

Closed-Form Solutions for Robust Inventory Management

Hamed Mamani,^a Shima Nassiri,^a Michael R. Wagner^a

^a Foster School of Business, University of Washington, Seattle, Washington 98195

Contact: hmamani@uw.edu (HM); shiman@uw.edu (SN); mrwagner@uw.edu (MRW)

Received: June 11, 2014

Revised: October 7, 2014; July 14, 2015;
October 13, 2015

Accepted: October 27, 2015

Published Online in Articles in Advance:
April 29, 2016

<https://doi.org/10.1287/mnsc.2015.2391>

Copyright: © 2016 INFORMS

Abstract. We propose and analyze robust optimization models of an inventory management problem, where cumulative purchase, inventory, and shortage costs over n periods are minimized for correlated nonidentically distributed demand. We assume that the means and covariance matrix of stochastic demand are known; the distributions are not needed. We derive *closed-form* ordering quantities for both symmetric and asymmetric uncertainty sets, under capacitated inventory constraints, in both static and dynamic settings. The behaviors of our robust strategies differ qualitatively depending on the symmetry of the uncertainty set. For instance, considering our simplest static problem, (1) if the uncertainty set is symmetric, then there is positive ordering in all periods, whereas (2) if the set is asymmetric, then there is a set of periods in the middle of the planning horizon with zero orders. We also link the symmetry of the uncertainty set to the symmetry of the demand distribution. Finally, we present encouraging computational results where our solution compares favorably to previously studied, more complex robust solutions.

History: Accepted by Yinyu Ye, optimization.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2015.2391>.

Keywords: inventory management • robust optimization • dynamic optimization • central limit theorem

1. Introduction

Intelligent inventory management is a highly important area of management science that has received significant attention for decades. However, even with the most advanced techniques, firms still face inventory problems. For example, during December 2006, Nintendo faced widespread shortages of its Wii video game system in Europe, as reported by *BBC News* (2006). Similarly, during January 2008, Microsoft faced shortages of its Xbox 360 system, as seen in the *Seattle Post-Intelligencer* (2008). Inventory surpluses and write-offs have also been costly. For example, Research in Motion wrote off \$485 million in unsold or discounted PlayBook tablets in 2011; see Connors and Cummins (2011) for further details. Similarly, Microsoft wrote off \$900 million in Surface tablet inventory; see Covert (2013). Inventory management models also serve as building blocks for more sophisticated material requirements planning (MRP), enterprise resource planning, and supply chain management systems; see Chap. 3 of Hopp and Spearman (2011) for further details.

In this paper we analyze a classic n -period inventory management problem for a single product with no fixed costs and backordering allowed. Demand is stochastic, correlated, and not identically distributed; we assume that the means and covariance matrix are available but not the distributions. We apply recent advances in robust optimization by utilizing uncertainty sets that are motivated by the central limit theorem (CLT). These sets, pioneered by Bandi and

Bertsimas (2012), provide an intuitive link between robust and stochastic optimization, which leads to tractable robust counterparts of stochastic optimization problems that suffer from the curse of dimensionality. These achievements strongly suggest that robust optimization should be the methodology of choice for decision making under uncertainty for a variety of difficult problems.

Our basic model is a static one, where all order quantities must be determined at time 0. For this problem we derive *closed-form* robust order quantities that require only knowledge of the means and covariance matrix of demand, and not the underlying distribution; we do so for both symmetric and asymmetric uncertainty sets. Furthermore, the simple nature of our results allows a decision maker to better understand the inventory management strategy he or she implements, in contrast to the computational solution to an optimization problem. We then extend our model in multiple directions. We first introduce capacitated inventories, and we again derive closed-form order quantities. We show that if the capacity is within a threshold value, the order quantities are reduced for an intermediate range of periods, and the remaining quantities are unchanged from the uncapacitated case; if the capacity is above the threshold, the uncapacitated results apply for all periods. We then study dynamic decision making, which allows the decision in each period to depend on the currently observed inventory position, under capacitated inventories; in computational results, our dynamic strategy performs

favorably compared with more complex robust optimization models presented in the literature.

A large portion of the existing literature on inventory management requires uncertain demand to be characterized by a stochastic distribution. However, as has been noted in the literature, solutions are highly sensitive to the distribution itself, and the application of an incorrect distribution can have a damaging impact; for example, see Isaacs (1963) for the implications of basing decisions on inaccurate probability distributions. In practice, choosing an appropriate distribution is not a trivial matter. In many cases, historical data are used to derive a distribution, which is premised on (1) history repeating itself and (2) the data having a high quality; both of these assumptions are suspect in realistic situations. Indeed, changing market trends and consumer behaviors can render past behavior, and past data, useless. When data are unavailable (e.g., fashion items, new product introductions), an alternative choice is to select a well-known theoretical distribution, whose match with reality is not always vetted. Alternatively, although it is possible to use management's expert opinion to create forecasts, these distributions are unlikely to be of the high quality necessary for quantitative models. Even if an accurate distribution is known, the resulting stochastic optimization problem might be intractable (i.e., the curse of dimensionality). Therefore, there is motivation to study inventory management problems with uncertain demand for which probabilistic distributions are not known. We next survey the relevant literature on alternative characterizations of uncertain demand.

1.1. Literature Review

One of the more popular approaches to study inventory management problems without probabilistic distributions is to assume that the mean and variance of demand are known and apply a worst-case analysis (i.e., a min-max approach) over all probabilistic distributions that have the given mean and variance. Scarf (1958) derives an ordering rule for the newsvendor problem under this scenario. Gallego and Moon (1993) gives a new compact proof of Scarf's ordering rule and also investigates the recourse case, where a secondary procurement decision is allowed. Considering the continuous and periodic review inventory models, Moon and Gallego (1994) studies the case where the distribution of the lead time is unknown, but the mean and variance of the lead time are given. Returning to the newsvendor model, Perakis and Roels (2008) studies the case where a demand distribution is unavailable but other information is available to partially characterize demand (as before, mean and variance, but also symmetry and/or unimodality of the distribution). Note that these models utilize a min-max approach to differing degrees, which implies that

the models result in decisions that are risk averse, in contrast to the risk neutrality of stochastic models that optimize expected values. Other examples of this min-max style of research include Ehrhardt (1979) and Natarajan et al. (2008). For other examples of risk-averse decision making in inventory management, see Lau (1980), Eeckhoudt et al. (1995), Chen et al. (2007b), Van Mieghem (2007), and the references therein.

Robust optimization, where uncertain parameters, rather than distributions, vary over a deterministic set, is another approach that can handle uncertain demand without a stochastic distribution. Generally speaking, robust optimization can be overly conservative, as a worst-case solution is derived. However, Bertsimas and Sim (2004) introduce the use of "budgets of uncertainty" to reduce the conservatism of robust models. Bertsimas and Thiele (2006) design a robust optimization model of inventory management that utilizes the budgets of uncertainty, resulting in excellent performance, but it comes at the cost of rather heavy optimization machinery, which can handle fixed costs, capacitated orders and inventory, and network topologies. Our paper is motivated by Bertsimas and Thiele (2006), and we study a similar and simpler problem (zero fixed costs and a single station), except that we use uncertainty sets whose structure is motivated by the CLT. Notably, we derive *closed-form* solutions under both symmetric and asymmetric uncertainty sets for correlated nonidentically distributed demand, which is only possible in Bertsimas and Thiele (2006) under a symmetric set with independent and identically distributed (i.i.d.) demand. Chen et al. (2007a) study how to create generic robust uncertainty sets allowing for asymmetry, resulting in a second-order cone counterpart (which does not generally lead to simple solutions). Bienstock and Özbay (2008) generalize Bertsimas and Thiele (2006) in multiple directions (period-dependent costs, more general uncertainty sets) and also analyze data-driven robust models, focusing on an algorithmic perspective to inventory management. See and Sim (2010) analyze a robust inventory management problem under a "factor-based" model of uncertainty, which in that case is also equivalent to a second-order cone program. Wagner (2010) studies a similar model to that considered in this paper, except that demand is only known to be nonnegative and revenues are incorporated to form a profit objective. Wagner (2011) analyzes a simplification without revenues.

To the best of our knowledge, the idea of combining robust optimization with the limit theorems of probability has its origins in the paper by Bertsimas et al. (2011), which analyzes queuing networks with a robust uncertainty set motivated by the probabilistic law of the iterated logarithm (Chung 2001). This paper is based on the Ph.D. thesis of Rikun (2011),

who also analyze robust inventory management with a simple CLT-inspired uncertainty set, which is solved using a simulated-annealing algorithm (i.e., no closed-form solutions are available). The CLT structure of our uncertainty sets is primarily motivated by Bandi and Bertsimas (2012), which provides an in-depth analysis of the combination of robust optimization with the limit theorems of probability. This connection intimately links robust and stochastic optimization in an intuitive manner. Taking this new approach, it is possible to formulate tractable robust counterparts of difficult stochastic optimization problems that suffer from the “curse of dimensionality.” These ideas have been applied to information theory (Bandi and Bertsimas 2011), option pricing (Bandi and Bertsimas 2014b), queueing theory (Bandi et al. 2012, 2015), and auction design (Bandi and Bertsimas 2014a). We continue the preliminary work of Rikun (2011), under a different fundamental model, providing an in-depth analysis of these ideas applied to inventory management.

We next discuss dynamic robust optimization, where some variable values can be set after uncertain parameters become known, which is useful in multiperiod decision making. Ben-Tal et al. (2004) introduce the adjustable robust optimization problem, which is shown to be NP-hard in general; this negative result motivates the analysis of an affinely adjustable robust counterpart as an approximation to this difficult robust problem, where the optimization is performed over a subset of (suboptimal) affine policies. These authors compare their approach, in a linear inventory model, with dynamic programming (DP) and conclude that if the curse of dimensionality is present, the robust approach is superior to the DP one, because of its computational advantage. Ben-Tal et al. (2005) utilize an affinely adjustable robust model to study a retailer-supplier flexible commitment problem, which analyzes inventory management in a supply chain. Bertsimas et al. (2010) prove the optimality of affine policies for a general class of multistage robust optimization models where independent random disturbances are constrained to lie in intervals. By contrast, our uncertainty sets have coupling constraints over multiple periods, which allow us to capture correlated disturbances (i.e., demands); in Bertsimas et al. (2010), the authors provide examples where coupling constraints render the affine policies suboptimal. Iancu et al. (2013) continue the study of affine policies, further characterizing the problem structures where affine policies are optimal in dynamic robust optimization. Chen et al. (2008) utilize second-order cone approximations to improve upon the linear decision rules in a generic multiperiod problem. Georghiou et al. (2015) apply a lifting technique to the primal and dual linear decision rule approximations (which usually provide loose bounds) and prove that the bounds that they introduce are tighter

than those derived from a linear decision rule. The authors numerically compare these bounds across different approximations in terms of their optimality gaps and running times. They admit that these approximations can be difficult to solve for different sets of parameters in inventory control problems. Bertsimas and Georghiou (2015) review the limitations of the decision rules used in the approximations above, because of their a priori design and not incorporating adaptive decisions. They propose a methodology for the optimal design of such decisions through a mixed-integer program. We provide our own closed-form solutions for both symmetric and asymmetric sets under a rolling horizon framework, whose simple intuitive structure is appealing from a decision maker’s viewpoint; furthermore, our strategy exhibits encouraging computational performance.

Gorissen and Hertog (2013) argue that the approach in Bertsimas and Thiele (2006) is conservative because it considers the worst-case scenario for each individual constraint as opposed to a global worst-case scenario. They propose a less conservative robust reformulation for optimization problems containing sums of maxima of linear functions. They compare various solution methodologies and propose an algorithm based on cutting-plane methods. Their numerical results show that affinely adjustable robust counterpart reformulations of inventory problems do not significantly improve on the approach provided in Bertsimas and Thiele (2006). Although more complex cutting-plane methods improve on costs, this comes at a relatively significant increase in computational time. They illustrate that after splitting up the constraint sums, the initial conservative approach can achieve similar costs as the cutting-plane-based algorithms without a significant increase in computational time.

1.2. Contributions

We study new robust variants of a classic single-echelon n -period inventory management problem for a single product, with purchasing, holding, and shortage costs. Our results make the following contributions to the operations management literature.

- We introduce a generic uncertainty set based on bounds on partial sums of demand as well as bounds on individual demands. We derive a closed-form *recursion* for the robust order quantities for this generic uncertainty set, for static and dynamic robust inventory models under capacitated inventory constraints. We then build on the recent work of Bandi and Bertsimas (2012), which introduces CLT-style uncertainty sets, to parameterize our uncertainty set. We derive closed-form solutions in a static setting with correlated nonidentically distributed demand with capacitated inventories; these simple and intuitive solutions allow a decision maker to more easily understand the inventory management strategy and more

readily implement it (e.g., no optimization software is needed).

- We identify uncertainty set *symmetry*, or the lack of it, as an important characteristic of our robust inventory management problems, which leads to qualitatively different ordering behaviors. If the robust uncertainty set is symmetric, our robust strategies order positive quantities in all periods. By contrast, if the uncertainty set is asymmetric, this results in a lull of zero ordering (for multiple periods) in the middle of the planning horizon; these closed-form solutions are, to the best of our knowledge, the first for an asymmetric uncertainty set.

