A new robust cycle-based inventory control policy

Yupeng Chen[1] Garud Iyengar[2]

Abstract

In this paper, we propose a new robust cycle-based control policy for single installation inventory models with non-stationary uncertain demand. The policy is simple, flexible, easily implementable and preliminary numerical experiments suggest that the policy has very promising empirical performance. The policy can be used both when the excess demand is backlogged as well as when it is lost; with non-zero fixed ordering cost, and also when lead time is non-zero. The policy decisions are computed by solving a collection of linear programs *even* when there is a positive fixed ordering cost. The policy extends in a very simple manner to the joint pricing and inventory control problem.

1 Introduction

Single installation inventory management is a classical research topic which has received persistent attention from various research communities. The standard approach to compute the optimal inventory control policy uses dynamic programming (DP). For simple inventory control problems, DP-based methods are able to theoretically *characterize* the structure of the optimal policies, e.g., Scarf [36] showed that the optimal policy for an inventory control problem with fixed ordering cost is a so-called (s, S)-policy. Many other researchers have also used DP to derive structural results of the optimal policies under various model assumptions [26, 41]. Recently, Song and Zipkin [39] and Sethi and Cheng [37] have extended the structural results to inventory control problems with a more general Markovian demand. In practice, however, computing the optimal parameters is often challenging. Since DP-based methods compute a value function for each possible state in each period, the computational complexity of DP is prohibitive when state space is large. This *curse of dimensionality* has severely limited the application of DP in practice.

The DP methodology demands that the demand distribution be *perfectly* known. In practice, the available historical data on demand distribution is limited, and moreover, the demand distribution

[1]IEOR Department, Columbia University, New York, NY. Email: yc2561@columbia.edu

[2]IEOR Department, Columbia University, New York, NY. Email: garud@ieor.columbia.edu Research partially supported by ONR grant N000140310514, NSF Grant DMS 10-16571 and DOE Grant DE-FG02-08-25856.

is very often non-stationary. Thus, the perfect information assumption almost never holds in practice. The performance of DP-based policies is often very sensitive to errors in the demand distribution [24, 32]. See Bertsimas and Thiele [14] for an empirical investigation. Consequently, practitioners prefer simple policies that can be computed using limited information about the demand distribution [43]. Scarf [35] formulated a min-max single-period inventory problem where only the mean and the variance of the demand are known. Moon and Gallego extended this approach to single-period news-vendor problems [20] and to single-stage periodic review inventory models with a fixed reorder quantity [29]. Gallego et al. [21] later extended the Scarf model to a finite horizon model under the assumption that the demand is a discrete random variable. However, the complexity of computing the optimal policy is exponential in the time horizon.

Recent research efforts have focused on developing computationally tractable robust solutions for the dynamic program. One such approach is to use the robust optimization (RO) methodology to solve inventory control problems. RO is a relatively new methodology for tractably computing good solutions for optimization problems with uncertain parameters in which the parameters are assumed to belong to a known bounded *uncertainty* set and decisions are chosen assuming worst case behavior of these uncertain parameters [3, 5, 6, 7, 12, 13]. Bertsimas and Thiele [14] introduced a robust inventory control model where the uncertain parameters are the realized demand, and showed that for a family of polyhedral uncertainty sets introduced in [13] the open-loop control policy can be computed by solving a linear program (LP) (resp. mixed integer program (MIP)) when fixed ordering cost is zero (resp. positive). Bienstock and Ozbay [15] showed that the optimal base-stock policy for a closely related model can be computed efficiently using a Benders' decomposition approach. Ben-Tal et al. [4] showed that the optimal affine policy for a robust inventory model can be efficiently computed for a large class of uncertainty sets. Recently, Bertsimas, Iancu and Parrilo [11] showed that affine policies are optimal for hypercube-like uncertainty sets. They also showed how to compute the optimal affine policies, and more generally optimal polynomial polices [10].

