{d}_t, D_{t,i}]) = V_{t+1}(x_{t+1}, [\mathbf{d}'_t, D_{t,i}])$ for any x_{t+1} and i . Moreover, we have $N_{t,i}(\mathbf{d}_t) = N_{t,i}(\mathbf{d}'_t)$ for any i , which implies $n_t(\mathbf{d}_t) = n_t(\mathbf{d}'_t)$ and $\mathcal{P}_t(\mathbf{d}_t) = \mathcal{P}_t(\mathbf{d}'_t)$. According to Theorem 4.1 and (5), the results hold for period t . \square

Suppose that we use the same bin intervals $[D_{t,i}, D_{t,i+1})$ for any period t in the planning horizon. Furthermore, let us assume that we update the histogram in time period t only based on the realized demand in periods 1 to $t - 1$, or, for example, given a fixed n , we update the histogram in time period t only based on realized demand in time periods $t - n$ through $t - 1$. Observe that these two scenarios do not allow any forecasting based on the just realized demand. From Proposition 3.2, it now follows that the number of different (s, S) levels at time t cannot exceed the number of bins to the power of t . This observation substantially reduces the computational burden.

4. Robust Models Based on the Chi-Square Test

The most widely used goodness-of-fit test is the chi-square test (see, e.g., Chernoff and Lehmann 1954) with the statistical test

$$\sum_i \frac{(N_{t,i}(\mathbf{d}_t) - n_t(\mathbf{d}_t)P_{t,i})^2}{n_t(\mathbf{d}_t)P_{t,i}} \leq \chi_t^2 \quad t = 1, \dots, T,$$

where parameter χ_t^2 controls how close the observed sample data is to the estimated expected number of observations according to the fitted distribution $(P_{t,i})_{i=1, \dots, M_t}$.

More specifically, suppose that k is the number of bins, c is the number of estimated parameters for the fitted distribution (e.g., $c = 2$ for normal distributions due to the mean and variance), and consider the null hypothesis H_0 that the observations are independent random samples drawn from the fitted distribution. Chernoff and Lehmann (1954) show that if H_0 is true, the test statistic converges to a distribution function that lies between the distribution functions of chi-square distributions with $k - 1$ and $k - c - 1$ degrees of freedom. Let α denote the significance level, and consider $\chi_{k-1, 1-\alpha}^2$ such that $F(\chi_{k-1, 1-\alpha}^2) = 1 - \alpha$, where $F(x)$ is the distribution function of the chi-square distribution with $k - 1$ degrees of freedom. It is often recommended that we reject the null hypothesis at the significance level α if the test statistic is greater than $\chi_{k-1, 1-\alpha}^2$ (see, e.g., Law and Kelton 2000). In our context, $k = M_t$ and α , whose interpretation is as above, is given by the decision maker.

Since $P_{t,i}$ should define a probability distribution, we have $\sum_i P_{t,i} = 1$ and $P_{t,i} \geq 0$. Let \mathbf{P}_t denote the vector of $(P_{t,i})_i$. The set of distributions that satisfy the chi-square test is

$$\mathcal{P}_t(\mathbf{d}_t) = \left\{ \mathbf{P}_t \left| \mathbf{A}_t \mathbf{P}_t = \mathbf{b}_t, \sum_i \frac{(N_{t,i}(\mathbf{d}_t) - n_t(\mathbf{d}_t)P_{t,i})^2}{n_t(\mathbf{d}_t)P_{t,i}} \leq \chi_t^2, \mathbf{P}_t \geq \mathbf{0} \right. \right\} \quad t = 1, \dots, T. \quad (8)$$

The linear constraints $\mathbf{A}_t \mathbf{P}_t = \mathbf{b}_t$ capture at least the fact that $\sum_i P_{t,i} = 1$. They can also be used to model more complicated properties of the distribution set, such as constraints on the expected value, any moment or desired percentiles of the distributions. It is straightforward to establish the compactness of $\mathcal{P}_t(\mathbf{d}_t)$.

We next give an alternative optimality equation that exploits the structure of (8). We first provide an alternative characterization of $\mathcal{P}_t(\mathbf{d}_t)$. We assume that every norm is the Euclidean norm.

Lemma 4.1. *The set of demand distributions $\mathcal{P}_t(\mathbf{d}_t)$ defined in (8) is equivalent to the projection of the set*

$$\left\{ (\mathbf{P}_t, \mathbf{Q}_t) \mid \mathbf{A}_t \mathbf{P}_t = \mathbf{b}_t, \sum_i N_{t,i}(\mathbf{d}_t)^2 Q_{t,i} - n_t(\mathbf{d}_t)^2 \leq n_t(\mathbf{d}_t) \chi_t^2, \left\| \begin{bmatrix} P_{t,i} - Q_{t,i} \\ 2 \end{bmatrix} \right\| \leq P_{t,i} + Q_{t,i} \right\}$$

on the space of \mathbf{P}_t .

Proof. Since $\sum_i P_{t,i} = 1$ and $\sum_i N_{t,i}(\mathbf{d}_t) = n_t(\mathbf{d}_t)$, we have

$$\begin{aligned} \sum_i \frac{(N_{t,i}(\mathbf{d}_t) - n_t(\mathbf{d}_t) P_{t,i})^2}{n_t(\mathbf{d}_t) P_{t,i}} &= \sum_i \frac{N_{t,i}(\mathbf{d}_t)^2}{n_t(\mathbf{d}_t) P_{t,i}} - \sum_i 2N_{t,i}(\mathbf{d}_t) + \sum_i n_t(\mathbf{d}_t) P_{t,i} \\ &= \sum_i \frac{N_{t,i}(\mathbf{d}_t)^2}{n_t(\mathbf{d}_t) P_{t,i}} - n_t(\mathbf{d}_t). \end{aligned}$$

As χ_t^2 and $n_t(\mathbf{d}_t)$ are finite, we have $P_{t,i} > 0$ for any i . Therefore,

$$\sum_i \frac{N_{t,i}(\mathbf{d}_t)^2}{n_t(\mathbf{d}_t) P_{t,i}} - n_t(\mathbf{d}_t) \leq \chi_t^2$$

is equivalent to

$$\sum_i N_{t,i}(\mathbf{d}_t)^2 Q_{t,i} - n_t(\mathbf{d}_t)^2 \leq n_t(\mathbf{d}_t) \chi_t^2, \frac{1}{P_{t,i}} \leq Q_{t,i}, P_{t,i}, Q_{t,i} > 0.$$

Obviously, we have

$$\frac{1}{P_{t,i}} \leq Q_{t,i}, P_{t,i}, Q_{t,i} > 0 \iff P_{t,i} Q_{t,i} \geq 1, P_{t,i}, Q_{t,i} \geq 0 \iff \begin{bmatrix} P_{t,i} & 1 \\ 1 & Q_{t,i} \end{bmatrix} \succeq 0.$$

Note that the eigenvalues of the matrix $\begin{bmatrix} P_{t,i} & 1 \\ 1 & Q_{t,i} \end{bmatrix}$ are $\frac{P_{t,i} + Q_{t,i} \pm \sqrt{(P_{t,i} - Q_{t,i})^2 + 4}}{2}$, therefore the positive semidefinite constraint is equivalent to

$$\frac{P_{t,i} + Q_{t,i} - \sqrt{(P_{t,i} - Q_{t,i})^2 + 4}}{2} \geq 0 \iff \left\| \begin{bmatrix} P_{t,i} - Q_{t,i} \\ 2 \end{bmatrix} \right\| \leq P_{t,i} + Q_{t,i},$$

which proves the proposition. \square

Lemma 4.1 shows that the set $\mathcal{P}_t(\mathbf{d}_t)$ can be defined by a set of linear and second order cone constraints (see, e.g., Lobo et al. 1998). Note that the second order cone constraints are a special class of positive semidefinite constraints and they have better computational properties than general positive semidefinite constraints. This alternative definition of the set $\mathcal{P}_t(\mathbf{d}_t)$ also suggests a compact optimality equation.

Proposition 4.1. *The optimality equation of the robust stochastic model (5) is equivalent to*

$$\begin{aligned}
V_t(x_t, \mathbf{d}_t) = & \min_{y_t, \mathbf{U}_t, \mathbf{p}_t, \mathbf{u}_t, \lambda_t} \quad K\mathbb{I}(y_t - x_t) + c_t(y_t - x_t) + \mathbf{p}_t^T \mathbf{b}_t - 2 \sum_i u_{t,i} N_{t,i}(\mathbf{d}_t) \\
& + \lambda_t (n_t(\mathbf{d}_t)^2 + n_t(\mathbf{d}_t) \chi_t^2) \\
\text{s.t.} \quad & \left\| \begin{bmatrix} \mathbf{p}_t^T - U_{t,i} \mathbf{A}_{t,i} - \lambda_t \\ 2u_{t,i} \end{bmatrix} \right\| \leq \mathbf{p}_t^T - U_{t,i} \mathbf{A}_{t,i} + \lambda_t \quad \text{for every } i \quad (9) \\
& U_{t,i} \geq h_t(y_t - D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \quad \text{for every } i \\
& U_{t,i} \geq b_t(y_t - D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \quad \text{for every } i \\
& y_t \geq x_t,
\end{aligned}$$

for any $t = 1, \dots, T$.

Proof. Please refer to the Online Supplement. \square

Note that this is not the standard optimality equation since $V_{t+1}(\cdot)$ is present in constraints and not the objective function. We use it later to obtain computationally tractable control policies.

4.1 Computation of (s, S) Levels

Next we give a computational approach to compute $s_t(\mathbf{d}_t)$ and $S_t(\mathbf{d}_t)$.

Theorem 4.1. *Let $S_t(\mathbf{d}_t)$ be an optimal solution to the minimization problem*

$$\begin{aligned}
\min_{y_t, \mathbf{U}_t, \mathbf{p}_t, \mathbf{u}_t, \lambda_t} \quad & c_t y_t + \mathbf{p}_t^T \mathbf{b}_t - 2 \sum_i u_{t,i} N_{t,i}(\mathbf{d}_t) + \lambda_t (n_t(\mathbf{d}_t)^2 + n_t(\mathbf{d}_t) \chi_t^2) \\
\text{s.t.} \quad & \left\| \begin{bmatrix} \mathbf{p}_t^T \mathbf{A}_{t,i} - U_{t,i} - \lambda_t \\ 2u_{t,i} \end{bmatrix} \right\| \leq \mathbf{p}_t^T \mathbf{A}_{t,i} - U_{t,i} + \lambda_t \quad \text{for every } i \\
& U_{t,i} \geq h_t(y_t - D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \quad \text{for every } i \\
& U_{t,i} \geq b_t(D_{t,i} - y_t) + \theta V_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \quad \text{for every } i,
\end{aligned}$$

and let $s_t(\mathbf{d}_t)$ be the smallest element of the set

$$\{s_t(\mathbf{d}_t) \mid s_t(\mathbf{d}_t) \leq S_t(\mathbf{d}_t), f_t(s_t, \mathbf{d}_t) = f_t(S_t, \mathbf{d}_t) + K\},$$

where $f_t(y_t, \mathbf{d}_t)$ is defined by (7).

A state dependent (s, S) policy is optimal for the robust stochastic model (5) with $\mathcal{P}_t(\mathbf{d}_t)$ defined by (8), and the (s, S) levels are given by $s_t(\mathbf{d}_t)$ and $S_t(\mathbf{d}_t)$ respectively. If there is no fixed cost, a state dependent base-stock policy is optimal, and the base-stock level given the realized demand \mathbf{d}_t is $S_t(\mathbf{d}_t)$.

Proof. The minimization problem to calculate $S_t(\mathbf{d}_t)$ follows from the alternative optimality equation (9). \square

Consider the models where the historical data used for period t is independent of the realized demand from periods 1 to $t - 1$, i.e., the number of observations $N_{t,i}$ in the i th bin and the total number of available observations n_t are constant for any realized demand \mathbf{d}_t . Therefore, the set of distributions that satisfy the chi-square test is defined by

$$\mathcal{P}_t = \left\{ \mathbf{P}_t \left| \mathbf{A}_t \mathbf{P}_t = \mathbf{b}_t, \sum_i \frac{(N_{t,i} - n_t P_{t,i})^2}{n_t P_{t,i}} \leq \chi_t^2, \mathbf{P}_t \geq \mathbf{0} \right. \right\} \quad t = 1, \dots, T. \quad (10)$$

In this case, the optimality equation of the robust model is reduced to

$$V_t(x_t) = \min_{y_t \geq x_t} \max_{\mathbf{P}_t \in \mathcal{P}_t} \left\{ \sum_i P_{t,i} \left(C_t(x_t, y_t, D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}) \right) \right\} \quad t = 1, \dots, T, \quad (11)$$

where \mathcal{P}_t and $C_t(x_t, y_t, D_{t,i})$ are defined by (10) and (1) respectively.