- We define a sequence of static problems whose collective solutions result in a *dynamic* ordering strategy. The order quantity in each period depends on the observable inventory position at the beginning of the period. These order quantities are also closed form and allow for correlated nonidentically distributed demand and capacitated inventories.

- We perform extensive numerical experiments to evaluate the effectiveness of the robust solutions derived in this paper. We find that our approach compares favorably to the robust solution developed in Bertsimas and Thiele (2006), which studies a similar problem, and to the dynamic affine policies studied in Bertsimas et al. (2010), which can be used for inventory problems. Our robust solutions perform particularly well when demand across periods is correlated and for high service-level values (above 90%). For these cases, our robust solutions indeed *lower* costs compared with the solutions in Bertsimas and Thiele (2006) and Bertsimas et al. (2010) in 71% and 66% of all the scenarios considered, respectively, and average cost reductions can be as high as 47% and 50%, respectively. Furthermore, our robust solution had a better worst-case cost than those in Bertsimas and Thiele (2006) and Bertsimas et al. (2010) in 60% and 55% of scenarios, respectively. These results are especially promising considering that the robust solutions developed in this paper require much less computational time as a result of their simple form.

1.3. Paper Outline

In Section 2 we discuss preliminaries, focusing on the CLT uncertainty sets. In Section 3 we introduce our static robust optimization model and detail our robust strategies. In Section 3.1 we derive the order quantities when the robust uncertainty set is symmetric, and in Section 3.2 we do the same for an uncertainty set that is asymmetric. Section 3.3 discusses the transition from the symmetric case to the asymmetric case, and Section 3.4 provides our closed-form ordering quantities for the capacitated inventory case. In Section 4 we derive analogous dynamic robust order quantities, which are functions of observable inventory positions. Computational experiments, focusing on our dynamic

strategies, and a discussion of the results are provided in Section 5. Managerial implications and concluding thoughts are given in Section 6. All proofs are given in the online appendix.

2. Preliminaries

2.1. Notations

We first detail the data for our models. We consider the n -period inventory management of a single product where the objective is to minimize total cost. The unit purchase cost is $c > 0$, the unit inventory holding cost is $h > 0$, and the unit inventory shortage cost is $s > 0$. In period $i = 1, \dots, n$, $d_i \geq 0$ are the stochastic demands; however, the distributions of the d_i are not known. We allow for period-dependent moments and correlated demand; we discuss details in Section 2.5. Note that the support of demand can be unbounded. In period i , $q_i \geq 0$ is the order quantity. We assume zero lead times.

2.2. Robust Optimization

Robust optimization is a methodology for optimization under uncertainty, where *sets* are utilized to characterize the uncertainty, rather than stochastic distributions. In stochastic optimization, an objective function $g(x, y)$ depends on a decision x and a random variable y with a *known* distribution F . By contrast, in robust optimization, the distribution F is not known, and the uncertainty in y is instead characterized by an “uncertainty set” U . In other words, $y \in U$, where U is not necessarily the support of the original random variable. Indeed, choosing an uncertainty set that is smaller than the support can lead to good performance; see, for example, the discussion on pp. 32–33 of Ben-Tal et al. (2009). There is a variety of approaches to define a robust counterpart to a stochastic program; we next discuss some of the most common ones.

One of the first studies of robust linear programs is Soyster (1973), where uncertain problem parameters are characterized by “interval uncertainty” (i.e., each parameter is constrained to be within a given interval). One shortcoming of this analysis is the overconservatism of the optimal solution, which restricted its practicality. In a series of papers (El-Ghaoui and Lebret 1997; El-Ghaoui et al. 1998; Ben-Tal and Nemirovski 1998, 1999, 2000), this issue of overconservatism is addressed via more sophisticated uncertainty sets, which results in more complex but solvable convex optimization problems (such as a second-order cone program). Bertsimas and Sim (2004) consider a different approach to address the overconservatism of Soyster (1973): they keep the interval uncertainty sets yet limit the total dispersion of uncertain parameters from nominal values by using the notion of a “budget of uncertainty.” By doing so, they retain the complexity of the original optimization model (e.g., a robust linear program is still a linear program). More recently, Bandi and

Bertsimas (2012) utilize an alternative method based on the limit theorems of probability, instead of the budgets of uncertainty, to reduce the conservatism of robust models; this work strongly influences the parameterization of our robust uncertainty sets.

2.3. Partial-Sum Uncertainty Sets

We utilize the following class of uncertainty sets in this paper:

$$\Omega = \left\{ (d_1, \dots, d_n): a_i \leq \sum_{j=1}^i d_j \leq b_i, \right. \\ \left. \ell_i \leq d_i \leq u_i, i = 1, \dots, n \right\}, \quad (1)$$

where the parameters $(a_i, b_i, \ell_i, u_i)_i$ are exogenous. Note that, for all i ,

$$a_i - b_{i-1} \leq d_i = \sum_{j=1}^i d_j - \sum_{j=1}^{i-1} d_j \leq b_i - a_{i-1}.$$

To guarantee that Ω is not empty, we assume $a_i \leq b_i$, $0 \leq \ell_i \leq u_i$, $a_i - b_{i-1} \leq b_i - a_{i-1}$, and $[a_i - b_{i-1}, b_i - a_{i-1}] \cap [\ell_i, u_i] \neq \emptyset$ for $i = 1, \dots, n$. In this paper we analyze multiple parameterizations of this uncertainty set, deriving closed-form ordering quantities in all cases.

We contrast our uncertainty set with that of Bertsimas and Thiele (2006), which studies a similar robust optimization approach to inventory management. In the Bertsimas and Thiele model, demand in period i is restricted to a symmetric set $d_i \in [\bar{d}_i - \hat{d}_i, \bar{d}_i + \hat{d}_i]$ for all i , where the parameters \bar{d}_i and \hat{d}_i are exogenous. Furthermore, these authors utilize “budgets of uncertainty” constraints, where, in period i , the constraints $\sum_{j=1}^i |(d_j - \bar{d}_j)/\hat{d}_j| \leq \gamma_i$ hold, where the γ_i are again exogenous. This uncertainty set can be written as

$$\Omega_{\text{BT}} = \left\{ (d_1, \dots, d_n): \sum_{j=1}^i \left| \frac{d_j - \bar{d}_j}{\hat{d}_j} \right| \leq \gamma_i, \right. \\ \left. \bar{d}_i - \hat{d}_i \leq d_i \leq \bar{d}_i + \hat{d}_i, i = 1, \dots, n \right\}.$$

We provide a specific parameterization of our uncertainty set to allow a clearer comparison. If we let $\ell_i = \bar{d}_i - \hat{d}_i$, $u_i = \bar{d}_i + \hat{d}_i$, $a_i = -\hat{d}_i \gamma_i + \sum_{j=1}^i \bar{d}_j$, and $b_i = \hat{d}_i \gamma_i + \sum_{j=1}^i \bar{d}_j$, our uncertainty set can be written as

$$\Omega' = \left\{ (d_1, \dots, d_n): \sum_{j=1}^i \frac{d_j - \bar{d}_j}{\hat{d}_j} \leq \gamma_i, \right. \\ \left. \bar{d}_i - \hat{d}_i \leq d_i \leq \bar{d}_i + \hat{d}_i, i = 1, \dots, n \right\}.$$

It turns out that in this comparison, the primary difference between these uncertainty sets is that Ω' has n absolute value of summation constraints, and Ω_{BT}

has n summation of absolute values constraints. Under similar operational settings (zero fixed costs and no capacities) and $c \leq s$, both the Bertsimas and Thiele (2006) study and our paper derive closed-form solutions; these solutions are similar yet not identical (see Remark 3 in Section 3). We also study the $c > s$ case, which Bertsimas and Thiele do not, and show how the ordering ceases prematurely in period $n - k + 1$, where $ks < c \leq (k + 1)s$ for some $k \in \{1, \dots, n - 1\}$; if $c > ns$, there is no ordering. Furthermore, the closed-form results in Bertsimas and Thiele rely on i.i.d. demand; we extend our results in a natural way to correlated nonidentically distributed demand.

Our paper further diverges in that we allow *asymmetric* uncertainty sets Ω as a result of the additional flexibility of the $(a_i, b_i, \ell_i, u_i)_i$ parameterization; for instance, we could modify the above definition of Ω' by letting $\ell_i = 0$ for all i . The symmetry of the uncertainty set is crucial in the Bertsimas and Thiele (2006) study, since the analysis is based on the work of Bertsimas and Sim (2004), which has symmetric random variables as a model primitive. We derive closed-form solutions under asymmetric uncertainty sets that have very different structures than those in Bertsimas and Thiele by appropriately choosing the $(a_i, b_i, \ell_i, u_i)_i$ parameterization.

2.4. Robust Optimization Models of Inventory Management

Our canonical model, motivated by that in Bertsimas and Thiele (2006), is formulated as the following robust optimization model:

$$\min_{(q_1, \dots, q_n) \geq 0} \sum_{i=1}^n (cq_i + y_i) \\ \text{s.t. } I_i = \sum_{j=1}^i (q_j - d_j), \quad i = 1, \dots, n, \quad (2) \\ y_i \geq hI_i, \quad i = 1, \dots, n, \forall (d_1, \dots, d_n) \in \Omega, \\ y_i \geq -sI_i, \quad i = 1, \dots, n, \forall (d_1, \dots, d_n) \in \Omega,$$

where I_i is the inventory position, and y_i captures the mismatch cost $\max\{hI_i, -sI_i\}$ in period i . In this static model, we assume zero initial inventory, though we relax this restriction in Section 4, where we analyze a dynamic rolling horizon strategy. Feasible values of variables q_i and y_i satisfy all constraints for every demand vector in Ω .

Remark 1. Note that our model is different than, and arguably suboptimal with respect to,

$$\min_{(q_1, \dots, q_n) \geq 0} \left\{ \max_{(d_1, \dots, d_n) \in \Omega} \left\{ \sum_{i=1}^n (cq_i + \max\{hI_i, -sI_i\}) \right\} \right\} \\ \text{s.t. } I_i = \sum_{j=1}^i (q_j - d_j), \quad i = 1, \dots, n,$$

which considers a single worst-case scenario. A variant of this model was analyzed by Rikun (2011), and a solution algorithm based on simulated annealing was developed; Gorissen and Hertog (2013) and Ardestani-Jaafari and Delage (2016) provide alternative approaches for a broader class of robust optimization problems involving sums of maxima and sums of piecewise linear functions, respectively. However, our aim is to derive closed-form ordering solutions that allow a decision maker to more easily understand an inventory control strategy and implement it in practice; this led us to formulate the model in problem (2). Furthermore, auxiliary computational studies (omitted) demonstrate that any suboptimality is minimal; in particular, our proposed closed-form solution results in an average 2% increase in cost relative to the single worst-case formulation, where the average is taken over various parameter (c, s, h, μ, σ) values.

The structure of Ω allows us to determine, in closed form, the minimum and maximum cumulative demands for the first k periods, $k = 1, \dots, n$, which we generically denote as

$$\underline{D}_k = \min_{(d_1, \dots, d_n) \in \Omega} \sum_{j=1}^k d_j \quad \text{and} \quad \overline{D}_k = \max_{(d_1, \dots, d_n) \in \Omega} \sum_{j=1}^k d_j, \quad (3)$$

respectively. Note that both \underline{D}_k and \overline{D}_k are increasing in k . The next lemma provides expressions for these partial sums for the uncertainty set defined in Equation (1).

Lemma 1. For $k = 1, \dots, n$,

$$\overline{D}_k = \min \left\{ \sum_{j=1}^k u_j, \min_{i < k} \left\{ b_i + \sum_{j=i+1}^k \min\{u_j, b_j - a_{j-1}\} \right\}, b_k, \min_{i > k} \left\{ b_i - \sum_{j=k+1}^i \max\{\ell_j, a_j - b_{j-1}\} \right\} \right\},$$

and

$$\underline{D}_k = \max \left\{ \sum_{j=1}^k \ell_j, \max_{i < k} \left\{ a_i + \sum_{j=i+1}^k \max\{\ell_j, a_j - b_{j-1}\} \right\}, a_k, \max_{i > k} \left\{ a_i - \sum_{j=k+1}^i \min\{u_j, b_j - a_{j-1}\} \right\} \right\}.$$

An appealing characteristic of our formulation is that we are able to find a robust ordering vector (q_1^*, \dots, q_n^*) where the cumulative orders $\sum_{j=1}^i q_j^*$ satisfy a newsvendor-type optimality condition, analogous to the case where cumulative demand up to period i is uniformly distributed on $[\underline{D}_i, \overline{D}_i]$ for all i . Furthermore, this characteristic persists in more complex environments, including a dynamic model, and we are also able to derive closed-form robust solutions in those cases. The next lemma provides an optimality condition for model (2).

Lemma 2. If $ks < c \leq (k + 1)s$ for some integer $0 \leq k \leq n - 1$, then the robust order quantities satisfy

$$\sum_{j=1}^i q_j^* = \frac{s\overline{D}_i + h\underline{D}_i}{s + h}, \quad i = 1, \dots, n - k, \quad \text{and} \\ q_i^* = 0, \quad i > n - k.$$

If $ns < c$, then $q_i^* = 0$ for all i .