The RO-based policies in the literature treat the inventory control problem as a special case of the more general problem of controlling a linear system. These policies typically do not exploit the special structure of inventory control problems. Since RO-based policies work only for linear dynamics, these policies are only applicable when the excess demand is fully backlogged, and not when the excess demand is lost. Moreover, the ordering decisions in the RO-based policies for inventory control problems with a positive fixed ordering cost have to be computed by solving an MIP. Since these policies are typically implemented in a rolling horizon fashion, one needs to solve several, typically large, MIPs. MIP solvers are neither as ubiquitous nor as stable as LP solvers; consequently, MIP-based control polices are less likely to be widely adopted.

In this paper, we propose a new RO-based control policy that is inspired in part by the Economic-Order-Quantity (EOQ) model [43] and the periodic review (R, T) or (S, T)-type policies [2, 33, 38, 43]. The EOQ model is an infinite horizon, continuous review inventory model with constant deterministic demand. In order to balance the inventory ordering cost, which includes a fixed cost, and the inventory holding cost and backlogging cost, EOQ optimal policy splits the planning horizon into fixed length *cycles* and chooses the cycle length and the order quantity jointly to minimize the average cost over a cycle. The periodic review (R, T)-type policies generalize the EOQ model to the case where the demand is a stationary stochastic process, specifically Poisson process. Under an (R, T) policy, the inventory position is reviewed every T periods, i.e., the planning horizon is split into length T cycles, and, if required, the policy raises the inventory position up to R. As in the EOQ model, in the (R, T) policy the cycle length T is an explicit decision variable. Although the (R, T) policy is generally sub-optimal, it is very attractive in practice because it specifies the reorder epochs upfront and is, thus, easy to implement. In particular, the (R, T) policy can substantially simplify the coordination within a more complex inventory system [33, 43]. Rao [33] showed that the long-run average cost function for the (R, T) policy is jointly convex in R and T, and provided performance guarantee for the (R, T) policy. See [2, 38] for a detailed study of the empirical performance of the (R, T) policy and further applications of the policy.

The EOQ model and the (R, T) policy are simple, yet powerful, control policies for a single installation inventory system. However, these policies have so far been restricted to models where the goal is to minimize the long-run average cost and the demand process is either stationary, deterministic, i.e., constant, or stationary, stochastic. Our RO-based approach extends the "cycle-based" perspective to a single installation, finite horizon inventory model with non-stationary uncertain demand. Specifically, the contributions of this paper are as follows.

(a) We propose a new robust cycle-based control policy for a single installation, finite horizon inventory model with non-stationary uncertain demand. Our policy dynamically breaks the planning horizon into *ordering cycles* and chooses an appropriate order quantity and cycle length at the beginning of each ordering cycle assuming the worst case behavior of the uncertain demand. Thus, our policy can be viewed as an *adaptive* version of the (R, T) policy for an inventory model with more general demand process. The policy has a simple intuitive economic interpretation and is very easy to implement. Consequently, we believe that our proposed policy has a higher likelihood of being adopted in practice.

(b) We show that our policy provides a unified and efficient approach for solving inventory control problems under various model assumptions. In particular, the decisions in our proposed robust cycle-based policy can be efficiently computed both when the excess demand is backlogged or lost. To the best of our knowledge, this is the first RO-based inventory control policy for the lost-sales model. The decisions are also efficiently computable when there is a non-zero fixed ordering cost and also when lead time is non-zero. More specifically, we show that for all these cases the decisions in our policy can be computed by solving a collection of LPs of modest size even when the fixed ordering cost is non-zero, i.e., we do not have to solve any MIPs. Consequently, the computation of our policy is extremely efficient (see Section 6). On the other hand, since LP solvers are significantly more easily accessible and stable when compared to MIP solvers, our policy should be an attractive candidate for practical implementation.

    Our policy also extends to the joint pricing and inventory control problem. In this case we need to solve a small collection of convex problems to compute the optimal decisions. These problems reduce to QPs in special cases.