Alternatively, it can be written as

$$\begin{aligned} V_t(x_t) = & \min_{y_t, \mathbf{U}_t, \mathbf{P}_t, \mathbf{u}_t, \lambda_t} K \mathbb{I}(y_t - x_t) + c_t(y_t - x_t) + \mathbf{p}_t^T \mathbf{b}_t - 2 \sum_i u_{t,i} N_{t,i} + \lambda_t (n_t^2 + n_t \chi_t^2) \\ \text{s.t.} & \left\| \left[\begin{array}{c} \mathbf{p}_t^T - U_{t,i} \mathbf{A}_{t,i} - \lambda_t \\ 2u_{t,i} \end{array} \right] \right\| \leq \mathbf{p}_t^T - U_{t,i} \mathbf{A}_{t,i} + \lambda_t \quad \text{for every } i \\ & U_{t,i} \geq h_t(y_t - D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}) \quad \text{for every } i \\ & U_{t,i} \geq b_t(D_{t,i} - y_t) + \theta V_{t+1}(y_t - D_{t,i}) \quad \text{for every } i \\ & y_t \geq x_t, \end{aligned} \quad (12)$$

for any $t = 1, \dots, T$.

The corresponding optimal (s, S) policy levels are also independent of the realized demand \mathbf{d}_t .

Theorem 4.2. *The (s, S) policy is optimal for the robust stochastic model (11). In particular, let S_t be the optimal solution to the minimization problem*

$$\begin{aligned} \min_{y_t, \mathbf{U}_t, \mathbf{P}_t, \mathbf{u}_t, \lambda_t} & c_t y_t + \mathbf{p}_t^T \mathbf{b}_t - 2 \sum_i u_{t,i} N_{t,i} + \lambda_t (n_t^2 + n_t \chi_t^2) \\ \text{s.t.} & \left\| \left[\begin{array}{c} \mathbf{p}_t^T \mathbf{A}_{t,i} - U_{t,i} - \lambda_t \\ 2u_{t,i} \end{array} \right] \right\| \leq \mathbf{p}_t^T \mathbf{A}_{t,i} - U_{t,i} + \lambda_t \quad \text{for every } i \\ & U_{t,i} \geq h_t(y_t - D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}) \quad \text{for every } i \\ & U_{t,i} \geq b_t(D_{t,i} - y_t) + \theta V_{t+1}(y_t - D_{t,i}) \quad \text{for every } i, \end{aligned}$$

and let s_t be the smallest element of the set

$$\{s_t \mid s_t \leq S_t, f_t(s_t) = f_t(S_t) + K\},$$

where

$$f_t(y_t) = c_t y_t + \max_{\mathbf{P}_t \in \mathcal{P}_t} \sum_i P_{t,i} [h_t(y_t - D_{t,i})^+ + b_t(y_t - D_{t,i})^- + \theta V_{t+1}(y_t - D_{t,i})].$$

The policy is to order $S_t - x_t$ units in period t if $x_t \leq s_t$, and no order is placed otherwise.

Without fixed procurement cost, a base-stock policy is optimal, that is, $S_t - x_t$ units are ordered in period t if $x_t \leq S_t$, and no order is placed otherwise.

4.2 Summary of Convergence Results

In this subsection, we explore the case of the bins in the histogram being defined by distinctive values in the sample data, and we study the performance of the robust model when the number of samples increase. The important results are summarized in the remaining part of this subsection, and the details of the analysis are presented in Appendix A.

The majority of the results are built on the following convergence property: as χ_t^2 approaches to 0, the cost-to-go function $V_t(x_t, \mathbf{d}_t)$ of the robust model converges to the corresponding cost-to-go function of the stochastic model where the demand distributions follow the empirical distribution defined by the histogram (c.f. Proposition A.1).

If there exists no fixed procurement, then based on Proposition A.1, if (i) χ_t^2 converges to 0 and (ii) the empirical distribution functions converge pointwise to the true distribution function, then the cost-to-go function of the robust model converges to that of the stochastic model with the true demand distribution (c.f. Proposition A.2). In particular, if the demand distributions for each time period follow independent continuous distributions, the convergence of the robust model holds as long as the sample size approaches infinity and χ_t^2 converges to 0 (c.f. Corollary A.1).

For the models with both fixed and variable procurement cost, we consider the case where the demands follow a discrete distribution over a finite set. Similar to Proposition A.2, the cost-to-go function of the robust model converges to that of the stochastic model with the true demand distribution under the conditions that (i) χ_t^2 converges to 0 and (ii) the empirical distribution converges to the true distribution (c.f. Proposition A.4). We have also identified a condition under which the convergence holds without χ_t^2 approaching 0 (c.f. Proposition A.5). Moreover, if the demand distributions are independent across different periods, as long as the sample size goes to infinity, the cost-to-go function of the robust model converges to that of the stochastic model with the true demand distribution, and the rate of convergence is $O(1/\sqrt{k})$, where k denotes the sample size (c.f. Corollary A.3).

The convergence study not only provides the asymptotic performance of the robust model when the sample size approaches infinitely, but also guarantees that the robust models with small bin sizes and small χ^2 values perform well in the presence of a significant number of samples.

5. Computational Results

In this section, we describe computational experiments and present numerical results to support the effectiveness of the minimax robust model based on the chi-square test. In particular, the robust model proposed in Section 4 is compared with (i) the approach which first fits the historical data and then solves the inventory optimization model using the fitted distribution and (ii) the robust model based on Delage and Ye (2010). These two comparisons are presented in the following two subsections, respectively.

5.1 Comparison with Separated Data-Fitting and Inventory Optimization

As we have mentioned in the previous sections, the traditional approach is to fit the historical data with a distribution and then apply stochastic inventory optimization using the fitted distribution. The main objective of our experiments is to compare performances of this separated approach and the studied minimax robust model with respect to optimality and robustness. At the same time, we would like to assess sensitivity of the robust model to the choices of the bin sizes and χ^2 parameters, and provide an empirical approach to choose these values.

We consider inventory control problems without fixed ordering costs. Following the notation in the previous sections, we let T denote the planning horizon and c_t, h_t, b_t denote the variable order cost, unit inventory holding cost, and backorder cost for any period $t, t = 1, \dots, T$, respectively. The demand distributions for any period t are assumed to be i.i.d. In the robust model, we restrict ourselves to the case of equal bin sizes and these, together with χ^2 , are the same for every period in the planning horizon. To simplify the notation, pair $\langle \epsilon, \chi^2 \rangle$ denotes the choice of the bin size ϵ and χ^2 in the robust model.

The procedure of the computational experiments is as follows.

Step 1. Suppose that the underlying demand distribution has support $\{0, 1, \dots, \bar{D}\}$. We randomly generate a distribution among all distributions whose support is a subset of $\{0, 1, \dots, \bar{D}\}$. In particular, we pick distribution

$$p_i = P(\tilde{D}_t = i) = \frac{U_i}{\sum_{i=0}^{\bar{D}} U_i},$$

for any $i = 0, 1, \dots, \bar{D}$, where U_i for all i are i.i.d. random variables uniformly distributed in the interval $[0, 1]$. We refer to the distribution $\mathbf{p} = \{p_i\}_i$ as the true distribution.

Step 2. Generate n random samples according to the true distribution selected in Step 1.

Step 3. Fit the samples obtained in Step 2 using Crystal Ball and then choose the l best-fitted distributions according to the χ^2 goodness-of-fit statistic.

Step 4. Solve the standard stochastic inventory control problem with distributions generated in Steps 1 and 3.

Step 5. Solve the robust inventory control model using a set of bin-size and χ^2 combinations.

Step 6. Evaluate the total expected cost with respect to the true distribution \mathbf{p} corresponding to the policies of the stochastic models and robust models computed in Steps 4 and 5. We use this step to investigate the optimality of the robust models.

Step 7. The n samples generated in Step 2 define the empirical distribution $\hat{\mathbf{p}}$ such that

$$\hat{p}_i = \frac{\text{the number of times value } i \text{ appears in the } n \text{ samples}}{n}$$

for any $i = 0, 1, \dots, \bar{D}$. Let $\delta = \mathbf{p} - \hat{\mathbf{p}}$. We generate m random permutations of vector δ and denote the j th permutation of the coordinates by δ^j . Vector $\hat{\mathbf{p}}^j = \hat{\mathbf{p}} + \delta^j$ also defines a distribution.² Note that $\hat{\mathbf{p}}^j$ is equal to \mathbf{p} if $\delta^j = \delta$, i.e., when δ^j is not permuted.

For each distribution defined by vector $\hat{\mathbf{p}}^j$, we can evaluate the corresponding cost for each policy computed in Steps 4 and 5. Therefore, we obtain m costs for each policy and we report the conditional value-at-risk³ (CVaR) at the 5% level of the m costs for each policy. The purpose of this step is to understand the robustness of different approaches.

Let us consider a 10-period problem. The support for the demand distribution is assumed to be the set $\{0, 1, \dots, 29\}$, i.e., $\bar{D} = 29$. The cost parameters c_t , h_t and b_t are generated independently according to uniform distributions within the intervals $[12, 15]$, $[2, 5]$ and $[22, 25]$, respectively. Following the computational procedure, we first draw $n = 20$ samples from the selected true distribution. Fitting the samples using Crystal Ball, the three best-fitted distributions according to the chi-square values are negative binomial, Poisson, and beta. The true distribution \mathbf{p} , sample frequency $\hat{\mathbf{p}}$ and the three distributions are displayed in Figure 1.

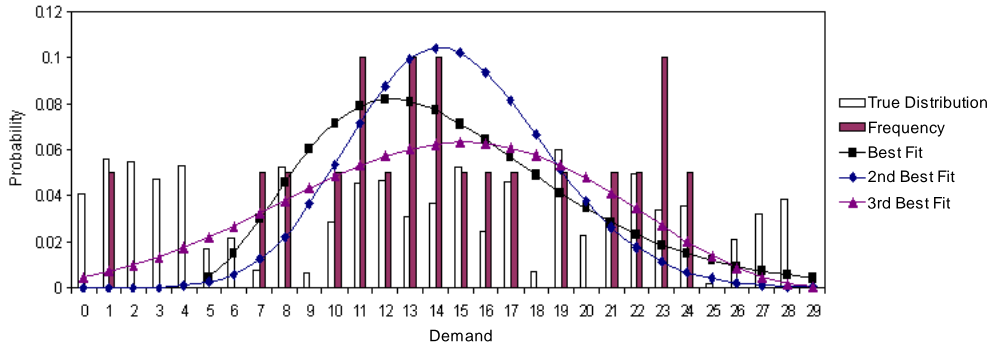


Figure 1: True Distribution, Frequency and Fitted Distributions with 20 Samples

In Steps 4 and 5 of our procedure, we compute the base-stock levels corresponding to different models: the stochastic model using the true distribution, the stochastic model using the three best-fitted distributions, and robust models with different bin-size and χ^2 value combinations. In particular, the following set of bin-size and χ^2 value combinations are considered: $\langle 3, 1 \rangle$, $\langle 3, 3 \rangle$, $\langle 3, 5 \rangle$, $\langle 5, 1 \rangle$, $\langle 5, 3 \rangle$, $\langle 5, 5 \rangle$.

As stated in our analysis, the robust model picks the demand distribution based on the on-hand inventory after the order is received, i.e., the order-up-to level y_t . Although we use the same histogram in each period, the demand distribution returned by the robust model depends also on t . We use the robust model with the bin-size/ χ^2 value $\langle 3, 3 \rangle$ to illustrate these properties.

In Figures 2 and 3, and Table 1, we use a simple representative sample of cost parameters. Figure 2 shows the robust distributions for the last period $t = 10$ and the first period $t = 1$ when

²If $\hat{\mathbf{p}}^j$ contains any negative component, we set $\hat{\mathbf{p}}^j$ to be the positive part of $\hat{\mathbf{p}}^j$ plus a random permutation of its negative part, and we repeat this process until $\hat{\mathbf{p}}^j \geq \mathbf{0}$.

³Given random variable X , the conditional value-at-risk at a quantile-level q is defined as $E[X|X < \mu]$ where μ is defined by $P(X < \mu) = q$.

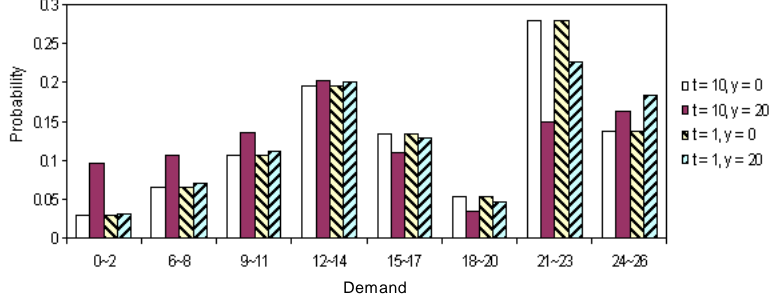


Figure 2: Demand Distributions Returned by the Robust Model with Bin Size = 3 and $\chi^2 = 3$

the inventory levels after receiving the order y_t are 0 and 20 respectively. For both periods, the distributions returned by the robust model for $y_t = 20$ have lower probabilities in the region 15 to 26 than those for $y_t = 0$. The intuition behind this observation is that the robust model picks a demand distribution maximizing the expected cost. For any possible value of the demand, we incur a certain cost corresponding to $U_{t,i}(y_t, \mathbf{d}_t)$ defined in (6). Therefore, the robust model chooses a lower probability for demand values with lower costs. Value $y_t = 20$ is very close to the demand when the demand falls in the region 15 to 26. The amount we over- or under-order is low and hence the corresponding over- or under-order cost is also low.⁴ Therefore, the corresponding costs associated with the demand values are lower than the costs corresponding to other demand values. As a result, the robust model assigns lower probabilities in these regions compared with the case when $y_t = 0$.