The reasoning behind this lemma balances the worst-case holding costs and backordering costs for the first $n - k$ periods (the entire horizon if $c \leq s$). In other words, the worst-case holding cost in period i for the smallest possible cumulative demand \underline{D}_i —namely, $h(\sum_{j=1}^i q_j - \underline{D}_i)$ —will equal the worst-case backordering cost for the largest possible cumulative demand \overline{D}_i —namely, $s(\overline{D}_i - \sum_{j=1}^i q_j)$. This intuition is similar in spirit to that of the dual balancing policies of Levi et al. (2007) and Levi et al. (2008), but it allows for closed-form ordering quantities in multiple environments.

Note that for the $ks < c \leq (k + 1)s$ case, since $(s\overline{D}_i + h\underline{D}_i)/(s + h)$ is nonnegative and increasing in i , this lemma provides a recursion for determining the robust ordering quantities:

$$q_i^* = \begin{cases} \frac{s\overline{D}_i + h\underline{D}_i}{s + h} - \sum_{j=1}^{i-1} q_j^* & i = 1, \dots, n - k, \\ 0 & i > n - k. \end{cases} \quad (4)$$

Although this recursion is valid for all parameterizations $(a_i, b_i, \ell_i, u_i)_i$ of the uncertainty set Ω in Equation (1), we are more interested in nonrecursive closed-form expressions for the ordering quantities q_i^* . In the next subsection, we discuss a class of parameterizations, motivated by the central limit theorem, which allows for such closed-form expressions.

2.5. Central Limit Theorem–Inspired Parameterizations

The central limit theorem is one of the most well-known, and powerful, results of probability theory. Consider the sum of n i.i.d. random variables X_i , $i = 1, \dots, n$, each with mean μ and standard deviation σ . In one of its simplest forms, the CLT states that

$$\lim_{n \rightarrow \infty} P \left(\frac{\sum_{i=1}^n X_i - n\mu}{\sqrt{n}\sigma} \leq t \right) = \Phi(t) \quad \forall t, \quad (5)$$

where $\Phi(t)$ is the cumulative distribution of the standard normal random variable. Note that if the distribution of X_i is reasonably symmetrical and well behaved, then summations of very few variables (e.g., $n = 5$) can be well approximated by a normal distribution.

Although the CLT is typically presented for i.i.d. random variables, as above, this is not necessary. The Lindeberg CLT allows for nonidentically distributed

(and independent) random variables. For instance, if demands are uniformly bounded, and the standard deviation of the cumulative demand diverges as $n \rightarrow \infty$, then the Lindeberg CLT holds; full technical details about this extension can be found in Feller (1968, Chap. X, Section 5). There also exist variants of the CLT that allow for correlation. These results hold if subsequences that are far apart are “almost” independent; see Billingsley (2012, Chap. 5, Section 27) for technical details.

The generality of the CLT motivates the structure of uncertainty sets for our robust optimization models, an approach pioneered by Bandi and Bertsimas (2012). These sets lead to tractable robust formulations of difficult stochastic optimization problems in the areas of auction design, option pricing, queueing theory, and information theory (Bandi and Bertsimas 2011, 2014a, b; Bandi et al. 2012, 2015), and we extend this work to inventory management.

We assume that the demand in period i , d_i , is a random variable with mean μ_i and standard deviation σ_i ; we assume that these statistics are known but the distributions are not. We also assume that the demand covariance matrix Σ is known. The partial sum $\sum_{j=1}^k d_j$ has mean $\sum_{j=1}^k \mu_j$ and standard deviation $\sqrt{\mathbf{e}'_k \Sigma_k \mathbf{e}_k}$, where Σ_k is the submatrix composed of the first k rows and columns of Σ and \mathbf{e}_k is the k -dimensional vector of all ones. Motivated by the CLT, we let $a_i = \sum_{j=1}^i \mu_j - \Gamma_i \sqrt{\mathbf{e}'_i \Sigma_i \mathbf{e}_i}$, $b_i = \sum_{j=1}^i \mu_j + \Gamma_i \sqrt{\mathbf{e}'_i \Sigma_i \mathbf{e}_i}$, $\ell_i = \max\{\mu_i - \hat{\Gamma}_i \sigma_i, 0\}$, and $u_i = \mu_i + \hat{\Gamma}_i \sigma_i$, so that our uncertainty set becomes

$$\Omega_{\text{CLT}} = \left\{ (d_1, \dots, d_n) : -\Gamma_i \leq \sum_{j=1}^i \frac{d_j - \mu_j}{\sqrt{\mathbf{e}'_i \Sigma_i \mathbf{e}_i}} \leq \Gamma_i, \right. \\ \left. \max\{\mu_i - \hat{\Gamma}_i \sigma_i, 0\} \leq d_i \leq \mu_i + \hat{\Gamma}_i \sigma_i, i = 1, \dots, n \right\}. \quad (6)$$

Remark 2. Our approach to modeling correlation information is different from that of Bandi and Bertsimas (2012), whose approach is motivated by a factor-based model that generates correlated random variables from a vector of i.i.d. random variables. Our initial attempt to model correlation used a similar set to that in Bandi and Bertsimas (2012), but we were unable to derive closed-form solutions (because of the complications from the additional vector of generating variables); we consequently devised our current approach. Furthermore, our approach is appealing since the correlation information is a model primitive, as opposed to the approach in Bandi and Bertsimas (2012), where the factors are related to a decomposition (e.g., Cholesky) of the covariance matrix.

The constraints $d_i \in [\max\{\mu_i - \hat{\Gamma}_i \sigma_i, 0\}, \mu_i + \hat{\Gamma}_i \sigma_i]$ preclude extreme demand movements that are unlikely to appear in practice and are similar to the interval

uncertainty in Bertsimas and Thiele (2006). Furthermore, these period-specific constraints can be viewed to hold with high probability, as a result of, for example, the Chebyshev inequality: $P(|d_i - \mu_i| \geq \hat{\Gamma}_i \sigma_i) \leq 1/\hat{\Gamma}_i^2$. The $-\Gamma_i \leq \sum_{j=1}^i ((d_j - \mu_j)/\sqrt{\mathbf{e}'_i \Sigma_i \mathbf{e}_i}) \leq \Gamma_i$ constraints are motivated by the CLT. The $\Gamma_i \geq 0$ and $\hat{\Gamma}_i \geq 0$ are tunable parameters that allow adjustment of the conservatism of the robust optimization approach, as pioneered in Bertsimas and Sim (2004). We next discuss the selection of these parameters.

Practically speaking, the Γ_i will be chosen as small constants, usually no more than 3; this is motivated by the fact that if Z is a standard normal random variable, then $P(-3 \leq Z \leq 3) \approx 99.7\%$. The simplest parameterization is to set $\Gamma_i = \Gamma$ for all i , where $\Gamma \in [2, 3]$. Another parameterization that we explore in our paper, one that emphasizes the limit aspect of the CLT, is to set $\Gamma_i \rightarrow \infty$ for $i < n$ and $\Gamma_n \in [2, 3]$.

The $\hat{\Gamma}_i$ parameters can be set using a qualitative description of the demand d_i in period i . For example, if demand in period i is assumed to be symmetric, we select $\hat{\Gamma}_i \leq \mu_i/\sigma_i$, so that the constraint $\mu_i - \hat{\Gamma}_i \sigma_i \leq d_i \leq \mu_i + \hat{\Gamma}_i \sigma_i$ is symmetric around μ_i . By contrast, if the distribution of demand in period i is asymmetric, a value of $\hat{\Gamma}_i > \mu_i/\sigma_i$ is required to have an asymmetric constraint $0 \leq d_i \leq \mu_i + \hat{\Gamma}_i \sigma_i$; this implies that $\hat{\Gamma}_i \in (\mu_i/\sigma_i, 3]$, which is feasible only when the coefficient of variation is $\sigma_i/\mu_i \geq 0.33$. Put differently, in this paper we study asymmetric demand distributions where the lower bound of demand is zero. Such demand distributions arise in inventory management and are used to model the demand for “slow-moving” or “low-volume” products where demand is intermittent (Chen and Yu 2005). For instance, many spare aircraft parts satisfy the bound on σ_i/μ_i and exhibit asymmetric demand patterns, earning descriptions of “lumpy” and “erratic” in the literature (Ghobbar and Friend 2003, Williams 1984). In particular, Ghobbar and Friend (2003) studies a variety of aircraft parts with coefficients of variation up to 1.28. Therefore, highly variable asymmetric demands warrant a value of $\hat{\Gamma}_i \in (\mu_i/\sigma_i, 3]$ and the resulting asymmetric uncertainty set.

2.6. Heavy-Tail Parameterizations

Although our paper focuses on the above CLT-inspired parameterizations, the general form of Equation (1) allows us to also capture demand distributions that are “heavy tailed” (e.g., distributions without finite variance). Although the CLT does not apply for these distributions, other *stable laws* do apply; for instance, if Y_i are i.i.d. random variables with mean μ but no finite variance, a typical stable law is $\lim_{n \rightarrow \infty} \sum_{i=1}^n (Y_i - \mu)/n^{1/\alpha} \sim Y$ for some random variable Y and $\alpha \in (1, 2]$; see Theorem 1 in Bandi and Bertsimas (2012) for details and further citations. This discussion motivates the

following robust uncertainty set that is expressible in our framework:

$$\Omega_{HT} = \left\{ (d_1, \dots, d_n): -\Gamma_i \leq \sum_{j=1}^i \frac{d_j - \mu_j}{n^{1/\alpha}} \leq \Gamma_i, \right. \\ \left. \ell_i \leq d_i \leq u_i, i = 1, \dots, n \right\}. \quad (7)$$

3. Robust Order Quantities for the Static Problem

In this section we derive closed-form ordering solutions for the static model in Equation (2) for various parameterizations of the Ω_{CLT} uncertainty set defined in Equation (6). These robust order quantities (q_1^*, \dots, q_n^*) must be completely specified at time 0. In Section 4 we define a dynamic analogue of this model, and we solve for optimal robust order quantities that are functions of observable inventory positions.

Although dynamic policies are perhaps the most useful in practice, there is value in considering the static policies as well. Specifying multiple order quantities in advance is an appropriate model for some supply chain contracting, such as “advanced booking discount” programs (Tang et al. 2004). Similarly, Özer and Wei (2006) study advanced purchase contracts and show their strategic value in information sharing. In addition, Chen et al. (2014) study supply chain contracts for short-life-cycle products (e.g., fashion, high technology), where there is usually a single ordering opportunity for the life cycle. They explore inventory management *within* the product’s life cycle. They provide examples where restrictions (e.g., space constraints at 7-Eleven Japan, cash flow constraints) preclude the entire order being received at once, which results in a staggered delivery of products and accompanying inventory management issues. However, as in our static model, the staggered deliveries (and the total order) must be specified in advance.

In Section 3.1, we consider the case where $\mu_i - \hat{\Gamma}_i \cdot \sigma_i \geq 0$ for all i , which results in a symmetric uncertainty set, and in Section 3.2, we analyze the case where $\mu_i - \hat{\Gamma}_i \sigma_i < 0$ for all i , which gives an asymmetric uncertainty set; we obtain closed-form solutions in both cases. If there exists a set $S \subset \{1, \dots, n\}$ where $\mu_i - \hat{\Gamma}_i \sigma_i \geq 0$ for $i \in S$, and $\mu_i - \hat{\Gamma}_i \sigma_i < 0$ otherwise, then Lemma 2 can be used to obtain a recursion that can be solved for the robust quantities. In Section 3.3, we consider the transition from symmetric to asymmetric uncertainty sets for i.i.d. demand, and in Section 3.4, we generalize our results to capacitated inventories.

3.1. Symmetric Uncertainty Set

In this subsection we consider symmetric uncertainty sets. We assume that $\hat{\Gamma}_i$ are selected so that $\mu_i - \hat{\Gamma}_i \sigma_i \geq 0$ for all i . Furthermore, we let $\Gamma_i \rightarrow \infty$ for $i < n$ to retain only a constraint on the full summation $\sum_{j=1}^n d_j$, in

the asymptotic spirit of the CLT. If finite values of Γ_i , $i < n$, are required, the ordering quantities can still be determined using recursion (4); unfortunately, we are unable to derive closed-form q_i^* for finite Γ_i , $i < n$. The uncertainty set that we consider in this subsection is

$$\Omega_{CLT}^{symmetric} = \left\{ (d_1, \dots, d_n): -\Gamma_n \leq \sum_{j=1}^n \frac{d_j - \mu_j}{\sqrt{\mathbf{e}'\Sigma\mathbf{e}}} \leq \Gamma_n, \right. \\ \left. \mu_i - \hat{\Gamma}_i \sigma_i \leq d_i \leq \mu_i + \hat{\Gamma}_i \sigma_i, i = 1, \dots, n \right\}, \quad (8)$$

where \mathbf{e} is the n -dimensional vector of ones. This uncertainty set also captures pooling effects in the sense that the demands in periods $1, \dots, j$, for large enough j , are not all allowed to simultaneously be at their maximum (or minimum) bounds. In other words, the total amount of uncertainty over these periods is reduced, and this is due to the global constraint of the uncertainty set, which is motivated by the CLT. If j is not large enough, then there is no pooling effect, and the intuition is that the “limit” of the CLT does not yet apply. The threshold value of j depends on many details, but, for example, a threshold for i.i.d. demand is $(n + \sqrt{n})/2$, as we shall see shortly.