(c) The robust cycle-based policy has very promising empirical performance. We conducted an extensive numerical experiment to test the performance of our proposed simple policy. We compared the performance of our policy with the rolling-horizon version of the robust policy in [14] when the excess demand is backlogged. For the lost-sales model, we compared the performance of our policy with a DP-based policy, the best base-stock policy in hindsight, and the set of heuristic policies considered in Zipkin [44]. In the joint pricing and inventory control model, we compared the performance of the simple extension of our policy with a DP-based policy. For the backlogging model, the lost-sales model without lead time and the joint pricing and inventory control model, the cost of the robust cycle-based policy is typically within ±3% of the cost of benchmark methods. For the lost-sales model with positive lead times, our policy incurs no more than 7% of the cost of the optimal policy in most scenarios and is very competitive with other heuristics. In most cases, the computational time required to compute the decisions in our policy is at least 5–10 times smaller than that required for the competing methods.

The rest of the paper is organized as follows. In Section 2, we introduce our basic inventory model with zero lead time. In Section 3, we propose our robust cycle-based inventory control policy, and generalize it to the case with non-zero lead time in Section 4. We briefly discuss the extension to a joint pricing and inventory control model in Section 5. In Section 6, we report the preliminary numerical results. In Section 7 we include some concluding remarks and avenues for further extensions.

2 Inventory model with zero lead time

We consider a single product single installation stochastic inventory problem over a finite discrete horizon of T periods. Let x_k denote the inventory available at the beginning of period k, u_k denote the quantity ordered at the beginning of period k, and d_k denote the realized demand in period k, for $k = 1, 2, \dots, T$. We assume that the lead time is zero, i.e., the quantity ordered at the beginning of any period is delivered immediately. In particular, we assume that the quantity u_k is delivered before the demand d_k is realized. We discuss constant positive lead time in Section 4.

Dynamics We consider the following two most prevalent system dynamics [9, 43]:

(a) Backlogging: The demand in excess of the inventory in any period is backlogged and carried over into next period, i.e., $x_{k+1} = x_k + u_k - d_k$, for all $1 \leq k \leq T - 1$. Note that x_k denotes the net inventory and can take negative values.

(b) Lost-sales: The demand in excess of the inventory in any period is permanently lost, i.e., $x_{k+1} = \max\{x_k + u_k - d_k, 0\}$, for all $1 \leq k \leq T - 1$. Here, x_k denotes the actual inventory at hand and must be nonnegative.

We assume that the installation has a positive capacity M (possibly $+\infty$). Since u_k arrives before the demand, d_k, is realized, we require that $x_k + u_k \leq M$, for all $1 \leq k \leq T$.

Cost function The ordering cost in period k is given by $K\mathbf{1}(u_k > 0) + c \cdot u_k$, where K denotes the fixed ordering cost, c denotes the per unit variable ordering cost, and $\mathbf{1}(\cdot)$ denotes an indicator function that takes value 1 when argument is true, and zero otherwise. The inventory holding cost at the end of period k is given by $h \cdot \max\{x_k + u_k - d_k, 0\}$, where h denotes the per unit holding cost, and the shortage cost at the end of period k is given by $b \cdot \max\{-(x_k + u_k - d_k), 0\}$, where b denotes the per unit shortage cost; i.e., the total inventory holding and shortage cost in period k is given by $\max\Big\{h \cdot (x_k + u_k - d_k), -b \cdot (x_k + u_k - d_k)\Big\}$. Thus, the total cost incurred in period k is given by $C_k = K\mathbf{1}(u_k > 0) + c \cdot u_k + \max\Big\{h \cdot (x_k + u_k - d_k), -b \cdot (x_k + u_k - d_k)\Big\}$.

Demand distribution We assume that the demand process $\mathfrak{D} = (D_1, D_2, \ldots, D_T)$ is a stochastic process that satisfies the following assumptions.