If we compare the robust distributions when $y_t = 20$ for period 10 and period 1, we observe that the probability for period 10 is higher for small demand values. This can also be explained by the tradeoff between the over- and under-order costs. In the last period, the over-order cost is $c_{10} + h_{10}$ and the under-order cost is b_{10} since we set $V_{T+1}(\cdot) = 0$. For any earlier period $t < 10$, the over-order costs are significantly lower as we can carry the inventory to the next period and save the order cost c_t , but the under-order cost is $b_t + c_{t+1}$ since we not only pay the backorder cost but also procure the product in period $t + 1$ to satisfy the unmet demand in period t . When $y_t = 20$, we pay the over-order costs when the demands are low (e.g., in the region 0 to 11), and the under-order costs are incurred when the demand are high (e.g., in the region 21 to 26). As the over-order costs are higher and the under-order costs are lower in the last period, it implies that the ratio between the costs for low demands and the costs for high demands is greater in period 10 than period 1. This is the reason why the robust model assigns higher probabilities for low demands in period 10.

On the other hand, the robust distributions when $y_t = 0$ are almost the same for the two periods with $t = 10$ and $t = 1$. In this case, we only have the under-order cost no matter if the demand is high or low. Although the under-order cost is higher in period 1 than period 10, the ratios between the costs for low and high demands are almost the same for periods 1 and 10. Therefore, the worst

⁴In this section, the over-order (under-order, respectively) cost includes not only the inventory holding cost h_t (backorder cost b_t , respectively) incurred in period t , but also the impact of over-order (under-order, respectively) in period t based on the cost-to-go function $V_{t+1}(\cdot)$.

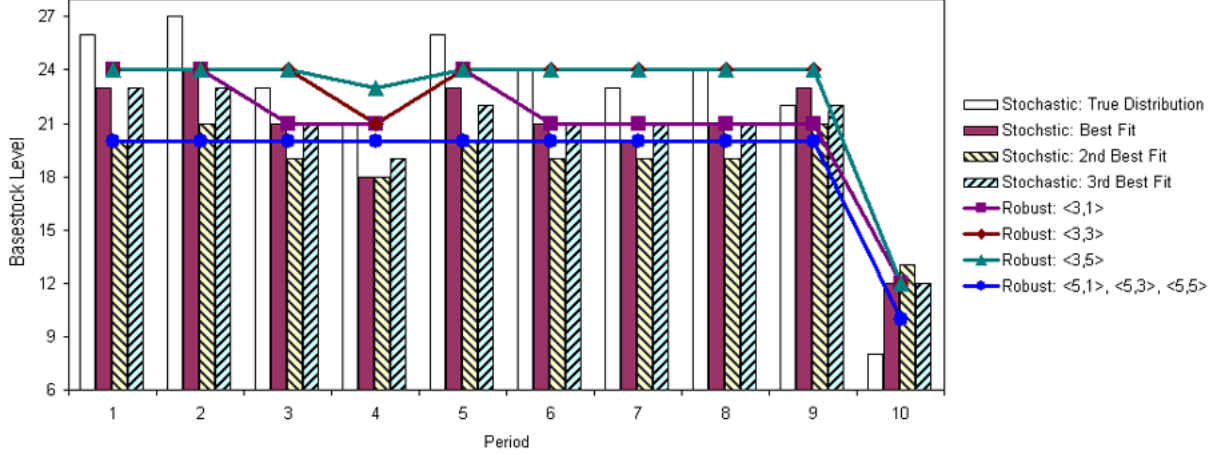


Figure 3: Base-Stock Levels Computed Using Different Models

case distributions are similar for these two periods.

The base-stock levels computed in Steps 4 and 5 are displayed in Figure 3. For any of the stochastic or robust models, the base-stock level for period 10 is significantly lower than the remaining periods. As explained before, this is caused by the fact that the overorder cost is much higher while the underorder cost is lower in period 10 because of $V_{T=1}(\cdot) = 0$, and thus we should order less in that period. In addition, the base-stock level for period 4 is slightly lower for most of the models since period 4 has the highest order and inventory holding cost while its backorder cost is relatively low.

For the three robust models with the bin-size 3, the base-stock levels are nondecreasing with respect to the χ^2 value, since the sets of distributions are inclusion-wise increasing in the χ^2 value. In our instances, the backorder cost is much higher than the inventory holding cost. Intuitively, the worst case distribution should assign higher probabilities for high demand values. Therefore, the larger the χ^2 value is, the higher the probabilities for high demand values in the worst case distribution, and hence we should order more to minimize the worst case expected cost. As a result, the base-stock levels are higher for the robust models with greater χ^2 values. However, if we set the bin-size to 5 for the robust models, the base-stock levels are the same when the χ^2 values are equal to 1, 3 and 5. This observation indicates that the base-stock levels are less sensitive to the χ^2 values when we have larger bins.

We use Steps 6 and 7 to understand the performance of different models. The results are summarized in Table 1. The first four columns correspond to the results for the stochastic models using true distribution \mathbf{p} and the three best-fitted distributions, respectively. The next four columns show the results for the robust models. Note that the last column corresponds to the robust models with bin-size 5 and χ^2 values 1, 3 and 5. These three robust models have the same performance for this example as they have the same base-stock levels. We show the expected cost for different models with respect to the true distribution in the first line, which corresponds to the output of Step 6. In the second line, we report the output of Step 7, i.e., the CVaR at 5% level for the costs of $m = 1000$ distributions generated by $\hat{\mathbf{p}}$ plus random permutations of $\mathbf{p} - \hat{\mathbf{p}}$. For the purpose of

comparison, the numbers in Table 1 are calculated by subtracting the cycle stock order cost, i.e., $\left(\sum_{t=1}^T c_t\right) \left(\sum_{i=1}^D ip_i\right)$, from the original cost or CVaR, and normalizing with respect to that of the stochastic model using true distribution.

	Stochastic Models				Robust Models			
	True Dist	Best Fit	2nd Best Fit	3rd Best Fit	$\langle 3, 1 \rangle$	$\langle 3, 3 \rangle$	$\langle 3, 5 \rangle$	$\langle 5, 1 \text{ or } 3 \text{ or } 5 \rangle$
Cost	1	1.0595	1.1834	1.0582	1.0415	1.0211	1.0249	1.1511
CVaR	1	1.0486	1.1802	1.0511	1.0356	0.9774	0.9739	1.1662

Table 1: Performance of Different Models for the Instance in Figure 1

Obviously, the stochastic model using the true distribution gives the lowest expected cost. The output of Step 7, CVaR, also indicates that this model is robust with respect to perturbations in the input distribution as it has the third lowest CVaR, which is only 2.61% higher than the lowest CVaR.

For the three stochastic models using fitted distributions, the models using the 1st and 3rd best-fitted distributions have a very similar performance. The best-fit case has the best performance among the fitted stochastic models as its CVaR is 0.25% better than the 3rd best-fit stochastic model and the cost is only 0.13% higher than that. The performance of the model using the 2nd best distribution is much worse compared with the other two. Its cost and CVaR values are at least 12% higher than those of the remaining two models.

The three robust models with bin-size 3 outperform all of the stochastic models using fitted distributions in terms of both optimality (cost) and robustness (CVaR). The robust models with bin-size 5 also have better values of the cost and CVaR than the stochastic model using the 2nd best-fitted distribution. In particular, the robust models with bin-size/ χ^2 value combinations of $\langle 3, 3 \rangle$ and $\langle 3, 5 \rangle$ are significantly better than the stochastic models using fitted distributions. They reduce the cost by more than 3% and CVaR by more than 7% when comparing with the fitted stochastic models. Among the robust models we prefer the model with bin-size/ χ^2 value combination $\langle 3, 3 \rangle$, since it improves the cost by 0.38% at the price of a 0.35% increase in CVaR.

There are mainly two reasons why the robust model outperforms the stochastic model with the best-fit distribution. First, the sample size is relatively small and hence the fitted distributions could be significantly different from the true distribution. The robust model, on the other hand, recognizes the difference between the samples and the true distribution and corrects the error by considering a set of distributions close to the empirical histogram. Second, although Crystal Ball considers 16 families of commonly used parametric distributions, it is still possible that the true distribution does not follow any one of these 16 families of distributions. In fact, most distribution encountered in practice cannot be described using parametric families such as the uniform or Poisson distributions. Therefore, as long as the true distribution does not belong to any parametric family of distributions, the best-fit distribution does not match the true distribution even if the sample size goes to infinity. The robust approach based on the chi-squared test, which does not make any assumption regarding the parametric family of the underlying distribution, is much more flexible in this perspective.

Next we repeated the experiment 10 times from Step 1 to Step 7, i.e., each time with a different true distribution, demand data and cost parameters. Table 2 shows the average and standard deviation of the cost and CVaR values for the 10 data samples for the stochastic model using the true distribution, the stochastic model using the best-fitted distribution as well as the 6 robust models already considered. All robust models have lower average and standard deviation of cost and CVaR compared with the stochastic model using the best-fitted distribution. In terms of both optimality (cost) and robustness (CVaR), the performance of our robust models is better on average (smaller average) and more stable (smaller standard deviation) than the stochastic model using the best-fitted distribution.

	Stochastic Models		Robust Models					
	True Dist	Best Fit	$\langle 3, 1 \rangle$	$\langle 3, 3 \rangle$	$\langle 3, 5 \rangle$	$\langle 5, 1 \rangle$	$\langle 5, 3 \rangle$	$\langle 5, 5 \rangle$
Cost Average	1	1.0902	1.0412	1.0361	1.0499	1.0725	1.0597	1.0570
Cost Std. Dev.	0	0.0893	0.0210	0.0302	0.0330	0.0736	0.0511	0.0519
CVaR Average	1	1.0894	1.0014	0.9745	0.9746	1.0648	1.0281	1.0212
CVaR Std. Dev.	0	0.1142	0.0507	0.0327	0.0355	0.0917	0.0639	0.0618

Table 2: Performance of Different Models in 10 Instances

The robust models with bin-size 3 have lower values of average and standard deviation of both measures than the robust models with bin-size 5. Moreover, the robust models with higher χ^2 values, e.g., when χ^2 is set to 3 or 5, have lower CVaR than those with χ^2 values set to 1. This observation agrees with our understanding that increasing χ^2 values can improve the robustness of the models. However, it may also affect the cost of the models, e.g., the average cost for the $\langle 3, 5 \rangle$ robust model is 0.8% higher than that of the $\langle 3, 1 \rangle$ robust model.

The robust model with bin-size 3 and χ^2 value 3 has the lowest average cost, lowest CVaR, and lowest standard deviation of CVaR among all robust models, and its standard deviation of the cost is the second lowest. This agrees with our suggestion drawn from Figure 1: the robust model with bin-size/ χ^2 value combination $\langle 3, 3 \rangle$ should be the best among the robust models.

Figure 4 shows the cost and CVaR values for the stochastic model using the best-fitted distribution and the $\langle 3, 3 \rangle$ robust model in each of the 10 instances. The cost values of the $\langle 3, 3 \rangle$ robust model are at least 7.5% lower than the stochastic model with the best-fitted distribution for instances 5, 8, 9 and 10. The improvement in instances 9 and 10 even exceeds 20%. The cost values of instances 2 and 3 are almost the same for both models. Instance 7 is the only case where the cost of the robust model is more than 2% (2.04% to be exact) higher than the cost of the stochastic model.

The values of CVaR for the robust model are less than one for 7 out of the 10 instances, they are very close to one (at most 0.04% higher than one) for the other 2 instances, and the largest value is 1.0150. On the other hand, the values of CVaR for the stochastic model with the best-fitted distribution is less than one only for 3 instances and the largest value is 1.2923. We conclude that the $\langle 3, 3 \rangle$ robust model is much more robust compared with the stochastic model using the true distribution.




Figure 4: The Stochastic Model Using Best-Fitted Distribution vs. the Robust Model with Parameters $\langle 3, 3 \rangle$ for 10 Instances

In order to understand the sensitivity of different models with respect to the number of samples drawn from the true distribution, we ran 10 additional experiments in which we generated $n = 40$ samples from the true distribution in Step 2.

Table 3 summarizes the main statistics of the stochastic model using the best-fitted distribution and our robust models. Similar to the result in Table 2 where we have 20 samples from the true distribution, all of the robust models outperform the stochastic model with the best-fitted distribution in both the average and standard deviation of the two measures.