Our first theorem characterizes the robust order quantities in closed form. The subsequent discussion focuses on the $c \leq s$ case; if $c > s$, the ordering pattern is preserved, with the simple modification that ordering stops before period n , per Lemma 2.

Theorem 1. *If $\mu_i - \hat{\Gamma}_i \sigma_i \geq 0$ for all i and $c \leq s$, the robust order quantities are*

$$q_i^* = \begin{cases} \mu_i + \hat{\Gamma}_i \sigma_i \left(\frac{s-h}{s+h} \right) & i \leq \kappa, \\ \mu_i + \left(\Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} + \sum_{j=\kappa+2}^n \hat{\Gamma}_j \sigma_j - \sum_{j=1}^{\kappa} \hat{\Gamma}_j \sigma_j \right) \cdot \left(\frac{s-h}{s+h} \right) & i = \kappa + 1, \\ \mu_i - \hat{\Gamma}_i \sigma_i \left(\frac{s-h}{s+h} \right) & i > \kappa + 1, \end{cases} \quad (9)$$

where $\kappa = \max\{i: \sum_{j=1}^i \hat{\Gamma}_j \sigma_j \leq (\Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} + \sum_{j=1}^n \hat{\Gamma}_j \sigma_j)/2\}$. If $\kappa s < c \leq (\kappa + 1)s$ for some integer $1 \leq \kappa \leq n - 1$, the orders in Equation (9) are applied for $i = 1, \dots, n - \kappa$ and $q_i^* = 0$ for $i = n - \kappa + 1, \dots, n$. If $\kappa s < c$, $q_i^* = 0$ for all i .

If $\Sigma = \sigma \mathbf{I}$, where \mathbf{I} is the identity matrix and $\mu_i = \mu$ for all i , we obtain the case of i.i.d. demand with mean μ and standard deviation σ . For simplicity, we also let $\Gamma_n = \hat{\Gamma}_i = \Gamma$ for all i . The uncertainty set in Equation (8) simplifies to

$$\left\{ (d_1, \dots, d_n): -\Gamma \leq \frac{\sum_{j=1}^n d_j - n\mu}{\sqrt{n}\sigma} \leq \Gamma, \right. \\ \left. \mu - \Gamma\sigma \leq d_i \leq \mu + \Gamma\sigma, i = 1, \dots, n \right\}.$$

The ordering quantities for i.i.d. demand are especially simple and are presented in the following corollary. The remainder of this subsection focuses on these simpler ordering quantities since they allow an intuitive discussion.

Corollary 1 (i.i.d. Demand). *If $\mu - \Gamma\sigma \geq 0$ and $c \leq s$, the robust order quantities are*

$$q_i^* = \begin{cases} \mu + \Gamma\sigma \left(\frac{s-h}{s+h} \right) & i \leq \lfloor \tau \rfloor, \\ \mu - \Gamma\sigma \left(\frac{s-h}{s+h} \right) (1 - 2\epsilon) & i = \lfloor \tau \rfloor + 1, \\ \mu - \Gamma\sigma \left(\frac{s-h}{s+h} \right) & i > \lfloor \tau \rfloor + 1, \end{cases} \quad (10)$$


where $\tau = (n + \sqrt{n})/2$ and $\epsilon = \tau - \lfloor \tau \rfloor$. If $ks < c \leq (k+1)s$ for some integer $1 \leq k \leq n-1$, the orders in Equation (10) are applied for $i = 1, \dots, n-k$ and $q_i^* = 0$ for $i = n-k+1, \dots, n$. If $ns < c$, $q_i^* = 0$ for all i .

Remark 3. Under the $c \leq s$ case and comparable conditions, Bertsimas and Thiele (2006) derive a similar order-up-to level $S_k = \mu + \Gamma\sigma((s-h)/(s+h)) \cdot (\gamma_k - \gamma_{k-1})$, where γ_k is the budget of uncertainty for period k (see Remark 3 on p. 156 of Bertsimas and Thiele 2006). However, this solution is different from ours for any sequence of γ_k since these budgets must be nondecreasing in k , which is not compatible with the third case of Equation (10). Furthermore, Bertsimas and Thiele (2006) do not analyze the $c > s$ case, as we do throughout our paper, and their result depends on identically distributed demand; our Theorem 1 provides a closed-form solution for correlated nonidentically distributed demand.

The robust strategy in Corollary 1 orders in every period; this will not be true when $\mu - \Gamma\sigma < 0$ (which we study in the next section). Note that if ordering too much is equivalent to ordering too little (i.e., $s = h$), then the robust strategy will always order the mean of demand μ , as in the classic newsvendor solution (when the median equals the mean). More generally, in Figure 1, we plot various ordering strategies and observe a clear pattern that depends on the relative size of s and h , conveniently represented as the service level $s/(s+h)$. For service levels that are at least 50%, the strategy first orders aggressively, until a threshold $\lfloor \tau \rfloor$, and then reduces the orders; the opposite behavior of ordering conservatively at first, and then aggressively, is observed for service levels below 50% (plots are omitted). We also observe that the standard deviation σ and the budget-of-uncertainty parameter Γ appear in the ordering strategy as the product $\Gamma\sigma$, which can be interpreted as the key variability metric in the i.i.d. case. Similar behavior is observed for other values of n, μ, σ , and Γ .

The threshold period induced by the value of τ , where the ordering behavior changes, can be explained

Figure 1. Illustration of Corollary 1 for $n = 30, \mu = 10, \sigma = 3, \Gamma = 3$, and Various Service Levels $s/(s+h)$



as follows: When $i \leq \lfloor \tau \rfloor$, the order quantity is driven by the period-dependent constraints $d_i \in [\mu - \Gamma\sigma, \mu + \Gamma\sigma]$, whereas when $i > \lfloor \tau \rfloor + 1$, the order quantity is driven by the CLT constraints $-\Gamma \leq (\sum_{i=1}^n d_i - n\mu)/(\sqrt{n}\sigma) \leq \Gamma$. When $i = \lfloor \tau \rfloor + 1$, both have an influence. Note that this threshold is independent of the values of μ, σ , and Γ ; this is not true in the next section, where $\mu - \Gamma\sigma < 0$ and the uncertainty set contains the zero vector.

Next we discuss the comparative statics, with respect to μ and σ , of the individual orders $q_i^*, i = 1, \dots, n$, and the total cumulative order $Q = \sum_{i=1}^n q_i^*$. Note that the order quantities' behavior as a function of Γ is identical to that of σ . The following corollaries summarize these behaviors.

Corollary 2. *For $c \leq s$,*

- (1) $\partial q_i^* / \partial \mu > 0$ for all i .
- (2) If $s > h$, then $\partial q_i^* / \partial \sigma > 0$ for $i \leq \lfloor \tau \rfloor$; $\partial q_i^* / \partial \sigma < 0$ otherwise.
- (3) If $s < h$, then $\partial q_i^* / \partial \sigma < 0$ for $i \leq \lfloor \tau \rfloor$; $\partial q_i^* / \partial \sigma > 0$ otherwise.

Note that all order quantities are increasing in μ , which is not surprising. However, the effect of σ is not monotone. Larger σ values increase the space of feasible demand vectors. In particular, increasing σ widens the range of possible demand in each period by increasing the upper bound of demand and decreasing the lower bound. If shortages are more expensive than holding inventory ($s > h$), then the robust order quantities primarily protect against shortages on the upper bound of demand, and in the earlier stages, orders increase. Subsequent order quantities are then decreased to compensate for increased inventory accumulated during the early periods. Alternatively, if shortages are less expensive ($s < h$), then the robust

quantities primarily protect against inventory holding costs on the lower bound of demand, and early orders are decreased. Subsequent quantities are then increased to compensate for larger backorders accrued in the early stages.

We next consider how the total order quantity $Q = \sum_{i=1}^n q_i^*$ behaves as a function of μ and σ . It can be easily shown, using basic algebra, that $Q = n\mu + \sqrt{n}\Gamma\sigma((s-h)/(s+h))$.

Corollary 3. For $c \leq s$, the total number of orders over n periods, $Q = n\mu + \sqrt{n}\Gamma\sigma((s-h)/(s+h))$, is increasing in μ . Furthermore, if $s > h$, then $\partial Q/\partial \sigma > 0$, and if $s < h$, then $\partial Q/\partial \sigma < 0$.

Corollary 3 can be interpreted in terms of the newsvendor model. Letting Φ^{-1} denote the inverse cumulative distribution function of the standard normal, if demand is normally distributed with mean μ and standard deviation σ , then the newsvendor solution is $\mu + \sigma z$, where $z = \Phi^{-1}(s/(s+h))$ is the z -score that depends on the overage and underage costs. In our model, the mean of cumulative demand is $D = \sum_{i=1}^n d_i$ is $\mu_D = n\mu$, the standard deviation is $\sigma_D = \sqrt{n}\sigma$, and the total robust order quantity can be written as $Q = \mu_D + \sigma_D \Gamma((s-h)/(s+h))$. Therefore, $\tilde{z} = \Gamma((s-h)/(s+h))$ can be interpreted as a robust z -score.

It is also insightful to note that, from Lemma 2, the optimality conditions are $Q_i = (s\bar{D}_i + h\underline{D}_i)/(s+h) \forall i$, where $Q_i = \sum_{j=1}^i q_j^*$, and \bar{D}_i (\underline{D}_i) is the maximum (minimum) cumulative demand up to period i that respects the uncertainty set. This condition is identical to the newsvendor solution for demand that is uniformly distributed on $[\underline{D}_i, \bar{D}_i]$ with overage and underage costs of h and s , respectively. Therefore, the derivation of the robust quantities can be viewed as a sequence of newsvendor solutions, where in the i th application, the costs up to period i are minimized for uniformly distributed demand, whose interval is determined from the uncertainty set. Note that this also shows why holding and shortage costs must be period independent; if not, the optimality condition would be generalized to $Q_i = (s_i\bar{D}_i + h_i\underline{D}_i)/(s_i + h_i)$, which is not necessarily nondecreasing in i , a requirement for $q_i = Q_i - Q_{i-1} \geq 0$.

3.2. Asymmetric Uncertainty Set

Here, we solve our static robust model, as described in model (2), under the assumption that $\mu_i - \hat{\Gamma}_i\sigma_i < 0$ for all i , which results in the bounds $d_i \in [0, \mu_i + \hat{\Gamma}_i\sigma_i]$ that are asymmetric around μ_i . We again let $\Gamma_i \rightarrow \infty$ for $i < n$, so that the uncertainty set considered in this subsection is

$$\Omega_{CLT}^{asymmetric} = \left\{ (d_1, \dots, d_n): -\Gamma_n \leq \sum_{j=1}^n \frac{d_j - \mu_j}{\sqrt{\mathbf{e}'\Sigma\mathbf{e}}} \leq \Gamma_n, \right. \\ \left. 0 \leq d_i \leq \mu_i + \hat{\Gamma}_i\sigma_i, i = 1, \dots, n \right\}. \quad (11)$$

The following theorem characterizes the robust order quantities for this case in closed form.

Theorem 2. If $\mu_i - \hat{\Gamma}_i\sigma_i < 0$ for all i and $c \leq s$, the robust order quantities are

$$q_i^* = \begin{cases} \frac{s}{s+h}(\mu_i + \hat{\Gamma}_i\sigma_i) & i \leq \kappa_1, \\ \frac{s}{s+h} \left(\sum_{j=1}^n \mu_j + \Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} - \sum_{j=1}^{\kappa_1} (\mu_j + \hat{\Gamma}_j\sigma_j) \right) & i = \kappa_1 + 1, \\ 0 & \kappa_1 + 1 < i \leq \kappa_2, \\ \frac{h}{s+h} \left(\sum_{j=1}^{\kappa_2+1} (\mu_j + \hat{\Gamma}_j\sigma_j) - \Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} - \sum_{j=1}^n \hat{\Gamma}_j\sigma_j \right) & i = \kappa_2 + 1, \\ \frac{h}{s+h}(\mu_i + \hat{\Gamma}_i\sigma_i) & i \geq \kappa_2 + 2, \end{cases} \quad (12)$$

where

$$\kappa_1 = \max \left\{ i: \sum_{j=1}^i (\mu_j + \hat{\Gamma}_j\sigma_j) \leq \sum_{j=1}^n \mu_j + \Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} \right\} \quad \text{and} \\ \kappa_2 = \max \left\{ i: \sum_{j=1}^i (\mu_j + \hat{\Gamma}_j\sigma_j) \leq \Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} + \sum_{j=1}^n \hat{\Gamma}_j\sigma_j \right\}.$$

If $ks < c \leq (k+1)s$ for some integer $1 \leq k \leq n-1$, the orders in Equation (12) are applied for $i = 1, \dots, n-k$ and $q_i^* = 0$ for $i = n-k+1, \dots, n$. If $ns < c$, $q_i^* = 0$ for all i .

We first explain the interval of zero ordering, which was not present in the previous subsection for a symmetric uncertainty set. For the asymmetric uncertainty set $\Omega_{CLT}^{asymmetric}$, demand can be zero, and the worst-case cumulative demands \bar{D}_i and \underline{D}_i are not necessarily strictly increasing in i . Indeed, the periods with zero ordering, $i \in (\kappa_1 + 1, \kappa_2]$, correspond exactly to the periods where \bar{D}_i and \underline{D}_i are both constant. Therefore, matching the worst-case costs, or, equivalently, $\sum_{j=1}^i q_j^* = (s\bar{D}_i + h\underline{D}_i)/(s+h)$, results in zero ordering.