Assumption 1. The random variables $\{D_k\}_{k=1}^{T}$ are mutually independent, but not necessarily identically distributed.

Assumption 2. Only the first two moments of the distribution are known, i.e., only $\mathrm{E}(D_k) = \bar{d}_k$ and $\mathrm{var}(D_k) = \sigma_k^2$ are known for all k; all other moments and the distribution are unknown.

Assumption 1 is common in inventory literature. It allows demand sequence to be non-stationary, e.g., seasonally cyclic. Assumption 2 is a departure from the classical inventory literature where one assumes perfect information about the demand distributions. The perfect information assumption leads to a tractable inventory control problem; however, it rarely holds in practice. In most practical settings, the inventory manager may have sufficient data to accurately estimate the first two moments of the demand distribution; however, higher moments or the shape of the distribution cannot be estimated accurately [21].

We denote the set of all distributions for the demand process $\mathfrak{D}$ which are consistent with Assumption 1-2 by the set $\mathcal{F}$. We denote a realization of $\mathfrak{D}$ by the lower case vector as $\mathbf{d} = (d_1, d_2, \ldots, d_T)$. For $1 \leq k \leq T$, $\mathbf{d}[k]$ denotes the sub-sequence $(d_k, d_{k+1}, \ldots, d_T)$, and for $1 \leq j \leq T - k + 1$, $\mathbf{d}[k, j]$ denotes the sub-sequence $(d_k, d_{k+1}, \ldots, d_{k+j-1})$.

Optimality criterion An inventory control policy $\pi = \big\{\pi_k : k \in \{1, 2, \ldots, T\}\big\}$ consists of ordering functions π_k that prescribe the order quantity u_k for period k. We will call a policy π admissible if π_k is non-anticipatory, i.e., the function π_k is measurable with the respect to the filtration generated by the demand up to (and including) period $k - 1$. Let Π denote the set of all admissible policies. The cost $J(\pi)$ of an admissible policy π is defined as

$$J(\pi) = \sup_{\mathfrak{D} \in \mathcal{F}} \mathbf{E}_{\mathfrak{D}} \Big[\sum_{k=1}^{T} \Big(K\mathbf{1}(\pi_k > 0) + c \cdot \pi_k + \max\big\{ h \cdot (x_k + \pi_k - d_k), -b \cdot (x_k + \pi_k - d_k) \big\} \Big) \Big], \quad (1)$$

where the expectation is taken with respect to the distribution of $\mathbf{d}$. We want to characterize $J^* = \inf_{\pi \in \Pi} J(\pi)$, and identify an optimal policy π^* that achieves J^*.

The optimality criterion (1) and the associated optimization problem J^* is not new. Scarf [35] considers a single period version of J^*. Gallego et al. [21] consider J^* for the setting with discrete demand distribution and backlogging of excess demand. They show that the optimal policy is of the (s, S)-type, and provide a dynamic programming (DP) recursion to compute the optimal parameters; the solution algorithm is, however, exponential in T.

Robust inventory control Robust inventory control was introduced in Bertsimas and Thiele [14] (see also [4, 15] and references therein). We briefly review the RO-based approach proposed in [14]. Instead of using (1), Bertsimas and Thiele define the *robust* cost $J_r(\pi)$ for policy π as

$$J_r(\pi) = \sup_{\mathbf{d} \in \mathcal{D}} \Big[\sum_{k=1}^{T} \Big(K\mathbf{1}(\pi_k > 0) + c \cdot \pi_k + \max\big\{ h \cdot (x_k + \pi_k - d_k), -b \cdot (x_k + \pi_k - d_k) \big\} \Big) \Big], \quad (2)$$

where $\mathcal{D}$ is an *uncertainty set* of demand realizations, i.e., in the robust approach the uncertainty set $\mathcal{F}$ of demand *distributions* is replaced by an uncertainty set $\mathcal{D}$ of demand *realizations.* The uncertainty set $\mathcal{D}$ is constructed using the available moment information and the optimal control policy is computed by solving min-max problem $J_r^* = \inf_{\pi \in \Pi} J_r(\pi)$. The min-max problem J_r^* is still a hard problem. By restricting attention to open-loop control policies, Bertsimas and Thiele are able to reformulate an appropriate relaxation of J_r^* as a mathematical program and solve it efficiently. They show that the empirical performance of the robust policy is close to optimal for the original min-max optimality criterion J^*.