	Stochastic Models		Robust Models					
	True Dist	Best Fit	$\langle 3, 1 \rangle$	$\langle 3, 3 \rangle$	$\langle 3, 5 \rangle$	$\langle 5, 1 \rangle$	$\langle 5, 3 \rangle$	$\langle 5, 5 \rangle$
Cost Average	1	1.0749	1.0278	1.0223	1.0364	1.0457	1.0317	1.0304
Cost Std. Dev.	0	0.0500	0.0232	0.0262	0.0346	0.0353	0.0297	0.0297
CVaR Average	1	1.0824	1.0146	0.9904	0.9832	1.0544	1.0170	1.0092
CVaR Std. Dev.	0	0.0768	0.0415	0.0272	0.0216	0.0445	0.0165	0.0126

Table 3: Performance of Different Models for 10 Instances and 40 Samples

As expected, the average cost of all robust models and the best fit stochastic model improves when the sample size increases from 20 to 40. The robust models with bin-size 5 have a slightly greater improvement than the remaining models. For the other three statistics, we also observe improvements for the stochastic model using the best-fitted distribution as well as the robust models with bin-size 5 when the sample size is increased to 40. Again, the robust models with bin-size 5 show slightly better improvements in these statistics.

If we compare the robust models with different bin sizes, those with bin-size 3 still perform better than those with bin-size 5. However, compared with the case of 20 samples, the differences are slightly smaller for all statistics, which suggests that the robust models with bin-size 5 improve faster as the sample size increases. Similar to the experiments with 20 samples, the increase in χ^2 values also helps to improve the robustness of the models, which is measured by CVaR. The improvements in robustness as χ^2 values increase are more significant for 40-sample experiments than those with 20 samples. In addition, the increased χ^2 may also increase the cost, e.g., the average cost increases from 1.0278 to 1.0364 if we increase the χ^2 value from 1 to 5 for the robust

models with bin-size 3.

The robust model with parameters $\langle 3, 3 \rangle$ has the lowest average cost, the second lowest average CVaR and the second lowest standard deviation of the cost. Besides, its standard deviation of CVaR is less than 3%. We still consider it as the most efficient model among all the robust models and the stochastic model using the best-fitted distribution.

To summarize the numerical results, the computational experiments show that the robust models outperform the stochastic models using fitted distributions in terms of both optimality and robustness. The robust models with a lower bin size perform better than those with a larger bin size, but an increase in sample size may decrease the difference in performance caused by the choice of the bin size. In addition, a higher χ^2 value helps to increase the robustness but it may sacrifice the cost of the robust models.

5.2 Comparison with the Robust Model Based on Delage and Ye (2010)

Delage and Ye (2010) propose a robust formulation for single-stage optimization problems, which minimizes the worst-case expectation over a set of distributions defined by the estimates of the mean μ_0 and variance-covariance matrix Σ_0 of the underlying distribution. Let ξ denote the κ -dimensional vector of random parameters in the optimization problem and f_ξ the distribution of ξ . The set of distributions considered in Delage and Ye (2010) is

$$\mathcal{D} = \{f_\xi \in \mathcal{M} \mid P(\xi \in \mathcal{S}) = 1, (E[\xi] - \mu_0^T)\Sigma_0^{-1}(E[\xi] - \mu_0^T) \leq \gamma_1, E[(\xi - \mu_0)(\xi - \mu_0)^T] \preceq \gamma_2\Sigma_0\}, \quad (13)$$

where \mathcal{M} is the set of probability measures in \mathbb{R}^κ , $\mathcal{S} \subseteq \mathbb{R}^\kappa$ is a closed convex set containing the support of ξ , and $\gamma_1 \geq 0$ and $\gamma_2 \geq 1$ are two parameters controlling the size of \mathcal{D} . When μ_0 and Σ_0 correspond to the sample mean and variance, Delage and Ye (2010) identify certain conditions for γ_1 , γ_2 , and the sample size under which the probability that the distribution of ξ lies in \mathcal{D} is greater than a given confidence level.

Obviously, the set defined in (13) is also applicable to the robust inventory model discussed in Section 2. Similarly to the computational settings in the previous part, suppose that the demand distribution is stationary and we have n samples drawn from the true distribution. Let $\hat{\mu}$ and $\hat{\sigma}^2$ denote the sample mean and variance, respectively. If the support of \tilde{D}_t is contained in the set $\{0, 1, \dots, \bar{D}\}$ for any t , the set of distributions in (13) can be written as follows:

$$\mathcal{P}^{DY} = \left\{ \mathbf{P} \left| \sum_{i=0}^{\bar{D}} P_i = 1, -\hat{\sigma}_t\sqrt{\gamma_1} \leq \sum_{i=0}^{\bar{D}} iP_i - \hat{\mu} \leq \hat{\sigma}_t\sqrt{\gamma_1}, \sum_{i=0}^{\bar{D}} (i - \hat{\mu})^2 P_i \leq \gamma_2\hat{\sigma}^2, \mathbf{P} \geq \mathbf{0} \right. \right\}. \quad (14)$$

Replacing $\mathcal{P}_t(\mathbf{d}_t)$ in (5) by \mathcal{P}^{DY} , we obtain a minimax robust inventory model. Without a fixed ordering cost, the optimal policy is a base-stock policy and the base-stock levels can be computed in a fashion similar to Theorem 4.2. Next, we apply this approach to all the instances generated in Section 5.1, and use Steps 6 and 7 in Section 5.1 to evaluate its performance.

An important issue when applying the robust model based on Delage and Ye (2010) is to determine the values of the parameters γ_1 and γ_2 . We choose 10 pairs of γ_1 and γ_2 using a statistical

analysis similar to that in the computational experiments in Delage and Ye (2010). In particular, we generated 1000 distributions as described in Step 1 in Section 5.1. Each of these distributions is referred to as the j th distribution, $j = 1, \dots, 1000$. The values of γ_1 and γ_2 are selected so that set \mathcal{P}^{DY} in (14) defined by γ_1, γ_2 and the mean and variance of the j th distribution contains the mean and variance of the $(j + 1)$ th distribution with a confidence level c^{DY} .⁵ Table 4 displays the values of γ_1 and γ_2 for different c^{DY} used in the computation.

c^{DY}	99%	90%	80%	70%	60%	50%	40%	30%	20%	10%
γ_1	0.1655	0.0668	0.0426	0.0261	0.0182	0.0115	0.0059	0.0028	0.0010	0.0002
γ_2	1.3572	1.1916	1.1220	1.07450	1.0395	1.0034	0.9693	0.9315	0.8897	0.8381

Table 4: Values of γ_1 and γ_2 for the Robust Model Based on Delage and Ye (2010)

For the example discussed in Figure 1, we apply the robust model based on Delage and Ye for different γ_1 and γ_2 corresponding to the values of c^{DY} in Table 4. Table 5 shows the normalized cost with respect to the true distribution in the second row and in the third row the values are the CVaR at 5% level of the costs out of 1,000 distributions generated by permutations described in Step 7 of Section 5.1.

c^{DY}	99%	90%	80%	70%	60%	50%	40%	30%	20%	10%
Cost	1.2142	1.1505	1.1403	1.1352	1.1191	1.1206	1.1122	1.1003	1.0951	1.0868
CVaR	1.0693	1.0461	1.0423	1.0421	1.0380	1.0425	1.0385	1.0385	1.0550	1.0600

Table 5: Performance of the Robust Model Based on Delage and Ye (2010) for the Instance in Figure 1

In general, the expected cost, which measures the optimality of the model, is decreasing in the confidence level c^{DY} , which is expected since the parameters γ_1 and γ_2 , and hence the size of the distributional set, are decreasing in c^{DY} . However, the value of CVaR, which measures the robustness of the model, is not increasing in c^{DY} . The largest CVaR is obtained when $c^{DY} = 99\%$, which implies that the corresponding distributional set is too large compared with the set of distributions generated by the permutations.

Compared with Table 1, for most parameters, except for $c^{DY} = 99\%$, the robust model based on Delage and Ye (2010), has lower expected cost and CVaR than our robust model with bin-size 5. However, our model with bin-size 3 outperforms all results shown in Table 5. In particular, the performance measures for the $\langle 3, 3 \rangle$ robust model are 1.0211 and 0.9774 respectively, while the lowest expected cost for the model based on Delage and Ye (2010) is 1.0868 (when $c^{DY} = 10\%$) and the lowest CVaR is 1.0380 (when $c^{DY} = 60\%$).

The robust model based on Delage and Ye (2010) with parameters in Table 5 is also applied to the 20- and 40-sample instances in Tables 2 and 3. The statistics corresponding to these are presented in Tables 6 and 7, respectively.

First, we compare the performance of our robust model in Section 4 with the robust model based on Delage and Ye (2010) for the 20-sample instances. The average cost of the latter is from

⁵In Delage and Ye (2010), the means and variances are computed from financial market data in 30 consecutive days, and the confidence level is 99% confidence level.

c^{DY}	99%	90%	80%	70%	60%	50%	40%	30%	20%	10%
Cost Average	1.135	1.125	1.122	1.115	1.117	1.114	1.110	1.112	1.114	1.114
Cost Std. Dev.	0.049	0.052	0.050	0.047	0.045	0.047	0.045	0.053	0.061	0.064
CVaR Average	0.957	0.953	0.953	0.951	0.955	0.956	0.957	0.962	0.965	0.969
CVaR Std. Dev.	0.119	0.115	0.115	0.115	0.120	0.122	0.123	0.130	0.134	0.137

Table 6: Performance of the Robust Model Based on Delage and Ye (2010) for the 20-Sample Instances in Table 2

c^{DY}	99%	90%	80%	70%	60%	50%	40%	30%	20%	10%
Cost Average	1.144	1.130	1.126	1.118	1.115	1.109	1.106	1.103	1.100	1.100
Cost Std. Dev.	0.037	0.031	0.030	0.029	0.027	0.027	0.028	0.026	0.025	0.025
CVaR Average	1.009	1.004	1.002	1.001	1.001	1.000	1.000	1.002	1.001	1.005
CVaR Std. Dev.	0.067	0.066	0.065	0.064	0.064	0.063	0.062	0.063	0.061	0.069

Table 7: Performance of the Robust Model Based on Delage and Ye (2010) for the 40-Sample Instances in Table 3

1.110 to 1.135 while the corresponding cost in Table 2 is less than 1.073. On the other hand, the model based on Delage and Ye (2010) have lower average CVaR, which varies from 0.951 to 0.969, than our robust model, whose average CVaR is from 0.975 to 1.065. As for the standard deviations, our robust model has a lower standard deviation in CVaR, and its standard deviation in cost is comparable to the model based on Delage and Ye (2010).

We now discuss the 40-sample instances. The average cost of the model based on Delage and Ye (2010) is almost the same as the 20-sample instances, but these statistics of our robust model are on average 1.9% lower in the 20-sample instances. Moreover, our robust model has similar values of CVaR for both 20- and 40-sample instances, while the CVaR of the robust model based on Delage and Ye (2010) increases 4.6% on average when the sample size increases to 40. The two types of robust models have similar standard deviation in cost, but our robust model has a lower standard deviation in CVaR.

The numerical results indicate that our robust model outperforms with respect to the solution quality, which is mainly due to the distributional set derived from the chi-square test carrying more information about the shape of the distribution compared with the distributional set (14) defined only by the sample mean and variance. The model based on Delage and Ye (2010) is more robust, but the difference in robustness of these two models decreases as we increase the sample size from 20 to 40.

In addition, we identified that the $\langle 3, 3 \rangle$ robust model based on the chi-square test has the best performance among all of the six combinations of the parameters. For 20-sample instances, its average cost is 7.3% lower than the lowest cost for the model based on Delage and Ye (2010) (when $c^{DY} = 40\%$) and the corresponding average CVaR is only 2.3% higher than the lowest CVaR of that model (when $c^{DY} = 70\%$). For 40-sample instances, the average cost and CVaR of the $\langle 3, 3 \rangle$ robust model are lower than the lowest values of the model based on Delage and Ye (2010) (achieved when $c^{DY} = 10\%$ and $c^{DY} = 50\%$, respectively) by 7.8% and 1.0%, respectively. We conclude that the $\langle 3, 3 \rangle$ robust model achieves a better tradeoff between optimality and robustness than the model

based on Delage and Ye (2010).

6. Conclusions and Extensions

In this paper, we propose a robust stochastic model for the multi-period lot sizing problem, in which the demand distribution is unknown and the only available information is historical data. The convergence results for the chi-square test based models suggest that the solutions to the robust approach are very close to the optimal stochastic programming solutions when the sample size is sufficiently large. When the sample size is relatively small, the extensive numerical results show that the robust model still obtains a close-to-optimal solution whose performance is insensitive to the disturbances in demand distributions. This robust framework based on historical data can be extended to many more general finite-horizon dynamic programming problems, and the convergence properties can also be extended to more general problems.

Although we consider backorder models, most of our results can be extended to lost sales models. In particular, for lost-sales models with only linear procurement cost and under the same technical assumptions, the optimal policy under the robust model is a state-dependent base-stock policy.

A. Appendix: Convergence of Robust Models Based on the Chi-Square Test

We assume that the demand random variables $\tilde{D}_1, \tilde{D}_2, \dots, \tilde{D}_T$ are subject to some multivariate distributions. Although the distribution may not be known, we assume that histograms pertaining to the robust model are obtained from samples from these distributions. We study the behavior of cost-to-go functions as the number of samples increases.