Next, as in the previous subsection, we now focus on the case of i.i.d. demand by setting $\Sigma = \sigma\mathbf{I}$ and $\mu_i = \mu$ for all i ; we also let $\Gamma_n = \hat{\Gamma}_i = \Gamma$ for all i . The uncertainty set in Equation (11) simplifies to

$$\left\{ (d_1, \dots, d_n): -\Gamma \leq \frac{\sum_{j=1}^n d_j - n\mu}{\sqrt{n}\sigma} \leq \Gamma, \right. \\ \left. 0 \leq d_i \leq \mu + \Gamma\sigma, i = 1, \dots, n \right\},$$

and the ordering quantities are presented in the following corollary.

Corollary 4 (i.i.d. Demand). If $\mu - \Gamma\sigma < 0$ and $c \leq s$, the robust order quantities are

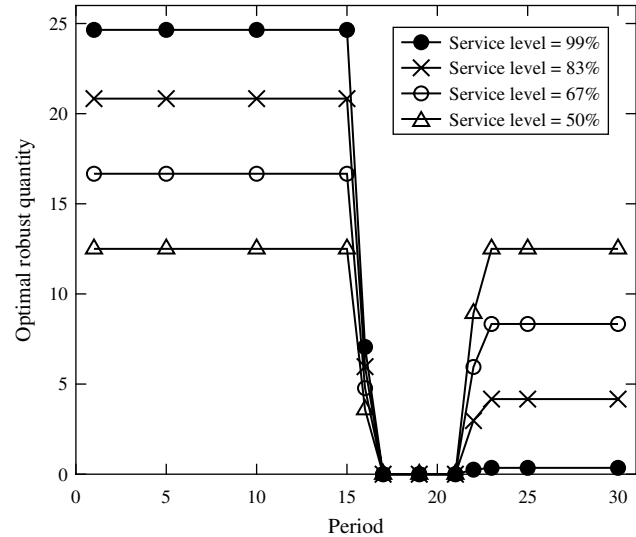
$$q_i^* = \begin{cases} \frac{s}{s+h}(\mu + \Gamma\sigma) & i \leq \lfloor \tau_1 \rfloor, \\ \epsilon \frac{s}{s+h}(\mu + \Gamma\sigma) & i = \lfloor \tau_1 \rfloor + 1, \\ 0 & \lfloor \tau_1 \rfloor + 1 < i \leq \lfloor \tau_2 \rfloor, \\ (1 - \epsilon) \frac{h}{s+h}(\mu + \Gamma\sigma) & i = \lfloor \tau_2 \rfloor + 1, \\ \frac{h}{s+h}(\mu + \Gamma\sigma) & i > \lfloor \tau_2 \rfloor + 1, \end{cases} \quad (13)$$

where $\tau_1 = (n\mu + \sqrt{n}\Gamma\sigma)/(\mu + \Gamma\sigma)$, $\tau_2 = (n + \sqrt{n})\Gamma\sigma/(\mu + \Gamma\sigma)$, $\epsilon = \tau_1 - \lfloor \tau_1 \rfloor$, and $\epsilon = \tau_2 - \lfloor \tau_2 \rfloor$. If $ks < c \leq (k+1)s$ for some integer $1 \leq k \leq n-1$, the orders in Equation (13) are applied for $i = 1, \dots, n-k$ and $q_i^* = 0$ for $i = n-k+1, \dots, n$. If $ns < c$, $q_i^* = 0$ for all i .

The different structure of the uncertainty set for the $\mu - \Gamma\sigma < 0$ case results in qualitatively different ordering behavior; in particular, there is a range of periods for which no ordering takes place ($q_i^* = 0$). There are now two thresholds, τ_1 and τ_2 , that dictate this lack of ordering, rather than the single threshold τ under the case where $\mu - \Gamma\sigma \geq 0$. It can be easily shown using basic algebra that $\tau_1 < \tau < \tau_2$. The drivers of the thresholds τ_1 and τ_2 , where the ordering behavior changes, can be explained as follows: When $i \leq \lfloor \tau_1 \rfloor$, the order quantity is driven by the period-dependent constraints $d_i \in [0, \mu + \Gamma\sigma]$, whereas when $i > \lfloor \tau_2 \rfloor + 1$, the order quantity is driven by the CLT constraints $-\Gamma \leq (\sum_{i=1}^n d_i - n\mu)/(\sqrt{n}\sigma) \leq \Gamma$. When $\lfloor \tau_1 \rfloor + 1 \leq i \leq \lfloor \tau_2 \rfloor + 1$, both constraints have an influence, as can be seen in the proof of Theorem 2. Unlike the threshold τ in the previous section, these thresholds depend on μ , σ , and Γ . As in the previous section, σ and Γ appear in the ordering strategy, as well as the threshold definitions, as the product $\Gamma\sigma$, which further supports using this product as the key variability metric in the i.i.d. case. Basic calculus shows that $\partial\tau_1/\partial\mu > 0$ and $\partial\tau_2/\partial\mu < 0$, which implies that, as the mean of demand increases, the range of intervals with zero ordering is reduced. Intuitively, this ordering strategy approaches that of the previous section. Conversely, we can show that $\partial\tau_1/\partial\sigma < 0$ and $\partial\tau_2/\partial\sigma > 0$, which implies that as the standard deviation of demand increases, the range of intervals with zero ordering is enlarged. Finally, note that it is possible to have $\tau_2 > n$ and the last regime of ordering is never reached; this occurs if $n < (\Gamma\sigma/\mu)^2$. In this scenario, ordering takes place only in early periods, which then ceases after period $\lfloor \tau_1 \rfloor + 1$.

The ordering behavior when $s = h$ also differs from the $\mu - \Gamma\sigma \geq 0$ case. Apart from not ordering in the middle of the horizon, the robust strategy will order a constant amount equal to half the maximum demand

Figure 2. Illustration of Corollary 4 for $n = 30$, $\mu = 10$, $\sigma = 5$, $\Gamma = 3$, and Various Service Levels $s/(s+h)$



$(\mu + \Gamma\sigma)/2$; note that this is in contrast to a newsvendor solution. However, this solution makes intuitive sense: since the costs of over- and underordering are the same, the robust strategy orders the average of the range of demand $[0, \mu + \Gamma\sigma]$.

In Figure 2 we plot various ordering strategies for service levels $s/(s+h) \geq 50\%$. We observe some similar patterns to that in Figure 1: the strategy first orders aggressively, then orders nothing, and finally orders conservatively; the opposite behavior is observed for service levels below 50% (plots are omitted).

Next we discuss the comparative statics, which are summarized in the following corollaries.

Corollary 5. For $c \leq s$, the robust order quantities q_i^* are increasing in μ and σ .

We compare and contrast this result with Corollary 2, which addresses the case of positive period-dependent constraints $d_i \in [\mu - \Gamma\sigma, \mu + \Gamma\sigma]$. In both cases, the robust order quantities are increasing in μ . By contrast, here, the robust order quantities are always increasing in σ , as opposed to those in Corollary 2, whose behaviors are not monotone in σ . This difference is due to σ increasing the upper bound of demand $(\mu + \Gamma\sigma)$ but not changing the lower bound (0). The rates of increase of q_i^* (with respect to either μ or σ) depend on the values of s and h , and they can be found by inspection.

We next consider how the total order quantity $Q = \sum_{i=1}^n q_i^*$ behaves as a function of μ and σ . Assuming that $n \geq (\Gamma\sigma/\mu)^2$, so that $\tau_2 \leq n$, it can be easily shown, using basic algebra, that $Q = n\mu + \sqrt{n}\Gamma\sigma((s-h)/(s+h))$. Note that this is the same expression as in the previous section, which implies that Corollary 3, and

its accompanying discussion, apply here as well; we formalize this observation in the following corollary.

Corollary 6. For $c \leq s$, if $n \geq (\Gamma\sigma/\mu)^2$, then the total number of orders over n periods, $Q = n\mu + \sqrt{n}\Gamma\sigma((s-h)/(s+h))$, are increasing in μ . Furthermore, if $s > h$, then $\partial Q/\partial\sigma > 0$, and if $s < h$, then $\partial Q/\partial\sigma < 0$.

Therefore, as long as n is not too small, the cumulative orders over n periods do not depend on the value of $\mu - \Gamma\sigma$. This is due to the lower bound on n , which ensures that the CLT applies equally in both the cases $\mu - \Gamma\sigma \geq 0$ and $\mu - \Gamma\sigma < 0$; note that in our discussion of the thresholds τ , τ_1 , and τ_2 , we elaborate on how the CLT dominates for $n > \tau$ and $n > \tau_2$. The major difference between the $\mu - \Gamma\sigma \geq 0$ and $\mu - \Gamma\sigma < 0$ cases—namely, the symmetry of the uncertainty set—does not factor into the cumulative ordering totals. However, if $n < (\Gamma\sigma/\mu)^2$, ordering under the $\mu - \Gamma\sigma < 0$ case ends prematurely in period $\lfloor \tau_1 \rfloor + 1$, and there are no orders for the remainder of the time horizon. In this scenario, the cumulative order is $Q = (s/(s+h))(n\mu + \sqrt{n}\Gamma\sigma)$, which increases in μ and σ ; we formalize this in the following corollary.


Corollary 7. For $c \leq s$, if $n < (\Gamma\sigma/\mu)^2$, then the total number of orders over n periods, $Q = (s/(s+h))(n\mu + \sqrt{n}\Gamma\sigma)$, are increasing in μ and σ .

3.3. Strategies as a Function of Γ

In this subsection, we discuss, for the i.i.d. case, the transition of our robust ordering strategies from the symmetric to the asymmetric cases, as a function of Γ . If $\Gamma < \mu/\sigma$, then the symmetric strategy of Corollary 1 applies, and if $\Gamma > \mu/\sigma$, then the asymmetric strategy of Corollary 4 is used. If $\Gamma = \mu/\sigma$, then these two strategies are equivalent since $\tau = \tau_1 = \tau_2$; the first cases of Corollaries 1 and 4 both simplify to an ordering level of $2\mu s/(s+h)$, the last cases both simplify to $2\mu h/(s+h)$, and the three middle cases of Corollary 4 simplify to the middle case of Corollary 1.

To understand the general dynamic, we present graphs for various values of $\Gamma \in [0, 3]$; we set $\mu = 15$ and $\sigma = 10$, so that the threshold $\mu/\sigma = 1.5$ is the midpoint of $[0, 3]$ and has a service level of 83% (an intermediate value considered in Figures 1 and 2; similar behaviors are observed for other service levels). In Figure 3 we consider $\Gamma \in \{0.25, 0.75, 1.50, 2.00, 2.75\}$, where the first two values lead to the symmetric strategy, the last two values result in the asymmetric one, and the middle value results in the equivalence of the two strategies. As Γ is increased from 0.25 to 1.50, the early (before period τ) order level increases, and the later (after period τ) order level decreases, expanding the range of ordering. Once the threshold μ/σ is surpassed and Γ is increased from 1.50 to 2.75, the early (before period τ_1) and later (after period τ_2) order levels both increase (early levels increasing faster if and only if $s > h$), and

Figure 3. Transition from Symmetric to Asymmetric Demand; $n = 30$, $\mu = 15$, $\sigma = 10$, Service Level = 83%, and Various Γ 's



the number of periods with no ordering also increases, since the size of the interval $[\tau_1, \tau_2]$ is increasing in Γ (τ_1 is decreasing and τ_2 is increasing). These periods with zero ordering are observed in practice, for demand that is “lumpy” and “erratic,” typically characterized by high coefficients of variation. For instance, the Steel Works, Inc., case presented in Simchi-Levi et al. (2008) discusses such an environment with frequent firm orders of zero (i.e., no ordering). As discussed in Section 2.5, this demand pattern corresponds to our asymmetric uncertainty set, which is the only set where zero orders (in intermediate periods) are optimal. See Ghobbar and Friend (2003) and Williams (1984) for further details regarding highly variable demand patterns.

3.4. Capacitated Problems

In this section we consider capacitated problems. In the first variant, we consider capacitated inventories and extend Lemma 2 to this case. In the second variant, we analyze capacitated orders and show that a natural extension of our results is not optimal, suggesting capacitated orders are more difficult than capacitated inventories in our modeling framework.

3.4.1. Capacitated Inventories. We introduce a capacity C that bounds the amount of inventory that can be held in each period. We add the constraints

$$\max_{(d_1, \dots, d_n) \in \Omega} I_i \leq C, \quad i = 1, \dots, n,$$

to our basic model in Equation (2). Note that these constraints are equivalent to $\sum_{j=1}^i q_j \leq C + \underline{D}_i$ for all i . The next result extends our closed-form recursive solutions to include capacitated inventories.

Lemma 3. If $c \leq s$, the robust order quantity for period i satisfies the following:

$$\sum_{j=1}^i q_j^* = \min \left\{ C + \underline{D}_i, \frac{s\bar{D}_i + h\underline{D}_i}{s+h} \right\}, \quad i = 1, \dots, n \quad (14)$$

for $i = 1, \dots, n$. If $ks < c \leq (k+1)s$ for some integer $1 \leq k \leq n-1$, the orders in Equation (14) are applied for $i = 1, \dots, n-k$ and $q_i^* = 0$ for $i = n-k+1, \dots, n$. If $ns < c$, then $q_i^* = 0$ for all i .