3 A robust cycle-based inventory control policy

In this section, we propose a simple robust cycle-based inventory control policy that provides a good solution for J^*. Intuitively, we expect our policy to show good *ex post* average cost across different demand distributions. We first motivate the main conceptual insight underlying our policy; next, we show how to compute the policy decisions in detail. We also show that the policy presented here can be applied to both backlogging dynamics and lost-sales dynamics.

Cycle-based policies in inventory management Consider an infinite horizon inventory control problem where the goal is to minimize the long-run average cost for flat deterministic demand, i.e., $D_k = d$ for all k. A good inventory control policy has to "balance" the ordering cost with other costs. On one extreme is the policy that orders the quantity d in each period; thereby, incurring no inventory holding and shortage cost but a large fixed ordering cost. A policy that places larger but less frequent orders incurs a smaller fixed ordering cost; but the policy either incurs inventory holding cost, or shortage cost, or both. The optimal policy breaks the horizon into fixed length *cycles* where an order is placed in the first period of each cycle. For any given cycle, the cycle length and the order quantity is chosen to minimize the *average* cost over the cycle. This type of analysis underlies the well known Economic-Order-Quantity (EOQ) model [43] and extends to the (R, T)-type policies with stationary stochastic demand processes [2, 33, 38, 43].

Drawing insights from the EOQ model and the (R, T)-type policies, we propose a robust cycle-based policy that is tailored for the finite horizon inventory model with non-stationary uncertain demand. Without loss of generality, assume the initial inventory $x_1 = 0$. At the beginning of period $\tau_1 = 1$, we make the following two decisions: order quantity u_{τ_1} and the *cycle length* ξ_1. The order quantity u_{τ_1} and the cycle length ξ_1 is chosen to minimize an appropriately defined worst-case average cost over the *ordering cycle* $(\tau_1, \ldots, \tau_1 + \xi_1 - 1)$. We do not order until period $\tau_2 = \tau_1 + \xi_1$. At the beginning of period τ_2, we choose order quantity u_{τ_2} and the next cycle length ξ_2 that minimizes the worst-case average cost over the next ordering cycle $(\tau_2, \ldots, \tau_2 + \xi_2 - 1)$. We proceed in this manner until the end of the planning horizon. In the rest of this section, we precisely define the worst-case average cost that defines our policy and show that the decisions (u_τ, ξ) can be computed efficiently in an *online* fashion.

Note that our proposed policy is different from the (r, q)-type policies [43]. In the (r, q) policy the decision maker chooses the reorder point r and order quantity q, and a new cycle is triggered when the inventory position hit the reorder point; thus, the cycle length is a function of r, q and the demand dynamics. Consequently, the cycle length of an (r, q) policy cannot be easily computed when the demand distribution is uncertain. In our cycle-based policy, the cycle length is an explicit decision variable, and once the cycle length is determined at the beginning of cycle, it is held constant. This alternative formulation leads to tractable approximation for both backlogging and lost-sales dynamics, even in the presence of positive lead time and uncertain demand. Moreover, the (r, q) policy was proposed for a continuous review inventory model with stationary demand, and the decision maker computes the optimal *stationary* pair (r^*, q^*) *offline*. In contrast, our policy computes (u_τ, ξ) in an *online* fashion and the (u_τ, ξ) are adapted to the sample path of demands, and are generally different across cycles. Thus, our cycle-based policy is capable of handling non-stationary demand.