Let $\tilde{\mathbf{d}}_t = [\tilde{D}_1, \dots, \tilde{D}_{t-1}]$, and let $F_t(D_t | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ denote the conditional cumulative distribution function of demand \tilde{D}_t given realized demand \mathbf{d}_t from periods 1 to $t - 1$. Assuming that the conditional distribution function is known for each t and \mathbf{d}_t , we can solve the corresponding dynamic programming problem, and obtain $\bar{V}(x_t, \mathbf{d}_t)$ for each period t ,

$$\bar{V}_t(x_t, \mathbf{d}_t) = \min_{y_t \geq x_t} \left\{ \int_{D_t} (C_t(x_t, y_t, D_t) + \theta \bar{V}_{t+1}(y_t - D_t, [\mathbf{d}_t, D_t])) dF_t(D_t | \tilde{\mathbf{d}}_t = \mathbf{d}_t) \right\},$$

where $t = 1, \dots, T$.

We investigate how accurately the value functions $V_t(x_t, \mathbf{d}_t)$ of our robust model approximate the true cost-to-go function $\bar{V}(x_t, \mathbf{d}_t)$ if histograms are based on samples.

We start by analyzing the convergence of the robust model as χ_t^2 converges to 0. Let $\hat{V}_t(x_t, \mathbf{d}_t)$ denote the cost-to-go function of the stochastic model with the distribution defined by

$$P(\tilde{D}_\tau = D_{\tau,i} | \tilde{\mathbf{d}}_\tau = \mathbf{d}_\tau) = \frac{N_{\tau,i}(\mathbf{d}_\tau)}{n_\tau(\mathbf{d}_\tau)}, \quad \tau = t, \dots, T. \quad (15)$$

Formally,

$$\hat{V}_t(x_t, \mathbf{d}_t) = \min_{y_t \geq x_t} \left\{ \sum_i \frac{N_{t,i}(\mathbf{d}_t)}{n_t(\mathbf{d}_t)} \left(C_t(x_t, y_t, D_{t,i}) + \theta \hat{V}_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right) \right\}, \quad t = 1, \dots, T.$$

Proposition A.1. *If $V_{T+1}(\cdot) = \hat{V}_{T+1}(\cdot)$, then for any x_t , \mathbf{d}_t , and t , we have*

$$\lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} V_t(x_t, \mathbf{d}_t) = \hat{V}_t(x_t, \mathbf{d}_t).$$

Proof. Please refer to the Online Supplement. □

Now suppose that for each period t , we have a sequence of samples. For any $k = 1, 2, \dots$ we have the set of m_t^k available samples for period t ,

$$\mathbf{d}_t^k = \left\{ \mathbf{d}_{t,1}^k, \dots, \mathbf{d}_{t,m_t^k}^k \right\}.$$

The samples are drawn from the distribution \tilde{D}_t conditioned on realized demand \mathbf{d}_t . Therefore, given realized demand \mathbf{d}_t from periods 1 to $t - 1$, and the k th sample set \mathbf{d}_t^k for period t , we can construct a histogram such that the total number of samples selected in the histogram is $n_t^k(\mathbf{d}_t)$, the boundaries of bins are $\left\{ D_{t,1}^k, \dots, D_{t,M_t^k}^k \right\}$, and the number of samples falling in the i th bin $[D_{t,i}^k, D_{t,i+1}^k)$ is denoted by $N_{t,i}^k(\mathbf{d}_t)$. This histogram naturally defines an empirical distribution with the conditional cumulative distribution function $F_t^k \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right)$ defined by

$$F_t^k \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right) = \frac{1}{n_t^k(\mathbf{d}_t)} \sum_{i=1}^{i^*(D_t)} N_{t,i}^k(\mathbf{d}_t),$$

where $i^* = i^*(D_t)$ is such that $D_{t,i^*}^k \leq D_t < D_{t,i^*+1}^k$.

Note that the k th set of samples \mathbf{d}_t^k , $t = 1, \dots, T$, also defines a robust model $V_t^k(x_t, \mathbf{d}_t)$ based on the just described parameters $n_t^k(\mathbf{d}_t)$ and $N_{t,i}^k(\mathbf{d}_t)$. In the remainder of this section we analyze under what conditions $V_t^k(x_t, \mathbf{d}_t)$ converge to $\bar{V}_t(x_t, \mathbf{d}_t)$, which denotes the cost-to-go function of the stochastic model with respect to true distributions. We always assume that the distribution of \tilde{D}_t has finite support $[0, D_t^{\max}]$ for any t , and $V_{T+1}(\cdot) = \hat{V}_{T+1}(\cdot) = \bar{V}_{T+1}(\cdot)$. We first study the case with general distributions, and we derive stronger results when the distributions are discrete.

A.1 General Distributions

We first show convergence under general distributions. We only need the distribution functions of samples to converge pointwise to the distribution function of the true distribution and $\chi_t^2 \rightarrow 0$.

Proposition A.2. *Suppose that for any \mathbf{d}_t and t we have*

$$\lim_{k \rightarrow \infty} F_t^k \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right) = F_t \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right)$$

for every D_t . If there is no fixed procurement cost, then

$$\lim_{k \rightarrow \infty} \lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} V_t^k(x_t, \mathbf{d}_t) = \bar{V}_t(x_t, \mathbf{d}_t).$$

Furthermore, the convergence is uniform with respect to k .

Proof. Please refer to the Online Supplement. □

In the proof, we need some concepts from measure theory and a known result. A sequence of measures μ_k converge to a measure μ *weakly* if $\int f d\mu_k \xrightarrow{k \rightarrow \infty} \int f d\mu$ for every continuous bounded function f . A sequence of measures μ_k converge to a measure μ *setwise* if $\mu_k(B) \xrightarrow{k \rightarrow \infty} \mu(B)$ for every measurable set B .

It is well known that convergence in distribution does not imply setwise convergence of the underlying probability measures. Indeed, convergence in distribution is equivalent to weak convergence. It is not difficult to see that setwise convergence implies weak convergence.

The following result can be found in Royden (1988).

Proposition A.3. *Let μ_k be a sequence of measures converging setwise to a measure μ . Let $\{f_k\}_k, \{g_k\}_k$ be two sequences of measurable functions converging pointwise to f and g respectively. Furthermore, let $|f_k| \leq g_k$ for every k and $\lim_k \int g_k d\mu_k = \int g d\mu < \infty$. Then*

$$\lim_k \int f_k d\mu_k = \int f d\mu.$$

If F_t is continuous, then the following result is obtained.

Corollary A.1. *If for any \mathbf{d}_t and t , $F_t^k(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ converge in distribution to $F_t(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ and $F_t(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ is continuous, then*

$$\lim_{k \rightarrow \infty} \lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} V_t^k(x_t, \mathbf{d}_t) = \bar{V}_t(x_t, \mathbf{d}_t).$$

Proof. Convergence in distribution implies that $F_t^k(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ converge to $F_t(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ at any point D_t where $F_t(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ is continuous. Since by assumption $F_t(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ is continuous, it follows that $F_t^k(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ converge pointwise to $F_t(\cdot | \tilde{\mathbf{d}}_t = \mathbf{d}_t)$ and thus we can apply Proposition A.2. \square

Now suppose that the demand distributions for each time period are independent, and let $F_t(D_t)$ denote the cumulative distribution function of \tilde{D}_t . Let $\{\mathbf{d}_{t,1}, \mathbf{d}_{t,2}, \dots\}$ denote a sequence of random samples drawn from the true distribution \tilde{D}_t . We can define the k th sample set for period t as $\mathbf{d}_t^k = \{\mathbf{d}_{t,1}, \dots, \mathbf{d}_{t,k}\}$. Consider the robust model independent of realized demand. The histogram for time period t is based on \mathbf{d}_t^k with the bins' boundaries being all distinct elements in this set. The corresponding empirical distribution is defined by

$$F_t^k(D_t) = \frac{1}{k} \times |\{\mathbf{d}_{t,j} : \mathbf{d}_{t,j} \leq D_t, j = 1, \dots, k\}|.$$

Let $V_t^k(x_t)$ denote the cost-to-function of the robust model defined by the histogram based on the k th sample set, and let $\bar{V}_t(x_t)$ denote the cost-to-go function corresponding to the stochastic model given distribution functions $F_t(D_t)$.

Corollary A.2. *If F_t is continuous and there is no fixed procurement cost, then*

$$\lim_{k \rightarrow \infty} \lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} V_t^k(x_t) = \bar{V}_t(x_t) \quad a.s.$$

Proof. As $k \rightarrow \infty$, the Glivenko-Cantelli theorem (see, e.g., Billingsley 1986) shows that $F_t^k(D_t)$ converges to $F_t(D_t)$ uniformly a.s. at every point D_t where $F_t(D_t)$ is continuous. The result follows immediately from Corollary A.1. \square

A.2 Discrete Distributions

Under the setting of Proposition A.2, consider the case of \tilde{D}_t being subject to a discrete distribution with finite support $\{D_{t,1}, \dots, D_{t,M_t}\} \subset [0, D_t^{\max}]$, and let $P(\tilde{D}_t = D_{t,i} \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t) = p_{t,i}(\mathbf{d}_t)$. Without loss of generality, we let this finite support be the boundaries of the bins for all the histograms associated with time period t . A result similar to Proposition A.2 is next proved for the robust stochastic model with both fixed and variable procurement cost.

Proposition A.4. *Suppose that for any \mathbf{d}_t , i and t , $N_{t,i}^k(\mathbf{d}_t)/n_t^k(\mathbf{d}_t)$ converge to $p_{t,i}(\mathbf{d}_t)$. Then with fixed and variable procurement cost, we have*

$$\lim_{k \rightarrow \infty} \lim_{\chi_t^2 \rightarrow 0, \forall \tau \geq t} V_t^k(x_t, \mathbf{d}_t) = \bar{V}_t(x_t, \mathbf{d}_t).$$

Proof. Please refer to the Online Supplement. \square

So far we assumed that χ_t^2 converges to zero. Next we establish a convergence result for any fixed χ_t^2 . We must require that the number of samples goes to infinity.

Proposition A.5. *Suppose that for any \mathbf{d}_t , i and t , $N_{t,i}^k(\mathbf{d}_t)/n_t^k(\mathbf{d}_t)$ converge to $p_{t,i}(\mathbf{d}_t)$, and that the number of samples $n_t^k(\mathbf{d}_t)$ converges to infinity. Then for any fixed χ_t^2 and with fixed and variable procurement cost we have*

$$\lim_{k \rightarrow \infty} V_t^k(x_t, \mathbf{d}_t) = \bar{V}_t(x_t, \mathbf{d}_t).$$

Proof. Please refer to the Online Supplement. \square

Now consider the setting of Corollary A.2, where the demand distributions are assumed to be independent, and the k th sample set is defined to be the first k elements in a sequence of independent random samples drawn from the true distribution. Using Proposition A.4, we can obtain a result analogous to Corollary A.2. We also establish the rate of convergence.

Corollary A.3. *With fixed and variable procurement cost, we have $\lim_{k \rightarrow \infty} V_t^k(x_t) = \bar{V}_t(x_t)$ a.s., and the rate of convergence is $O(1/\sqrt{k})$.*

Proof. Please refer to the Online Supplement. \square

Acknowledgement

The authors thank the senior editor and two anonymous referees for their comments to improve the paper. The authors were supported by NSF contract CMMI-0758069, NASA's Interplanetary Supply Chain Management and Logistics Architectures project, and a grant from SAP Research Center, Palo Alto. The last author was also supported by Hong Kong RGC GRF grant (HKU 717111E).

References

- Ahmed, S., U. Cakmak, A. Shapiro. 2007. Coherent risk measures in inventory problems. *European Journal of Operational Research* **182**(1), 226-238.
- Ben-Tal, A., A. Nemirovski. 1998. Robust convex optimization. *Mathematics of Operations Research* **23**(4), 769-805.
- Ben-Tal, A., A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming* **88**(3), 411-424.
- Bertsimas, D., M. Sim. 2003. Robust discrete optimization and network flows. *Mathematical Programming* **98**(1-3), 49-71.
- Bertsimas, D., M. Sim. 2004. The price of robustness. *Operations Research* **52**(1), 35-53.
- Bertsimas, D., A. Thiele. 2006. A robust optimization approach to inventory theory. *Operations Research* **54**(1), 150-168.
- Bienstock, D., N. Özbay. 2008. Computing robust basestock levels. *Discrete Optimization* **5**(2), 389-414.
- Billingsley, P. 1986. *Probability and Measure*, 2nd ed. John Wiley and Sons, New York, NY.
- Chen, X., D. Simchi-Levi. 2004. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: the finite horizon case. *Operations Research* **52**(6), 887-896.
- Chernoff, H., E. L. Lehmann. 1954. The use of maximum likelihood estimates in χ^2 tests for goodness-of-fit. *The Annals of Mathematical Statistics* **25**(3), 579-586.
- Delage, E., Y. Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problem. *Operations Research* **58**(3), 596-612.
- Gallego, G., I. Moon. 1993. The distribution free newsboy problem: review and extensions. *The Journal of the Operational Research Society* **44**(8), 825-834.
- Guan, Y., A. J. Miller. 2008. Polynomial-time algorithms for stochastic uncapacitated lot-sizing problems. *Operations Research* **56**(5), 1172-1183.
- Halman, N., D. Klabjan, M. Mostagir, J. Orlin, D. Simchi-levi. 2009. A fully polynomial-time approximation scheme for single-item stochastic inventory control with discrete demand. *Mathematics of Operations Research* **34**(3), 674-685.
- Huh, W. T., P. Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* **34**(1), 103-123.
- Iglehart, D. L. 1964. The dynamic inventory problem with unknown demand distribution. *Management Science* **10**(3), 429-440.