This lemma allows us to extend Theorems 1 and 2 for capacitated inventories. The impact of a capacity can be explained relatively simply: (1) if the capacity is high enough, the uncapacitated order quantities are optimal; (2) otherwise, there exist two indices $\ell < \ell'$ such that the ordering quantities before period ℓ and after period ℓ' are not affected by the capacity, whereas for the intermediate periods, ordering is reduced. Our first result, which generalizes Theorem 1, is for the symmetric uncertainty set $\Omega_{CLT}^{symmetric}$, defined in Equation (8).

Theorem 3. If $\mu_i - \hat{\Gamma}_i \sigma_i \geq 0$ for all i and $c \leq s$, we have the following:

- If $C \geq (2s/(s+h)) \max\{\sum_{j=1}^k \hat{\Gamma}_j \sigma_j, \Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} + \sum_{j=k+2}^n \hat{\Gamma}_j \sigma_j\}$, then the order quantities of Theorem 1 are optimal.
- Otherwise,

$$q_i^* = \begin{cases} \mu_i + \hat{\Gamma}_i \sigma_i \frac{s-h}{s+h} & i = 1, \dots, \ell, \\ C + \underline{D}_{\ell+1} - \sum_{j=1}^{\ell} \left(\mu_j + \hat{\Gamma}_j \sigma_j \frac{s-h}{s+h} \right) & i = \ell + 1, \\ \underline{D}_i - \underline{D}_{i-1} & \ell + 1 < i < \ell' - 1, \\ \frac{s\bar{D}_{\ell'} + h\underline{D}_{\ell'}}{s+h} - (C + \underline{D}_{\ell'-1}) & i = \ell', \\ \mu_i - \hat{\Gamma}_i \sigma_i \frac{s-h}{s+h} & i > \ell', \end{cases}$$

where

$$\ell = \max \left\{ i \leq \kappa: \frac{2s}{s+h} \sum_{j=1}^i \hat{\Gamma}_j \sigma_j \leq C \right\} \quad \text{and} \\ \ell' = \min \left\{ i > \kappa: \frac{2s}{s+h} \left(\Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} + \sum_{j=i+1}^n \hat{\Gamma}_j \sigma_j \right) \leq C \right\}.$$

If $ks < c \leq (k+1)s$ for some integer $1 \leq k \leq n-1$, the orders are applied for $i = 1, \dots, n-k$ and $q_i^* = 0$ for $i = n-k+1, \dots, n$. If $ns < c$, $q_i^* = 0$ for all i .

Our second result, which generalizes Theorem 2, is for the asymmetric uncertainty set $\Omega_{CLT}^{asymmetric}$, defined in Equation (11).

Theorem 4. If $\mu_i - \hat{\Gamma}_i \sigma_i < 0$ for all i and $c \leq s$, we have the following:

- If $C \geq (s/(s+h)) (\sum_{j=1}^n \mu_j + \Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}})$, then the order quantities of Theorem 2 are optimal.

- Otherwise,

$$q_i^* = \begin{cases} \frac{s}{s+h} (\mu_i + \hat{\Gamma}_i \sigma_i) & i = 1, \dots, \ell, \\ C + \underline{D}_{\ell+1} - \sum_{j=1}^{\ell} \frac{s}{s+h} (\mu_j + \hat{\Gamma}_j \sigma_j) & i = \ell + 1, \\ \underline{D}_i - \underline{D}_{i-1} & \ell + 1 < i < \ell' - 1, \\ \frac{s\bar{D}_{\ell'} + h\underline{D}_{\ell'}}{s+h} - (C + \underline{D}_{\ell'-1}) & i = \ell', \\ \frac{h}{s+h} (\mu_i + \hat{\Gamma}_i \sigma_i) & i > \ell', \end{cases}$$

where

$$\ell = \max \left\{ i \leq \kappa_1: \frac{s}{s+h} \sum_{j=1}^i (\mu_j + \hat{\Gamma}_j \sigma_j) \leq C \right\} \quad \text{and} \\ \ell' = \min \left\{ i > \kappa_2: \frac{s}{s+h} \left(\sum_{j=i+1}^n (\mu_j + \hat{\Gamma}_j \sigma_j) + 2\Gamma_n \sqrt{\mathbf{e}'\Sigma\mathbf{e}} \right) \leq C \right\}.$$

If $ks < c \leq (k+1)s$ for some integer $1 \leq k \leq n-1$, the orders are applied for $i = 1, \dots, n-k$ and $q_i^* = 0$ for $i = n-k+1, \dots, n$. If $ns < c$, then $q_i^* = 0$ for all i .

3.4.2. Capacitated Orders. We introduce a capacity B that bounds the order quantity in any period. We add the constraints

$$q_i \leq B, \quad i = 1, \dots, n,$$

to our basic model in Equation (2). We conjectured that the optimal orders satisfy the natural recursion

$$q_i^* = \min \left\{ B, \frac{s\bar{D}_i + h\underline{D}_i}{s+h} - \sum_{j=1}^{i-1} q_j^* \right\}, \quad i = 1, \dots, n,$$

but unfortunately, this solution is not optimal. For example, when $c = 1$, $s = 2$, $h = 5$, $n = 50$, $\mu = 15$, $\sigma = 5$, and $\Gamma = 3$, this conjectured solution gives a cost that is 1% higher than the optimal cost (as determined using Matlab's `linprog` function). In extensive computational tests, we discovered gaps exceeding 2% for large values of n . These gaps appeared only for $h > s$, and it does seem that the conjectured solution is indeed optimal for $s \leq h$. We leave a proof of this to future work.

4. Robust Policies for a Rolling Horizon Problem

In this section we analyze a dynamic robust model based on reoptimization, where in period k , past demands and order quantities are known, which are denoted by $(\hat{d}_1, \dots, \hat{d}_{k-1})$ and $(q_1^*, \dots, q_{k-1}^*)$, respectively. Consequently, the inventory position at the end of period $k-1$, denoted by $\hat{I}_{k-1} = \sum_{j=1}^{k-1} (q_j^* - \hat{d}_j)$, is also known.

We define the uncertainty set in period k , Ω_{CLT}^k , by the intersection of the original uncertainty set Ω_{CLT} , defined in Equation (6), with the knowledge of realized demand, or

$$\begin{aligned} \Omega_{\text{CLT}}^k &= \Omega_{\text{CLT}} \cap \{d_1 = \hat{d}_1, \dots, d_{k-1} = \hat{d}_{k-1}\} \\ &= \left\{ (d_k, \dots, d_n) : -\Gamma_i \leq \frac{\sum_{j=1}^{k-1} \hat{d}_j + \sum_{j=k}^i d_j - \sum_{j=1}^i \mu_j}{\sqrt{\mathbf{e}_i' \Sigma_i \mathbf{e}_i}} \leq \Gamma_i, \right. \\ &\quad \left. \max\{\mu_i - \hat{\Gamma}_i \sigma_i, 0\} \leq d_i \leq \mu_i + \hat{\Gamma}_i \sigma_i, i = k, \dots, n \right\}; \end{aligned}$$

we assume that Ω_{CLT}^k is not empty. Note that, in contrast to the results of Section 3, we allow finite values of Γ_i for $i < n$; the reoptimization aspect of our approach allows us to also obtain closed-form ordering quantities under this more general uncertainty set. We next provide a sequence of optimization models that collectively define a dynamic version of model (2). These models are used to find the robust order quantity in period k , q_k^* , for $k = 1, \dots, n$, which is a function of the inventory position \hat{I}_{k-1} . In period $k = 1, \dots, n$, we leverage the known information by solving the following model:

$$\begin{aligned} \min_{(q_k, \dots, q_n) \geq 0} & \sum_{i=k}^n (c q_i + y_i) \\ \text{s.t. } & I_i = \hat{I}_{k-1} + \sum_{j=k}^i (q_j - d_j), \quad i = k, \dots, n, \\ & y_i \geq h I_i, \quad i = k, \dots, n, \forall (d_k, \dots, d_n) \in \Omega_{\text{CLT}}^k, \\ & y_i \geq -s I_i, \quad i = k, \dots, n, \forall (d_k, \dots, d_n) \in \Omega_{\text{CLT}}^k, \end{aligned} \tag{15}$$

where I_i is the inventory position in period $i \geq k$. Note that in period k , only the solution q_k^* will be implemented, as the final determination of q_j^* will depend on the realized \hat{I}_{j-1} for $j > k$.

Remark 4. Note that problem (15) is not in the dynamic programming style, as each period’s optimization does not factor in future optimizations, and is suboptimal. However, we believe this relaxation is the reason we are able to derive closed-form solutions that can still be applied in a dynamic setting and that exhibit excellent computational performance (see Section 5). Alternatively, there are papers that explicitly combine the dynamic programming philosophy with robust uncertainty sets. For example, Bertsimas et al. (2010) and Iancu et al. (2013) successfully analyze such approaches. Although these papers are unable to determine simple and intuitive closed-form solutions, their solutions are tractable (solvable as convex optimization problems), and their models more general than ours (e.g., allowing multiple inventory locations).

The definitions of \underline{D}_i and \overline{D}_i must be redefined slightly in this section to accommodate the past demand realizations:

$$\begin{aligned} \underline{D}_i &= \min_{(d_k, \dots, d_n) \in \Omega_{\text{CLT}}^k} \sum_{j=k}^i d_j, \quad i = k, \dots, n \quad \text{and} \\ \overline{D}_i &= \max_{(d_k, \dots, d_n) \in \Omega_{\text{CLT}}^k} \sum_{j=k}^i d_j, \quad i = k, \dots, n. \end{aligned} \tag{16}$$

The main structural difference between the optimization in period k , $k = 1, \dots, n$, and the static problem of the previous section is the possibility of nonzero initial inventory. Therefore, we generalize Lemma 2 to accommodate nonzero initial inventory, including backlogs, which we then leverage in subsequent results. The presentation of the new lemma is facilitated if we consider the full horizon $i = 1, \dots, n$, with initial inventory I_0 , and we introduce some conditions:

Condition 1. $ks < c \leq (k+1)s$ for some integer $0 \leq k \leq n-1$.

Condition 2. $(s\overline{D}_m + h\underline{D}_m)/(s+h) < I_0 \leq (s\overline{D}_{m+1} + h\underline{D}_{m+1})/(s+h)$ for some integer $0 \leq m \leq n-1$, where $\underline{D}_0 = \overline{D}_0 = 0$.

Condition 3. $m < n - k$.

Lemma 4. *The robust order quantities satisfy the following:*

(1) *If $I_0 > 0$ and Conditions 1–3 hold,*

(a) $q_i^* = 0, i = 1, \dots, m$;

(b) $\sum_{j=m+1}^i q_j^* = (s\overline{D}_i + h\underline{D}_i)/(s+h) - I_0, i = m+1, \dots, n-k$; and

(c) $q_i^* = 0, i = n-k+1, \dots, n$.

Otherwise, $q_i^ = 0$ for all i .*

(2) *If $I_0 < 0$ and Condition 1 holds,*

(a) $\sum_{j=1}^i q_j^* = (s\overline{D}_i + h\underline{D}_i)/(s+h) - I_0, i = 1, \dots, n-k$; and

(b) $q_i^* = 0, i = n-k+1, \dots, n$.

Otherwise, $q_i^ = 0$ for all i .*

The following result only requires the minimum and maximum demands in the *first* period, which are \overline{D}_i and \underline{D}_i , from Equation (16), evaluated at $i = k$:

$$\underline{d}_k = \min_{(d_k, \dots, d_n) \in \Omega_{\text{CLT}}^k} d_k \quad \text{and} \quad \overline{d}_k = \max_{(d_k, \dots, d_n) \in \Omega_{\text{CLT}}^k} d_k. \tag{17}$$

This simplification is due to the reoptimization nature of our approach, since in period k we only need to solve for q_k^* , which requires only \overline{d}_k and \underline{d}_k . This fact allows us to utilize the more general uncertainty set Ω_{CLT} , with finite Γ_i for $i < n$, and still obtain closed-form ordering quantities, in contrast to the static closed-form ordering quantities of Section 3, which required $\Gamma_i \rightarrow \infty$ for $i < n$. The dynamic robust ordering strategy is presented in the next theorem.