- Iyengar, G. N. 2005. Robust dynamic programming. *Mathematics of Operations Research* **30**(2), 257-280.
- Kouvelis, P., Yu. G. 1996. *Robust Discrete Optimization and Its Applications*. Kluwer Academic Publishers, Dordrecht, Germany.
- Law, A. M., W. D. Kelton. 2000. *Simulation Modeling and Analysis*, 3rd ed. McGraw-Hill, New York, NY.
- Liyanage, L. H., J. G. Shanthikumar. 2005. A practical inventory control policy using operational statistics. *Operations Research Letters* **33**(4), 341-348.
- Lobo, M., L. Vandenbergh, S. Boyd, H. Lebert. 1998. Applications of second-order cone programming. *Linear Algebra and its Applications* **284**(1-3), 193-228.
- Natarajan, K., M. Sim, J. Uichancoz. 2008. Asymmetry and ambiguity in newsvendor models. Working paper, National University of Singapore, Singapore.
- Nilim, A., L. El Ghaoui. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research* **53**(5), 780-798.
- Notzon, E. 1970. Minimax inventory and queueing models. Technical Report No. 7, Department of Operations Research, Stanford University, Stanford, CA.
- Perakis, G., G. Roels. 2008. Regret in the newsvendor model with partial information. *Operations Research* **56**(1), 188-203.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.
- Rockafellar, R. T. 1996. *Convex Analysis*. Princeton University Press, Princeton, NJ.
- Royden, H. L. 1988. *Real Analysis*, 3rd ed. Prentice Hall, Englewood Cliffs, NJ.
- Scarf, H. 1958. A min-max solution of an inventory problem. *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, Stanford, CA.
- Scarf, H. 1960. The optimality of (s, S) policies in the dynamic inventory problems. *Mathematical Methods in the Social Sciences*. Stanford University Press, Stanford, CA.
- See, C. T., M. Sim. 2010. Robust approximation to multi-period inventory management. *Operations Research* **58**(3), 583 - 594.
- Soyster, A. L. 1973. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research* **21**(5), 1154-1157.
- von Neumann, J. 1928. Zur theorie der gesellschaftsspiele. *Mathematische Annalen* **100**, 295-320.
- Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill, New York, NY.

Online Supplement

Proof of Lemma 3.1. Consider the function

$$g(\mathbf{U}, \mathbf{d}) = \max_{\mathbf{P} \in \mathcal{P}(\mathbf{d})} \mathbf{U}^T \mathbf{P}$$

Note that $f(y, \mathbf{d}) = g(\mathbf{U}(y, \mathbf{d}), \mathbf{d})$.

We first show that $g(\mathbf{U}, \mathbf{d})$ is an increasing function of \mathbf{U} for any given \mathbf{d} . Suppose that $\mathbf{U}_1 \leq \mathbf{U}_2$ and $g(\mathbf{U}_1, \mathbf{d}) = \mathbf{U}_1^T \mathbf{P}_1^*$. Since $\mathbf{P}_1^* \geq 0$,

$$g(\mathbf{U}_1, \mathbf{d}) = \mathbf{U}_1^T \mathbf{P}_1^* \leq \mathbf{U}_2^T \mathbf{P}_1^* \leq g(\mathbf{U}_2, \mathbf{d}).$$

Consider now the value of $g(\mathbf{U} + K\mathbf{e}, \mathbf{d})$, where \mathbf{e} is the vector with all entries of 1. Let \mathbf{P}^* denote the maximizer of $g(\mathbf{U} + K\mathbf{e}, \mathbf{d})$. We have

$$g(\mathbf{U} + K\mathbf{e}, \mathbf{d}) = (\mathbf{U} + K\mathbf{e})^T \mathbf{P}^* = \mathbf{U}^T \mathbf{P}^* + K\mathbf{e}^T \mathbf{P}^* \leq g(\mathbf{U}, \mathbf{d}) + K, \quad (16)$$

where the last inequality follows from $g(\mathbf{U}, \mathbf{d}) \geq \mathbf{U}^T \mathbf{P}^*$ and $\mathbf{e}^T \mathbf{P}^* = \sum_i P_i^* = 1$ as \mathbf{P}^* defines a distribution.

For any $y_1 \leq y_2$ and $\lambda \in [0, 1]$, since $U_i(y, \mathbf{d})$ is K -convex in y for any given \mathbf{d} , we have

$$U_i((1 - \lambda)y_1 + \lambda y_2, \mathbf{d}) \leq (1 - \lambda)U_i(y_1, \mathbf{d}) + \lambda U_i(y_2, \mathbf{d}) + \lambda K.$$

As $g(\mathbf{U}, \mathbf{d})$ is increasing in \mathbf{U} ,

$$g(\mathbf{U}((1 - \lambda)y_1 + \lambda y_2, \mathbf{d}), \mathbf{d}) \leq g((1 - \lambda)\mathbf{U}(y_1, \mathbf{d}) + \lambda\mathbf{U}(y_2, \mathbf{d}) + \lambda K\mathbf{e}, \mathbf{d}).$$

It is straightforward to show that $g(\mathbf{U}, \mathbf{d})$ is a convex function of \mathbf{U} , as it is the maximum of linear functions of \mathbf{U} . Therefore,

$$g((1 - \lambda)\mathbf{U}(y_1, \mathbf{d}) + \lambda\mathbf{U}(y_2, \mathbf{d}) + \lambda K\mathbf{e}, \mathbf{d}) \leq (1 - \lambda)g(\mathbf{U}(y_1, \mathbf{d}), \mathbf{d}) + \lambda g(\mathbf{U}(y_2, \mathbf{d}) + K\mathbf{e}, \mathbf{d}).$$

According to (16) we have

$$g(\mathbf{U}(y_2, \mathbf{d}) + K\mathbf{e}, \mathbf{d}) \leq g(\mathbf{U}(y_2, \mathbf{d}), \mathbf{d}) + K.$$

As a result, it follows

$$g(\mathbf{U}((1 - \lambda)y_1 + \lambda y_2, \mathbf{d}), \mathbf{d}) \leq (1 - \lambda)g(\mathbf{U}(y_1, \mathbf{d}), \mathbf{d}) + \lambda g(\mathbf{U}(y_2, \mathbf{d}), \mathbf{d}) + \lambda K,$$

and therefore $f(y, \mathbf{d}) = g(\mathbf{U}(y, \mathbf{d}), \mathbf{d})$ is a K -convex function in y . □

Proof of Proposition 4.1. The optimality equation defined in (5) is equivalent to

$$\begin{aligned} V_t(x_t, \mathbf{d}_t) &= \min_{y_t \geq x_t, \mathbf{U}_t} K\mathbb{I}(y_t - x_t) + c_t(y_t - x_t) + \max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \sum_i P_{t,i} U_{t,i} \\ \text{s.t.} \quad &U_{t,i} \geq h_t(y_t - D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \quad \text{for every } i \\ &U_{t,i} \geq b_t(y_t - D_{t,i}) + \theta V_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \quad \text{for every } i. \end{aligned} \quad (17)$$

According to Lemma 4.1, the maximization problem $\max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \sum_i P_{t,i} U_{t,i}$ is the second order cone problem and hence it is equivalent to its Lagrangian dual

$$\begin{aligned} \min_{\mathbf{p}_t, \mathbf{u}_t, \lambda_t} \quad & \mathbf{p}_t^T \mathbf{b}_t - 2 \sum_i u_{t,i} N_{t,i}(\mathbf{d}_t) + \lambda_t (n_t(\mathbf{d}_t)^2 + n_t(\mathbf{d}_t) \chi_t^2) \\ \text{s.t.} \quad & \left\| \begin{bmatrix} \mathbf{p}_t^T - U_{t,i} \mathbf{A}_{t,i} - \lambda_t \\ 2u_{t,i} \end{bmatrix} \right\| \leq \mathbf{p}_t^T - U_{t,i} \mathbf{A}_{t,i} + \lambda_t \quad \text{for every } i, \end{aligned} \quad (18)$$

where $\mathbf{A}_{t,i}$ denotes the i th row of matrix \mathbf{A}_t (see, e.g., Lobo et al. 1998).

Note that (9) is obtained by replacing the maximization problem in (17) by (18). Therefore, the proposition is equivalent to proving that problem (9) is equivalent to problem (17). Let z_1^* and z_2^* denote the optimal values of problems (17) and (9), respectively.

We first show that $z_1^* \geq z_2^*$. Let y_t^* , \mathbf{U}_t^* and \mathbf{P}_t^* denote an optimal solution for problem (17). Problem $\max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \sum_i P_{t,i} U_{t,i}$ has a finite optimal value if we set $U_{t,i}$ to $U_{t,i}^*$. Therefore, there exists an optimal solution \mathbf{p}_t^* , \mathbf{u}_t^* , and λ_t^* for its dual, problem (18), and the corresponding optimal value is $z_1^* - K\mathbb{I}(y_t^* - x_t) - c_t(y_t^* - x_t)$. Obviously y_t^* , \mathbf{U}_t^* , \mathbf{p}_t^* , \mathbf{u}_t^* , and λ_t^* is a feasible solution to (9) with the objective value z_1^* , and therefore we have $z_1^* \geq z_2^*$.

It remains to show $z_1^* \leq z_2^*$. Let y^* , \mathbf{U}^* , \mathbf{p}^* , \mathbf{u}^* and λ^* be an optimal solution for problem (9). Problem (18) with $U_{t,i} = U_{t,i}^*$ has a finite optimal value, and therefore the problem $\max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \sum_i P_{t,i} U_{t,i}^*$ has an optimal solution \mathbf{P}^* with the optimal cost $z_2^* - K\mathbb{I}(y_t^* - x_t) - c_t(y_t^* - x_t)$. Since y^* , \mathbf{U}^* , and \mathbf{P}^* give a feasible solution to problem (17) and the corresponding objective value is z_2^* , we have $z_1^* \leq z_2^*$. \square

Proof of Proposition A.1. The proposition clearly holds for $t = T + 1$. Suppose that it holds for any τ such that $\tau > t$. To simplify notation, let

$$U_{t,i}(x_t, y_t) = C_t(x_t, y_t, D_{t,i}) + \theta V_{t+1}((y_t - D_{t,i})^+, [\mathbf{d}_t, D_{t,i}]) \quad (19)$$

and therefore

$$V_t(x_t, \mathbf{d}_t) = \min_{y_t \geq x_t} \max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \left\{ \sum_i P_{t,i} U_{t,i}(x_t, y_t) \right\} \quad t = 1, \dots, T.$$

According to the definition of $\mathcal{P}_t(\mathbf{d}_t)$ in (8), we have

$$\mathcal{P}_t(\mathbf{d}_t) \subset \left\{ \mathbf{P}_t \left| \frac{(N_{t,i}(\mathbf{d}_t) - n_t(\mathbf{d}_t) P_{t,i})^2}{n_t(\mathbf{d}_t) P_{t,i}} \leq \chi_t^2 \text{ for every } i \right. \right\}.$$

Let $\underline{P}_{t,i}$ and $\overline{P}_{t,i}$ correspond to the solutions of $\frac{(N_{t,i} - n_t P_{t,i})^2}{n_t P_{t,i}} = \chi_t^2$, i.e.,

$$\begin{aligned} \underline{P}_{t,i}(\mathbf{d}_t) &= \frac{N_{t,i}(\mathbf{d}_t)}{n_t(\mathbf{d}_t)} + \frac{\chi_t^2}{2n_t(\mathbf{d}_t)} - \frac{\sqrt{4N_{t,i}(\mathbf{d}_t)\chi_t^2 + (\chi_t^2)^2}}{2n_t(\mathbf{d}_t)}, \\ \overline{P}_{t,i}(\mathbf{d}_t) &= \frac{N_{t,i}(\mathbf{d}_t)}{n_t(\mathbf{d}_t)} + \frac{\chi_t^2}{2n_t(\mathbf{d}_t)} + \frac{\sqrt{4N_{t,i}(\mathbf{d}_t)\chi_t^2 + (\chi_t^2)^2}}{2n_t(\mathbf{d}_t)}. \end{aligned}$$

It follows directly that

$$\mathcal{P}_t(\mathbf{d}_t) \subset \overline{\mathcal{P}}_t(\mathbf{d}_t) = \left\{ \mathbf{P}_t(\mathbf{d}_t) \mid \underline{P}_{t,i}(\mathbf{d}_t) \leq P_{t,i} \leq \overline{P}_{t,i}(\mathbf{d}_t) \text{ for every } i \right\}.$$

Therefore we obtain

$$\begin{aligned} \max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \left\{ \sum_i P_{t,i} U_{t,i}(x_t, y_t) \right\} &\leq \max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \left\{ \sum_i P_{t,i} U_{t,i}(x_t, y_t) \right\} \\ &= \sum_{i: U_{t,i}(x_t, y_t) \leq 0} \underline{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t) + \sum_{i: U_{t,i}(x_t, y_t) > 0} \bar{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t). \end{aligned}$$

Minimizing both sides over $\{x_t | y_t \geq x_t\}$ yields

$$\begin{aligned} V_t(x_t, \mathbf{d}_t) &= \min_{y_t \geq x_t} \max_{\mathbf{P}_t \in \mathcal{P}_t(\mathbf{d}_t)} \left\{ \sum_i P_{t,i} U_{t,i}(x_t, y_t) \right\} \\ &\leq \min_{y_t \geq x_t} \left\{ \sum_{i: U_{t,i}(x_t, y_t) \leq 0} \underline{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t) + \sum_{i: U_{t,i}(x_t, y_t) > 0} \bar{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t) \right\}. \end{aligned}$$

Let $y_t^* \geq x_t$ be a minimizer of $\hat{V}_t(x_t, \mathbf{d}_t)$. Then

$$V_t(x_t, \mathbf{d}_t) \leq \sum_{i: U_{t,i}(x_t, y_t^*) \leq 0} \underline{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t^*) + \sum_{i: U_{t,i}(x_t, y_t^*) > 0} \bar{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t^*). \quad (20)$$

Taking the limit on both sides yields

$$\lim_{\substack{\chi_\tau^2 \rightarrow 0, \\ \forall \tau \geq t}} V_t(x_t, \mathbf{d}_t) \leq \lim_{\substack{\chi_\tau^2 \rightarrow 0, \\ \forall \tau \geq t}} \left\{ \sum_{i: U_{t,i}(x_t, y_t^*) \leq 0} \underline{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t^*) + \sum_{i: U_{t,i}(x_t, y_t^*) > 0} \bar{P}_{t,i}(\mathbf{d}_t) U_{t,i}(x_t, y_t^*) \right\}.$$

Note that

$$\lim_{\chi_t^2 \rightarrow 0} \underline{P}_{t,i}(\mathbf{d}_t) = \lim_{\chi_t^2 \rightarrow 0} \bar{P}_{t,i}(\mathbf{d}_t) = \frac{N_{t,i}(\mathbf{d}_t)}{n_t(\mathbf{d}_t)},$$

and by the induction assumption

$$\lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} U_{t,i}(x_t, y_t^*) = C_t(x_t, y_t^*, D_{t,i}) + \theta \hat{V}_{t+1}((y_t^* - D_{t,i})^+, [\mathbf{d}_t, D_{t,i}]).$$

By the definition of y_t^* ,

$$\lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} V_t(x_t, \mathbf{d}_t) \leq \hat{V}_t(x_t, \mathbf{d}_t).$$

Since the distribution defined by (15) is in $\mathcal{P}_t(\mathbf{d}_t)$, it is easy to verify that $\hat{V}_t(x_t, \mathbf{d}_t) \leq V_t(x_t, \mathbf{d}_t)$. Therefore, we have

$$\lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} V_t(x_t, \mathbf{d}_t) = \hat{V}_t(x_t, \mathbf{d}_t). \quad \square$$

Proof of Proposition A.2. Let $\hat{V}_t^k(x_t, \mathbf{d}_t)$ denote the cost-to-go function of the stochastic model with respect to the empirical distribution $F_\tau^k(D_\tau | \tilde{\mathbf{d}}_\tau = \mathbf{d}_\tau)$ for any $\tau \geq t$. As shown in Proposition A.1 we have

$$\begin{aligned} \lim_{\chi_\tau^2 \rightarrow 0, \forall \tau \geq t} V_t^k(x_t, \mathbf{d}_t) &= \hat{V}_t^k(x_t, \mathbf{d}_t) \\ &= \min_{y_t \geq x_t} \left\{ \int_0^{D_t^{\max}} \left(C_t(x_t, y_t, D_t) + \theta \hat{V}_{t+1}^k(y_t - D_t, [\mathbf{d}_t, D_t]) \right) dF_t^k(D_t | \tilde{\mathbf{d}}_t = \mathbf{d}_t) \right\}. \end{aligned}$$

Therefore, it is sufficient to show that

$$\lim_{k \rightarrow \infty} \hat{V}_t^k(x_t, \mathbf{d}_t) = \bar{V}_t(x_t, \mathbf{d}_t).$$

Let us fix x_0 . Then under an optimal policy the inventory is always within

$$\left[x_0 - \sum_{\tau=1}^T D_\tau^{\max}, x_0 + \sum_{\tau=1}^T D_\tau^{\max} \right]. \quad (21)$$

Therefore we can assume that y_t and $x_t = y_t - \tilde{D}_t$ are always within this range for any t . It is easy to show by induction that $|\hat{V}_t^k(x_t, \mathbf{d}_t)| \leq M(x_0) < \infty$, where $M(x_0)$ is a constant depending only on x_0 .

Note that as $V_{T+1}(\cdot) = \hat{V}_{T+1}(\cdot) = \bar{V}_{T+1}(\cdot)$, the proposition holds for period $T+1$. Suppose that for any $\tau > t$, $\hat{V}_\tau^k(x_\tau, \mathbf{d}_\tau) \rightarrow \bar{V}_\tau(x_\tau, \mathbf{d}_\tau)$ pointwise.

Let

$$\begin{aligned} f_t^k(x_t, y_t, \mathbf{d}_t) &= \int_0^{D_t^{\max}} \left(C_t(x_t, y_t, D_t) + \theta \hat{V}_{t+1}^k(y_t - D_t, [\mathbf{d}_t, D_t]) \right) dF_t^k \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right), \\ f_t(x_t, y_t, \mathbf{d}_t) &= \int_0^{D_t^{\max}} \left(C_t(x_t, y_t, D_t) + \theta \bar{V}_{t+1}(y_t - D_t, [\mathbf{d}_t, D_t]) \right) dF_t \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right). \end{aligned} \quad (22)$$

Note that $\hat{V}_t^k(x_t, \mathbf{d}_t) = \min_{y_t \geq x_t} f_t^k(x_t, y_t, \mathbf{d}_t)$ and $\bar{V}_t(x_t, \mathbf{d}_t) = \min_{y_t \geq x_t} f_t(x_t, y_t, \mathbf{d}_t)$.

Let $\mu_{t,k}^{\mathbf{d}_t}$ be the Lebesgue-Stieltjes measure based on $F_t^k \left(\cdot \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right)$,⁶ and we define similarly $\mu_t^{\mathbf{d}_t}$ with respect to $F_t \left(\cdot \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right)$. By assumption, F_t^k converge pointwise to F_t at any point. It is now easy to see that as a result $\mu_{t,k}^{\mathbf{d}_t}$ converge setwise to $\mu_t^{\mathbf{d}_t}$.

Let now $g_{t,k}(D_t) = M(x_0)$, i.e., a sequence of constant functions, and $f_{t,k}^{\mathbf{d}_t, y_t}(D_t) = \hat{V}_{t+1}^k(y_t - D_t, [\mathbf{d}_t, D_t])$. By definition we have $|f_{t,k}^{\mathbf{d}_t, y_t}(D_t)| \leq g_{t,k}(D_t)$. Let also $g_t(D_t) = M(x_0)$. Clearly $g_{t,k}$ converge pointwise to g_t and by the induction assumption $f_{t,k}^{\mathbf{d}_t, y_t}(D_t)$ converge pointwise to $f_t^{\mathbf{d}_t, y_t}(D_t)$ defined by $f_t^{\mathbf{d}_t, y_t}(D_t) = \bar{V}_{t+1}(y_t - D_t, [\mathbf{d}_t, D_t])$.

Furthermore, $\int g_{t,k} d\mu_{t,k}^{\mathbf{d}_t} = M(x_0) = \int g_t d\mu_t^{\mathbf{d}_t}$. Thus we can apply Proposition A.3, which implies

$$\begin{aligned} & \lim_{k \rightarrow \infty} \int_0^{D_t^{\max}} \hat{V}_{t+1}^k(y_t - D_t, [\mathbf{d}_t, D_t]) dF_t^k \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right) \\ &= \int_0^{D_t^{\max}} \bar{V}_{t+1}(y_t - D_t, [\mathbf{d}_t, D_t]) dF_t \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right). \end{aligned}$$

Note that this holds at every y_t and \mathbf{d}_t .

Since setwise convergence implies weak convergence and since $C_t(x_t, y_t, D_t)$ is continuous and bounded, by weak convergence we obtain

$$\lim_{k \rightarrow \infty} \int_0^{D_t^{\max}} C_t(x_t, y_t, D_t) dF_t^k \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right) = \int_0^{D_t^{\max}} C_t(x_t, y_t, D_t) dF_t \left(D_t \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t \right).$$

⁶ $\mu_{t,k}^{\mathbf{d}_t}([a, b]) = F_t^k(b \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t) - F_t^k(a \mid \tilde{\mathbf{d}}_t = \mathbf{d}_t)$, and then $\mu_{t,k}^{\mathbf{d}_t}$ is extended by the Riesz representation theorem.

Therefore, we have

$$\lim_{k \rightarrow \infty} f_t^k(x_t, y_t, \mathbf{d}_t) = f_t(x_t, y_t, \mathbf{d}_t).$$

Note that if finite convex functions $f_k(\mathbf{x}) \rightarrow f(\mathbf{x})$ pointwise, then $f_k(\mathbf{x}) \rightarrow f(\mathbf{x})$ uniformly on each compact subset of the domain (see, e.g., Rockafellar 1996). Since there are no fixed costs, both $f_t^k(x_t, y_t, \mathbf{d}_t)$ and $f_t(x_t, y_t, \mathbf{d}_t)$ are convex in both x_t and y_t . Therefore, given \mathbf{d}_t , $f_t^k(x_t, y_t, \mathbf{d}_t) \rightarrow f_t(x_t, y_t, \mathbf{d}_t)$ pointwise implies that $f_t^k(x_t, y_t, \mathbf{d}_t) \rightarrow f_t(x_t, y_t, \mathbf{d}_t)$ uniformly.

Let $y_t^k(x_t, \mathbf{d}_t)$ and $y_t^*(x_t, \mathbf{d}_t)$ denote the minimizers of $\hat{V}_t^k(x_t, \mathbf{d}_t)$ and $\bar{V}_t(x_t, \mathbf{d}_t)$, respectively. Clearly, $\hat{V}_t^k(x_t, \mathbf{d}_t) = f_t^k(x_t, y_t^k(x_t, \mathbf{d}_t), \mathbf{d}_t)$ and $\bar{V}_t(x_t, \mathbf{d}_t) = f_t(x_t, y_t^*(x_t, \mathbf{d}_t), \mathbf{d}_t)$.

According to uniform convergence, for any $\epsilon > 0$, there exists a positive integer K such that

$$\left| f_t^k(x_t, y_t, \mathbf{d}_t) - f_t(x_t, y_t, \mathbf{d}_t) \right| < \epsilon$$

for any x_t, y_t , and $k > K$. Therefore,

$$f_t(x_t, y_t^k(x_t, \mathbf{d}_t), \mathbf{d}_t) - \epsilon < f_t^k(x_t, y_t^k(x_t, \mathbf{d}_t), \mathbf{d}_t) = \hat{V}_t^k(x_t, \mathbf{d}_t).$$

Note that $f_t(x_t, y_t^k(x_t, \mathbf{d}_t), \mathbf{d}_t) \geq f_t(x_t, y_t^*(x_t, \mathbf{d}_t), \mathbf{d}_t) = \bar{V}_t(x_t, \mathbf{d}_t)$ and therefore we have

$$\bar{V}_t(x_t, \mathbf{d}_t) - \epsilon = f_t(x_t, y_t^k(x_t, \mathbf{d}_t), \mathbf{d}_t) - \epsilon < \hat{V}_t^k(x_t, \mathbf{d}_t).$$

Also note that

$$f_t^k(x_t, y_t^*(x_t, \mathbf{d}_t), \mathbf{d}_t) \geq f_t^k(x_t, y_t^k(x_t, \mathbf{d}_t), \mathbf{d}_t) = \hat{V}_t^k(x_t, \mathbf{d}_t)$$

and

$$f_t^k(x_t, y_t^*(x_t, \mathbf{d}_t), \mathbf{d}_t) < f_t(x_t, y_t^*(x_t, \mathbf{d}_t), \mathbf{d}_t) + \epsilon = \bar{V}_t(x_t, \mathbf{d}_t) + \epsilon.$$

Thus

$$\hat{V}_t^k(x_t, \mathbf{d}_t) < \bar{V}_t(x_t, \mathbf{d}_t) + \epsilon.$$

As a result, for any x_t and $k > K$, we have

$$\left| \hat{V}_t^k(x_t, \mathbf{d}_t) - \bar{V}_t(x_t, \mathbf{d}_t) \right| < \epsilon,$$

i.e., $\hat{V}_t^k(x_t, \mathbf{d}_t) \rightarrow \bar{V}_t(x_t, \mathbf{d}_t)$ uniformly for any given \mathbf{d}_t , which completes the induction step. \square

Proof of Proposition A.4. Consider $f_t^k(x_t, y_t, \mathbf{d}_t)$ and $f_t(x_t, y_t, \mathbf{d}_t)$ defined in (22). Following the proof of Proposition A.2, it is sufficient to show that $f_t^k(x_t, y_t, \mathbf{d}_t) \rightarrow f_t(x_t, y_t, \mathbf{d}_t)$ uniformly for any fixed \mathbf{d}_t , under the induction assumption that $\hat{V}_\tau^k(x_\tau, \mathbf{d}_\tau) \rightarrow \bar{V}_\tau(x_\tau, \mathbf{d}_\tau)$ uniformly for any given \mathbf{d}_τ and $\tau > t$.