Theorem 5. If $c \leq s(n - k + 1)$, the robust order quantities in period $k = 1, \dots, n$ are

$$q_k^* = \max \left\{ \frac{s\bar{d}_k + h\underline{d}_k}{s + h} - \hat{I}_{k-1}, 0 \right\}, \quad (18)$$

where

$$\begin{aligned} \bar{d}_k = \min & \left\{ \mu_k + \hat{\Gamma}_k \sigma_k, \min_{i=k, \dots, n} \left\{ \sum_{j=1}^i \mu_j + \Gamma_i \sqrt{\mathbf{e}'_i \Sigma_i \mathbf{e}_i} - \sum_{j=1}^{k-1} \hat{d}_j \right. \right. \\ & \left. \left. - \sum_{j=k+1}^i \max \{ \mu_j - \hat{\Gamma}_j \sigma_j, 0, \mu_j - \Gamma_j \sqrt{\mathbf{e}'_j \Sigma_j \mathbf{e}_j} \right. \right. \\ & \left. \left. - \Gamma_{j-1} \sqrt{\mathbf{e}'_{j-1} \Sigma_{j-1} \mathbf{e}_{j-1}} \right\} \right\} \end{aligned}$$

and

$$\begin{aligned} \underline{d}_k = \max & \left\{ \max \{ \mu_k - \hat{\Gamma}_k \sigma_k, 0 \}, \max_{i=k, \dots, n} \left\{ \sum_{j=1}^i \mu_j - \Gamma_i \sqrt{\mathbf{e}'_i \Sigma_i \mathbf{e}_i} \right. \right. \\ & \left. \left. - \sum_{j=1}^{k-1} \hat{d}_j - \sum_{j=k+1}^i \min \{ \mu_j + \hat{\Gamma}_j \sigma_j, \mu_j + \Gamma_j \sqrt{\mathbf{e}'_j \Sigma_j \mathbf{e}_j} \right. \right. \\ & \left. \left. + \Gamma_{j-1} \sqrt{\mathbf{e}'_{j-1} \Sigma_{j-1} \mathbf{e}_{j-1}} \right\} \right\}. \end{aligned}$$

If $c > s(n - k + 1)$, then $q_k^* = 0$ for $k = 1, \dots, n$.

Note that, as in the static case, we again have closed-form solutions. By contrast, in this dynamic case q_k^* depends on the currently available inventory position \hat{I}_{k-1} , allowing feedback and improved decision making.

Remark 5. Note that the order quantities derived in Theorem 5 are state-dependent base stock policies where the base stock for period $k = 1, \dots, n$ is $(s\bar{d}_k + h\underline{d}_k)/(s + h)$; the dependence on state is from the definitions of \bar{d}_k and \underline{d}_k in Equation (17), which are functions of Ω_{CLT}^k , a set that depends on the past demand realizations $\hat{d}_1, \dots, \hat{d}_{k-1}$. When the demand distribution is available, base stock policies are known to minimize the expected total inventory-related costs in a periodic review model with no fixed ordering cost. When the demand distribution is not known, Bertsimas and Thiele (2006) show that the state-dependent base stock policy is the optimal robust inventory policy under their setting. Our results extend the Bertsimas and Thiele conclusion to a different class of demand uncertainty sets, including, for the first time, asymmetric sets.

Finally, we can extend our dynamic results to the capacitated inventory case. Lemmas 3 and 4 (and their proofs) imply the following result.

Corollary 8. If $c \leq s(n - k + 1)$, the robust order quantities in period $k = 1, \dots, n$ are

$$q_k^* = \min \left\{ C + \underline{d}_k, \max \left\{ \frac{s\bar{d}_k + h\underline{d}_k}{s + h} - \hat{I}_{k-1}, 0 \right\} \right\},$$

where \bar{d}_k and \underline{d}_k are defined in Theorem 5. If $c > s(n - k + 1)$, then $q_k^* = 0$ for $k = 1, \dots, n$.

5. Computational Experiments

In this section we present numerical experiments that respond to the motivating questions in the introduction and evaluate the effectiveness of the closed-form robust solutions derived in this paper. We do so by comparing the average total costs of the robust solution obtained from our model to that of Bertsimas and Thiele (2006), hereafter BT, as well as the dynamic affine policies in Bertsimas et al. (2010), hereafter AP, via Monte Carlo simulation. Given that the BT and AP models are dynamic implementations, in this section we present the results for our dynamic robust policy of Section 4 to create a more equitable comparison.

As argued previously, a firm may set its inventory ordering policy solely based on the means and covariance matrix of demand because it may not have the full knowledge of the demand distribution. That said, we assume that the actual demand follows a distribution F that is unknown to the firm. We therefore assess the average total cost of all models (BT, AP, and our paper's) in terms of F .

In summary, Section 5.1 compares the average total cost of BT and the robust policies developed in this paper when outcomes are evaluated based on the true demand distribution F . We find that although our model is simpler than the one in BT, which allows us to derive easy to implement closed-form solutions that can handle asymmetric uncertainty sets, it performs remarkably well and outperforms the BT solution in a majority of scenarios. Our model performs particularly well when the value of the service level is high, demand across periods is correlated, per-period demand variability is large, and/or demand uncertainty sets are asymmetric. We note that the model proposed in BT performs very well compared with the dynamic programming policies (Bertsimas and Thiele 2006); therefore, we expect our model's performance to be comparable as well.

Section 5.2 compares the outcomes of AP and the robust policies developed in this paper with respect to the average total costs based on the true demand distribution F . We find that whereas, by definition, the AP solution outperforms our solution when the demands in different periods are independent and have the same support as the uncertainty sets defined in Bertsimas et al. (2010), our robust solution leads to a lower cost when either one of these assumptions is not valid. More specifically, we show that our policy outperforms the AP solution in the majority of the scenarios considered, if F follows a multivariate normal distribution.

Table 1 lists the parameter values used in our experiments. We normalize the holding cost at $h = 1$ and vary shortage costs in the specified range to evaluate the effects of relatively low to high service levels ($s/(s + h) = 75\%$ to $s/(s + h) = 98\%$). We choose demand across periods to follow a multivariate normal distribution.

Table 1. Parameter Values for Numerical Experiments

n	{3, 10}
h	1
s	{3, 5, 20, 40}
c	{0.1, 0.5, 1, 2}
Γ	{1.0, 1.5, 2.0, 2.5, 3.0}
μ	5
σ	{0.1 μ , 0.3 μ , 0.5 μ , 0.8 μ , 1 μ , 1.5 μ , 2 μ }

We assume per-period demand distributions have an average of $\mu = 5$ and standard deviations that are varied in the range specified in Table 1. The results of our simulations are averaged across 70 randomly selected symmetric positive definite matrices as the covariance matrix of the multivariate normal distribution.

We ran $n_{\text{sims}} = 1,000$ simulation trials, where in each trial k we generated a demand vector $(\hat{d}_1^k, \dots, \hat{d}_n^k)$ from the multivariate normal distribution F . In particular, letting

$$G(q_1, \dots, q_n, d_1, \dots, d_n) = \sum_{i=1}^n \left[c_i q_i + h_i \max \left\{ \sum_{j=1}^i (q_j - d_j), 0 \right\} + s_i \max \left\{ - \sum_{j=1}^i (q_j - d_j), 0 \right\} \right]$$

in Section 5.1, we compare the estimated expected costs for the (dynamic) robust and BT quantities, namely,

$$E[C_R] = \frac{\sum_{k=1}^{n_{\text{sims}}} G(q_1^*, \dots, q_n^*, \hat{d}_1^k, \dots, \hat{d}_n^k)}{n_{\text{sims}}} \quad \text{and} \quad (19)$$

$$E[C_{\text{BT}}] = \frac{\sum_{k=1}^{n_{\text{sims}}} G(q_1^{\text{BT}}, \dots, q_n^{\text{BT}}, \hat{d}_1^k, \dots, \hat{d}_n^k)}{n_{\text{sims}}},$$

respectively. In Section 5.2 we compare the estimated expected costs for the (dynamic) robust and AP quantities, namely,

$$E[C_R] = \frac{\sum_{k=1}^{n_{\text{sims}}} G(q_1^*, \dots, q_n^*, \hat{d}_1^k, \dots, \hat{d}_n^k)}{n_{\text{sims}}} \quad \text{and} \quad (20)$$

$$E[C_{\text{AP}}] = \frac{\sum_{k=1}^{n_{\text{sims}}} G(q_1^{\text{AP}}, \dots, q_n^{\text{AP}}, \hat{d}_1^k, \dots, \hat{d}_n^k)}{n_{\text{sims}}},$$

respectively.

Note that the robust order quantities q_i^* are given in Theorem 5, whereas the BT quantities q_i^{BT} are provided in Theorem 3.2 of Bertsimas and Thiele (2006) and the AP quantities q_i^{AP} are provided in Theorem 3.1 of Bertsimas et al. (2010). In calculating BT quantities, one needs to define the so-called budgets of uncertainty, which ensure that the total variation of the parameters cannot exceed a certain threshold. We use period-dependent budgets of uncertainty based on the

algorithm in Bertsimas and Thiele (2006), though our results are not sensitive to small variations of these parameters. Furthermore, when $\mu - \Gamma\sigma < 0$, we modified the demand intervals in BT accordingly to make the two approaches commensurable.

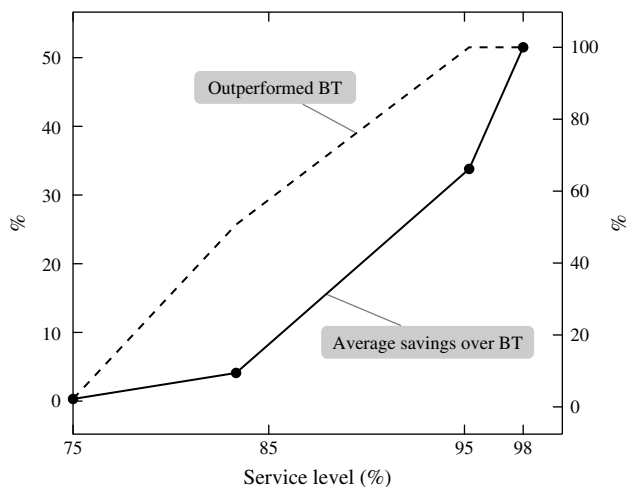
Before presenting our results, we note that in all of our experiments, the simulations generate fairly stable outcomes. More specifically, the standard errors of the BT, the AP, and our robust solutions were, on average, less than 1% of the average total cost of each policy for $n_{\text{sims}} = 1,000$ trials. Furthermore, the maximum standard error observed in our experiments was less than 2% of the corresponding average total cost.

5.1. BT Comparison

In this section we compare the average total cost of an inventory system when the robust quantities are selected based on the framework developed in this paper versus that of Bertsimas and Thiele (2006). Interestingly, when true demand is multivariate normally distributed, our results indicate that our robust solution can indeed outperform the BT solution's average cost in 70% of the cases in the entire test set. Furthermore, for scenarios where our robust solution leads to a lower cost than the BT policy, the average cost savings for using our robust policy is more than 45%. Put differently, the average percentage improvement in total cost for using our solution over the BT solution ($(E[C_{\text{BT}}] - E[C_R])/E[C_{\text{BT}}]$) for these scenarios was more than 45%. For those scenarios where the BT solution leads to a lower cost than our robust policy, the average cost savings for using the BT policy ($(E[C_R] - E[C_{\text{BT}}])/E[C_R]$) is less than 10%. It is noteworthy to point out that our robust policy achieves such an improvement in performance while being less complicated and extremely easy to implement, is insightful because of its closed-form nature, and can be run in a shorter time.

Figure 4 presents the simulation results for 70 randomly selected multivariate normal distributions and parameter values chosen from Table 1. Figure 4 plots the average percentage cost improvement of BT over our closed-form robust solution on the left axis, whereas the right axis shows the percentage of cases where our solution yields a lower average total cost than BT, as a function of the service level ($s/(s+h)$). We observe that our model performs considerably better for larger values of the service level. For instance, for all service level values above 95%, our robust solution yields a lower average cost than BT in all scenarios of our simulation. We note that, in general, our model performs best for larger values of c and σ , moderate values of Γ , and smaller values of n (graphs not provided here). Furthermore, our robust solution performs better when the uncertainty set Ω is asymmetric and demand across periods is correlated.

Figure 4. Performance of BT Compared with Our Robust Solution for Randomly Selected Multivariate Normal Distributions



Notes. The average percentage improvements of our policy over BT for all scenarios are shown on the left axis. Percentages of scenarios where our robust solution outperforms the BT policy are shown on the right axis.

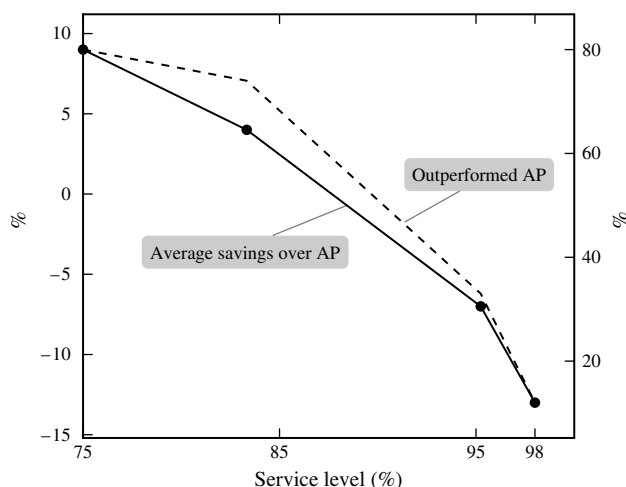
Finally, we also investigate the case where demand is correlated but not normally distributed. In particular, we utilize uniform, not normal, random variables as the primitive to generate correlated demand with a given covariance matrix. Our robust solution still outperforms the BT solution's average in 65% of the cases over the entire test set. For the scenarios where our robust solution leads to a lower cost than the BT policy, the average cost savings from using our robust policy is more than 46%. For those scenarios where the BT solution leads to a lower cost than our robust policy, the average cost savings for the BT policy is less than 19%.