As $k \rightarrow \infty$, $N_{t,i}^k(\mathbf{d}_t)/n_t^k(\mathbf{d}_t) \rightarrow p_{t,i}(\mathbf{d}_t)$ uniformly with respect to i (note that there are only finitely many i 's). That is, for any $\epsilon > 0$, there exists a positive integer K_1 such that

$$\left| \frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} - p_{t,i}(\mathbf{d}_t) \right| < \epsilon$$

for any i and $k > K_1$.

The induction assumption implies that for any $\epsilon > 0$, there exists a positive integer K_2 such that

$$\left| \hat{V}_{t+1}^k(x_{t+1}, \mathbf{d}_{t+1}) - \bar{V}_{t+1}(x_{t+1}, \mathbf{d}_{t+1}) \right| < \epsilon$$

for any x_{t+1} and $k > K_2$.

Consider $k > \max\{K_1, K_2\}$. Given \mathbf{d}_t , for any x_t and y_t we have

$$\begin{aligned} & \left| f_t^k(x_t, y_t, \mathbf{d}_t) - f_t(x_t, y_t, \mathbf{d}_t) \right| \\ &= \left| \sum_{i=1}^{M_t} \frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} \left(h_t(y_t - D_{t,i})^+ + b_t(y_t - D_{t,i})^- + \theta \hat{V}_{t+1}^k(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right) \right. \\ & \quad \left. - \sum_{i=1}^{M_t} p_{t,i}(\mathbf{d}_t) \left(h_t(y_t - D_{t,i})^+ + b_t(y_t - D_{t,i})^- + \theta \bar{V}_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right) \right| \\ &\leq \left| \sum_{i=1}^{M_t} \left(h_t(y_t - D_{t,i})^+ + b_t(y_t - D_{t,i})^- \right) \left(\frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} - p_{t,i}(\mathbf{d}_t) \right) \right| \\ & \quad + \theta \left| \sum_{i=1}^{M_t} \left(\frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} \hat{V}_{t+1}^k(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) - p_{t,i}(\mathbf{d}_t) \bar{V}_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right) \right|. \end{aligned}$$

According to (21), $|h_t(y_t - D_{t,i})^+ + b_t(y_t - D_{t,i})^-| \leq M'(x_0) < \infty$ where $M'(x_0)$ is a constant depending on the initial net inventory x_0 . Also note that $\left| \frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} - p_{t,i}(\mathbf{d}_t) \right| < \epsilon$, and hence

$$\left| \sum_{i=1}^{M_t} \left(h_t(y_t - D_{t,i})^+ + b_t(y_t - D_{t,i})^- \right) \left(\frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} - p_{t,i}(\mathbf{d}_t) \right) \right| < \sum_{i=1}^{M_t} M'(x_0) \epsilon = M_t M'(x_0) \epsilon.$$

Since y_t and $D_{t,i}$ are bounded again by (21), $|\hat{V}_{t+1}^k(x_{t+1}, \mathbf{d}_{t+1})| \leq M(x_0) < \infty$. Also note that $\sum_{i=1}^{M_t} p_{t,i}(\mathbf{d}_t) = 1$, $|\hat{V}_{t+1}^k(x_{t+1}, \mathbf{d}_{t+1}) - \bar{V}_{t+1}(x_{t+1}, \mathbf{d}_{t+1})| < \epsilon$, and $\left| \frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} - p_{t,i}(\mathbf{d}_t) \right| < \epsilon$. We obtain

$$\begin{aligned} & \left| \sum_{i=1}^{M_t} \left(\frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} \hat{V}_{t+1}^k(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) - p_{t,i}(\mathbf{d}_t) \bar{V}_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right) \right| \\ &\leq \sum_{i=1}^{M_t} \left| \hat{V}_{t+1}^k(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right| \left| \frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} - p_{t,i}(\mathbf{d}_t) \right| \\ & \quad + \sum_{i=1}^{M_t} p_{t,i}(\mathbf{d}_t) \left| \hat{V}_{t+1}^k(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) - \bar{V}_{t+1}(y_t - D_{t,i}, [\mathbf{d}_t, D_{t,i}]) \right| \\ &< \sum_{i=1}^{M_t} M(x_0) \epsilon + \epsilon = (M_t M(x_0) + 1) \epsilon. \end{aligned}$$

As a result,

$$\left| f_t^k(x_t, y_t, \mathbf{d}_t) - f_t(x_t, y_t, \mathbf{d}_t) \right| < M_t M'(x_0) \epsilon + \theta (M_t M(x_0) + 1) \epsilon,$$

and hence $f_t^k(x_t, y_t, \mathbf{d}_t)$ converge uniformly to $f_t(x_t, y_t, \mathbf{d}_t)$. \square

Proof of Proposition A.5. Since $V_{T+1}^k(\cdot) = \bar{V}_{T+1}(\cdot)$, the proposition holds for period $T+1$. Suppose that for any $\tau > t$, $\hat{V}^k(x_\tau, \mathbf{d}_\tau) \rightarrow \bar{V}_\tau(x_\tau, \mathbf{d}_\tau)$ uniformly for any fixed \mathbf{d}_τ . Consider period t .

According to the definition of the distribution set $\mathcal{P}_t^k(\mathbf{d}_t)$ for $V_t^k(x_t, \mathbf{d}_t)$, similar to the proof of Proposition A.1, we have

$$\mathcal{P}_t^k(\mathbf{d}_t) \subset \bar{\mathcal{P}}_t^k(\mathbf{d}_t) = \left\{ \mathbf{P}_t \mid \underline{P}_{t,i}^k(\mathbf{d}_t) \leq P_{t,i} \leq \bar{P}_{t,i}^k(\mathbf{d}_t) \text{ for every } i \right\}$$

where

$$\begin{aligned} \underline{P}_{t,i}^k(\mathbf{d}_t) &= \frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} + \frac{\chi_t^2}{2n_t^k(\mathbf{d}_t)} - \frac{\sqrt{4N_{t,i}^k(\mathbf{d}_t)\chi_t^2 + (\chi_t^2)^2}}{2n_t^k(\mathbf{d}_t)}, \\ \bar{P}_{t,i}^k(\mathbf{d}_t) &= \frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} + \frac{\chi_t^2}{2n_t^k(\mathbf{d}_t)} + \frac{\sqrt{4N_{t,i}^k(\mathbf{d}_t)\chi_t^2 + (\chi_t^2)^2}}{2n_t^k(\mathbf{d}_t)}. \end{aligned} \quad (23)$$

Consider $\hat{V}_t^k(x_t, \mathbf{d}_t)$ as defined in Propositions A.2 and A.4, which denotes the cost-to-go function of the stochastic model under the empirical distribution. Note that $\hat{V}_t^k(x_t, \mathbf{d}_t) \leq V_t^k(x_t, \mathbf{d}_t)$, and hence

$$\lim_{k \rightarrow \infty} \hat{V}_t^k(x_t, \mathbf{d}_t) \leq \lim_{k \rightarrow \infty} V_t^k(x_t, \mathbf{d}_t).$$

Proposition A.4 shows that $\lim_{k \rightarrow \infty} \hat{V}_t^k(x_t, \mathbf{d}_t) = \bar{V}_t(x_t, \mathbf{d}_t)$ whenever $\frac{N_{t,i}^k(\mathbf{d}_t)}{n_t^k(\mathbf{d}_t)} \rightarrow p_{t,i}(\mathbf{d}_t)$, and we obtain

$$\bar{V}_t(x_t, \mathbf{d}_t) \leq \lim_{k \rightarrow \infty} V_t^k(x_t, \mathbf{d}_t).$$

Also note that $N_{t,i}^k(\mathbf{d}_t)/n_t^k(\mathbf{d}_t) \rightarrow p_{t,i}(\mathbf{d}_t)$ and $\frac{\chi_t^2}{2n_t^k(\mathbf{d}_t)} \pm \frac{\sqrt{4N_{t,i}^k(\mathbf{d}_t)\chi_t^2 + (\chi_t^2)^2}}{2n_t^k(\mathbf{d}_t)} \rightarrow 0$ as $k \rightarrow \infty$. Therefore,

$$\lim_{k \rightarrow \infty} \underline{P}_{t,i}^k(\mathbf{d}_t) = \lim_{k \rightarrow \infty} \bar{P}_{t,i}^k(\mathbf{d}_t) = p_{t,i}(\mathbf{d}_t).$$

Following the same argument as in the proof of Proposition (A.1), it is easy to see that

$$\lim_{k \rightarrow \infty} V_t^k(x_t, \mathbf{d}_t) \leq \bar{V}_t(x_t, \mathbf{d}_t),$$

which completes the proof. \square

Proof of Corollary A.3. The convergence follows from the Glivenko-Cantelli theorem (see, e.g., Billingsley 1986) and Proposition A.5.

Since $V_{T+1}^k(\cdot) = \bar{V}_{T+1}(\cdot)$, the rate of convergence holds for time period $T+1$. Suppose that it holds for any time period $\tau > t$.

Consider the set of distributions \mathcal{P}_t^k defined for the robust model $V_t^k(x_t)$. The definition in (8) shows that

$$\mathcal{P}_t^k \supset \left\{ \mathbf{P}_t \mid \mathbf{A}_t \mathbf{P}_t = \mathbf{b}_t, \frac{(N_{t,i}^k - kP_{t,i})^2}{kP_{t,i}} \leq \frac{\chi_t^2}{M_t} \forall i, \mathbf{P}_t \geq \mathbf{0} \right\}.$$

The inequality $\frac{(N_{t,i}^k - kP_{t,i})^2}{kP_{t,i}} \leq \frac{\chi_t^2}{M_t}$ is equivalent to

$$P_{t,i} \in \left[\frac{N_{t,i}^k}{k} + \frac{\chi_t^2/M_t}{2k} - \frac{\sqrt{4N_{t,i}^k\chi_t^2/M_t + (\chi_t^2/M_t)^2}}{2k}, \frac{N_{t,i}^k}{k} + \frac{\chi_t^2/M_t}{2k} + \frac{\sqrt{4N_{t,i}^k\chi_t^2/M_t + (\chi_t^2/M_t)^2}}{2k} \right].$$

Therefore, the rate at which \mathcal{P}_t^k shrinks to the single point $\mathbf{P}_t = \{P_{t,1} = N_{t,1}^k/k, \dots, P_{t,M_t} = N_{t,M_t}^k/k\}$ is $O(1/\sqrt{k})$. According to the law of large numbers, $N_{t,i}^k/k$ converge to $p_{t,i} = P(\tilde{D}_t = D_{t,i})$ exponentially (see, e.g., Billingsley 1986). As a result, for sufficiently large k , \mathcal{P}_t^k contains vector \mathbf{p}_t a.s., and hence $\bar{V}_t(x_t) \leq V_t^k(x_t)$ a.s.

As shown in (20),

$$\bar{V}_t(x_t) \leq V_t^k(x_t) \leq \sum_{i:U_{t,i}^k(x_t, y_t^*) \leq 0} \underline{P}_{t,i}^k U_{t,i}^k(x_t, y_t^*) + \sum_{i:U_{t,i}^k(x_t, y_t^*) > 0} \bar{P}_{t,i}^k U_{t,i}^k(x_t, y_t^*) \quad \text{a.s.},$$

where $U_{t,i}^k(x_t, y_t)$ is defined in the same way as (19), and y_t^* denotes an optimal solution to $\bar{V}_t(x_t)$.

Note that

$$\lim_{k \rightarrow \infty} U_{t,i}^k(x_t, y_t^*) = C_t(x_t, y_t^*, D_{t,i}) + \theta \bar{V}_{t+1}((y_t^* - D_{t,i})^+).$$

The rate of convergence of $U_{t,i}^k(x_t, y_t^*)$ is determined by the convergence rate of $\bar{V}_{t+1}^k(\cdot)$, and hence it is in the order of $O(1/\sqrt{k})$.

According to the definition of $\underline{P}_{t,i}^k$ and $\bar{P}_{t,i}^k$ in (23), both $\underline{P}_{t,i}^k$ and $\bar{P}_{t,i}^k$ converge to $p_{t,i}$ at the convergence rate of $O(1/\sqrt{k})$, since $N_{t,i}^k/k$ converges to $p_{t,i}$ exponentially. Finally, note that

$$\bar{V}_t(x_t) = \lim_{k \rightarrow \infty} \left\{ \sum_{i:U_{t,i}^k(x_t, y_t^*) \leq 0} \underline{P}_{t,i}^k U_{t,i}^k(x_t, y_t^*) + \sum_{i:U_{t,i}^k(x_t, y_t^*) > 0} \bar{P}_{t,i}^k U_{t,i}^k(x_t, y_t^*) \right\}.$$

We conclude that the rate of convergence of $V_t^k(x_t)$ is in the order of $O(1/\sqrt{k})$. \square