5.2. AP Comparison

In this section we compare the total cost of our robust solution to that of the dynamic affine policies developed in Bertsimas et al. (2010), modified to account for the uncertainty set Ω_{CLT} in (6). Whereas the dynamic affine policies in Bertsimas et al. (2010) are shown to be optimal under some assumptions, we observe that our robust solution outperforms the AP solution's average cost in more than 51% of scenarios in our simulations when some of those assumptions are violated. Furthermore, for the scenarios where our robust solution leads to a lower cost than the AP solution, the average cost reduction for using our robust policy $((E[C_{AP}] - E[C_R])/E[C_{AP}])$ is more than 8%. For those scenarios where the AP solution leads to a lower cost than our robust policy, the average cost savings for using the AP solution $((E[C_R] - E[C_{AP}])/E[C_R])$ is less than 10%.

Figure 5 presents the simulation results for 70 randomly selected multivariate normal distributions and

Figure 5. Performance of AP Compared with Our Robust Solution for Randomly Selected Multivariate Normal Distributions



Notes. The average percentage improvements of our policy over AP for all scenarios are shown on the left axis. Percentages of scenarios where our robust solution outperforms the AP policy are shown on the right axis.

parameter values chosen from Table 1. It plots the average percentage cost improvement of AP over our closed-form robust solution on the left axis, whereas the right axis shows the percentage of cases where our solution yields a lower average total cost than AP, as a function of the service level.

Not surprisingly, the solution from the dynamic affine robust policy developed in Bertsimas et al. (2010) outperforms the robust solution in Bertsimas and Thiele (2006). Interestingly enough, however, although the dynamic affine policies in Bertsimas et al. (2010) are shown to be optimal, we observe that our robust solution outperforms the AP solution in the majority of scenarios. The reason for this observation is twofold: (1) the AP solution performs particularly well when the support of the demand vector coincides with the uncertainty set, which is not the case for normally distributed demand; and (2) the AP solution assumes demand across different periods are independent, whereas the demand in our simulation is correlated across different periods. In fact, if we repeat the numerical experiments with independent uniform demand in each period, with the support being identical to the range of the uncertainty set, then the AP solution outperforms our robust solution in all scenarios.

These observations suggest that the main advantage of our robust solution is in its considerations of general demand scenarios where demand across periods can be correlated and/or the support of the demand vectors are not the same as the range of the uncertainty sets and/or for higher values of the service level.

6. Managerial Implications and Conclusion

The main contribution of our paper is deriving *closed-form* robust order quantities for a generic inventory management problem, which greatly facilitates their adoption into existing systems (e.g., MRP software). To derive these robust quantities, we only require the means and covariance matrix of demand to be known, but not the distributions. Our results are applicable to static and dynamic models with capacitated inventories. We believe our robust strategies can be successfully deployed in practice as a result of these minimal requirements. Indeed, a manager, using historical data, can easily estimate the means and covariance matrix of demand using standard spreadsheet software. In particular, we derive closed-form ordering quantities for correlated nonidentically distributed demand for both symmetric and asymmetric uncertainty sets, under capacitated inventory constraints, in both static and dynamic settings.

An intriguing result of our paper is that the symmetry of the uncertainty set is important. Our analytical results show that the symmetric uncertainty set $\Omega_{CLT}^{\text{symmetric}}$ drives qualitatively different behavior than the asymmetric set $\Omega_{CLT}^{\text{asymmetric}}$. The values of $\mu_i - \hat{\Gamma}_i \sigma_i$ determine the symmetry, as a result of the embedded constraints $\max\{\mu_i - \hat{\Gamma}_i \sigma_i, 0\} \leq d_i \leq \mu_i + \hat{\Gamma}_i \sigma_i$. Finally, note that $\hat{\Gamma}_i$ are parameters chosen by a user of our model to influence the degree of conservatism embedded in the robust optimization; therefore, a user can influence whether or not the uncertainty set is symmetric for a given set of μ_i and σ_i values. We provide advice for selecting appropriate values of $\hat{\Gamma}_i$.

We also present very encouraging computational results. While significantly reducing the computational time, our robust solution performs favorably compared to more complicated robust policies studied in the literature, which examine a similar problem and are shown to often outperform the dynamic programming-based policies. In particular, when demand across periods is correlated and the demand uncertainty sets are not the same as the support of demand, our robust solution leads to lower average costs for the majority of numerical scenarios considered in this paper. Our computational results are based on multivariate normal distributions of demand.

Finally, we discuss the assumptions of our basic model. We assume zero discounting, which turns out to be an innocuous assumption. Our model is applicable under discounting, and, if all costs are discounted at the same fixed rate, then our robust strategy will not change and will not depend on the discount rate. We have also assumed zero fixed ordering costs. Adding fixed costs into our model will likely change the structure of our robust strategy, potentially eliminating our ability to find closed-form solutions; we leave the resolution of this question to future research.

Acknowledgments

The first author gratefully acknowledges the support of a Premera Faculty Fellowship. The third author gratefully acknowledges the support of a Neal and Jan Dempsey Faculty Fellowship.

References

- Ardestani-Jaafari A, Delage E (2016) Robust optimization of sums of piecewise linear functions with application to inventory problems. *Oper. Res.* 64(2):474–494.
- Bandi C, Bertsimas D (2011) Network information theory via robust optimization. Working paper, Kellogg School of Management, Evanston, IL.
- Bandi C, Bertsimas D (2012) Tractable stochastic analysis in high dimensions via robust optimization. *Math. Programming* 134(1): 23–70.
- Bandi C, Bertsimas D (2014a) Optimal design for multi-item auctions: A robust optimization approach. *Math. Oper. Res.* 39(4): 1012–1038.
- Bandi C, Bertsimas D (2014b) Robust option pricing. *Eur. J. Oper. Res.* 239(3):842–853.
- Bandi C, Bertsimas D, Youssef N (2012) Tractable analysis of multi-server queues in the transient domain. Working paper, Kellogg School of Management, Evanston, IL.
- Bandi C, Bertsimas D, Youssef N (2015) Robust queueing theory. *Oper. Res.* 63(3):676–700.
- BBC News (2006) Wii shortages frustrating gamers. (December 8), <http://news.bbc.co.uk/1/hi/technology/6161717.stm>.
- Ben-Tal A, Nemirovski A (1998) Robust convex optimization. *Math. Oper. Res.* 23(4):769–805.
- Ben-Tal A, Nemirovski A (1999) Robust solutions to uncertain programs. *Oper. Res. Lett.* 25(1):1–13.
- Ben-Tal A, Nemirovski A (2000) Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming* 88(3):411–424.
- Ben-Tal A, El Ghaoui L, Nemirovski A (2009) *Robust Optimization* (Princeton University Press, Princeton, NJ).
- Ben-Tal A, Golany B, Nemirovski A (2005) Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing Service Oper. Management* 7(3):248–271.
- Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Programming* 99(2):351–376.
- Bertsimas D, Georghiou A (2015) Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Oper. Res.* 63(3):610–627.
- Bertsimas D, Sim M (2004) Price of robustness. *Oper. Res.* 52(1): 35–53.
- Bertsimas D, Thiele A (2006) A robust optimization approach to inventory theory. *Oper. Res.* 54(1):150–168.
- Bertsimas D, Gamarnik D, Rikun A (2011) Performance analysis of queueing networks via robust optimization. *Oper. Res.* 59(2): 455–466.
- Bertsimas D, Iancu D, Parrilo P (2010) Optimality of affine policies in multi-stage robust optimization. *Math. Oper. Res.* 35(2):363–394.
- Bienstock D, Özbay N (2008) Computing robust basestock levels. *Discrete Optim.* 5(2):389–414.
- Billingsley P (2012) *Probability and Measure*, Anniversary ed. (John Wiley & Sons, Hoboken, NJ).
- Chen F, Yu B (2005) Quantifying the value of leadtime information in a single-location inventory system. *Manufacturing Service Oper. Management* 7(2):144–151.
- Chen S, Lee H, Moizadeh K (2014) Coordination of a two-level supply chain with inventory subsidizing contracts. Working paper, University of Washington, Seattle.
- Chen X, Sim M, Sun P (2007a) A robust optimization perspective on stochastic programming. *Oper. Res.* 55(6):1058–1071.
- Chen X, Sim M, Simchi-Levi D, Sun P (2007b) Risk aversion in inventory management. *Oper. Res.* 55(5):828–842.

- Chen X, Sim M, Sun P, Zhang J (2008) A linear decision-based approximation approach to stochastic programming. *Oper. Res.* 56(2):344–357.
- Chung K (2001) *A Course in Probability Theory*, 3rd ed. (Academic Press, San Diego).
- Connors W, Cummins C (2011) RIM takes PlayBook hit. *Wall Street Journal* (December 3), <http://online.wsj.com/news/articles/SB10001424052970204012004577073932113176106>.
- Covert A (2013) Microsoft sinks 11% on earnings miss and huge Surface write-down. *CNN Money* (July 9), <http://money.cnn.com/2013/07/18/technology/microsoft-earnings/>.
- Eeckhoudt L, Gollier C, Schlesinger H (1995) The risk-averse (and prudent) newsboy. *Management Sci.* 41(5):786–794.
- Ehrhardt R (1979) The power approximation for computing (s, S) inventory policies. *Management Sci.* 25(8):777–786.
- El-Ghaoui L, Lebret H (1997) Robust solutions to least-square problems to uncertain data matrices. *SIAM J. Matrix Anal. Appl.* 18(4):1035–1064.
- El-Ghaoui L, Oustry F, Lebret H (1998) Robust solutions to uncertain semidefinite programs. *SIAM J. Optim.* 9(1):33–52.
- Feller W (1968) *An Introduction to Probability Theory and Its Applications*, 3rd ed., Vol. 1 (John Wiley & Sons, Hoboken, NJ).
- Gallego G, Moon I (1993) The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.* 44(8):825–834.
- Georghiou A, Wiesemann W, Kuhn D (2015) Generalized decision rule approximations for stochastic programming via liftings. *Math. Programming Ser. A* 152(1):301–338.
- Ghobbar A, Friend C (2003) Evaluation of forecasting methods for intermittent parts demand in the field of aviation: A predictive model. *Comput. Oper. Res.* 30(14):2097–2114.
- Gorissen B, Hertog D (2013) Robust counterparts of inequalities containing sums of maxima of linear functions. *Eur. J. Oper. Res.* 227(1):30–43.
- Hopp W, Spearman M (2011) *Factory Physics*, 3rd ed. (Waveland, Long Grove, IL).
- Iancu D, Sharma M, Sviridenko M (2013) Supermodularity and affine policies in dynamic robust optimization. *Oper. Res.* 61(4):941–956.
- Isaacs H (1963) Sensitivity of decisions to probability estimation errors. *Oper. Res.* 11(4):536–552.
- Lau H (1980) The newsboy problem under alternative optimization objectives. *J. Oper. Res. Soc.* 31(6):525–535.
- Levi R, Janakiraman G, Nagarajan M (2008) A 2-approximation algorithm for stochastic inventory control models with lost sales. *Math. Oper. Res.* 33(2):351–374.
- Levi R, Pal M, Roundy R, Shmoys D (2007) Approximation algorithms for stochastic inventory control models. *Math. Oper. Res.* 32(2):284–302.
- Moon I, Gallego G (1994) Distribution free procedures for some inventory models. *J. Oper. Res. Soc.* 45(6):651–658.
- Natarajan K, Sim M, Uichanco J (2008) Asymmetry and ambiguity in newsvendor models. Working paper, National University of Singapore, Singapore.
- Özer Ö, Wei W (2006) Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Sci.* 52(8):1238–1257.
- Perakis G, Roels G (2008) Regret in the newsvendor model with partial information. *Oper. Res.* 56(1):188–203.
- Rikun A (2011) Applications of robust optimization to queueing and inventory systems. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge.
- Scarf H (1958) A min-max solution of an inventory problem. Arrow K, Karlin S, Scarf H, eds. *Studies in the Mathematical Theory of Inventory and Production* (Stanford University Press, Stanford, CA), 201–209.
- Seattle Post-Intelligencer* (2008) Supply shortages hurt Xbox sales last month. (February 13), http://seattlepi.nwsource.com/business/351201_tbrfs14.html.
- See C, Sim M (2010) Robust approximation to multi-period inventory management. *Oper. Res.* 58(3):583–594.
- Simchi-Levi D, Kaminsky P, Simchi-Levi E (2008) *Designing and Managing the Supply Chain: Concepts, Strategies and Case Studies*, 3rd ed. (McGraw-Hill Irwin, New York).
- Soyster A (1973) Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.* 21(5):1154–1157.
- Tang C, Rajaram K, Alptekinoglu A, Ou J (2004) The benefits of advance booking discount programs: Model and analysis. *Management Sci.* 50(4):465–478.
- Van Mieghem J (2007) Risk mitigation in newsvendor networks: Resource diversification, flexibility, sharing, and hedging. *Management Sci.* 53(8):1269–1288.
- Wagner M (2010) Fully distribution-free profit maximization: The inventory management case. *Math. Oper. Res.* 35(4):728–741.
- Wagner M (2011) Online lot-sizing problems with ordering, holding and shortage costs. *Oper. Res. Lett.* 39(2):144–149.
- Williams T (1984) Stock control with sporadic and slow-moving demand. *J. Oper. Res. Soc.* 35(10):939–948.