Convexity of inventory holding and shortage costs Consider an ordering cycle $(\tau, \tau + 1, \ldots, \tau + \xi - 1)$ that starts from period τ and has cycle length ξ. Let x denote available inventory at the beginning of the ordering cycle, u denote the order placed at the beginning of the cycle, and let $H(\tau, x, \xi, u, \mathbf{d}[\tau, \xi])$ denote the inventory holding and shortage cost incurred over the ordering cycle as a function of τ, x, ξ, u and the realized demand sequence $\mathbf{d}[\tau, \xi]$.

Lemma 1. For fixed $\big(\tau, x, \xi, \mathbf{d}[\tau,\xi]\big)$, $H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi])$ is a piecewise linear convex function of u with $\xi + 1$ linear pieces.

Proof. First consider backlogging dynamics. Since $u_t = 0$, for all $\tau + 1 \le t \le \tau + \xi - 1$, $x_t = x + u - \sum_{\ell=\tau}^{t-1} d_\ell$, for $\tau + 1 \le t \le \tau + \xi - 1$; thus, the inventory holding and shortage cost incurred over the cycle is given by

$$H\big(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]\big) = \sum_{t=\tau}^{\tau+\xi-1} \max\Big\{h\big(x + u - \sum_{\ell=\tau}^{t} d_\ell\big), -b\big(x + u - \sum_{\ell=\tau}^{t} d_\ell\big)\Big\}. \tag{3}$$

Since $x + u - \sum_{\ell=\tau}^{t} d_\ell$ is affine function of u, $\max\big\{h(x + u - \sum_{\ell=\tau}^{t} d_\ell), -b(x + u - \sum_{\ell=\tau}^{t} d_\ell)\big\}$, and consequently, H is a piecewise linear convex function of u. From (3) it follows that

(i) $u \in (-\infty, d_\tau - x]$: $H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = f^0(u) \triangleq -\sum_{t=\tau}^{\tau+\xi-1} b(x + u - \sum_{\ell=\tau}^{t} d_\ell)$.

(ii) $u \in \Big[\sum_{\ell=\tau}^{\tau+r-1} d_\ell - x, \sum_{\ell=\tau}^{\tau+r} d_\ell - x\Big]$, for $1 \le r \le \xi - 1$:
$H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = f^r(u) \triangleq \sum_{t=\tau}^{\tau+r-1} h(x + u - \sum_{\ell=\tau}^{t} d_\ell) - \sum_{t=\tau+r}^{\tau+\xi-1} b(x + u - \sum_{\ell=\tau}^{t} d_\ell)$.

(iii) $u \in \Big[\sum_{\ell=\tau}^{\tau+\xi-1} d_\ell - x, +\infty\Big)$: $H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = f^\xi(u) \triangleq \sum_{t=\tau}^{\tau+\xi-1} h(x + u - \sum_{\ell=\tau}^{t} d_\ell)$.

From this partition, it follows that for all $u \ge 0$,

$$H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = \max\Big\{f^0(u), f^1(u), \ldots, f^\xi(u)\Big\}. \tag{4}$$

Next, we consider the case with lost-sales dynamics. Since $u_t = 0$, for all $\tau + 1 \le t \le \tau + \xi - 1$, $x_t = \max\{x + u - \sum_{\ell=\tau}^{t-1} d_\ell, 0\}$, for $\tau + 1 \le t \le \tau + \xi - 1$. Therefore,

$$\begin{aligned} H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) &= \max\Big\{h(x + u - d_\tau), -b(x + u - d_\tau)\Big\} \\ &\quad + \sum_{t=\tau+1}^{\tau+\xi-1} \max\Bigg\{h\Big(\max\{x + u - \sum_{\ell=\tau}^{t-1} d_\ell, 0\} - d_t\Big), -b\Big(\max\{x + u - \sum_{\ell=\tau}^{t-1} d_\ell, 0\} - d_t\Big)\Bigg\}. \end{aligned} \tag{5}$$

From (5), it follows that

(i) $u \in (-\infty, d_\tau - x]$: $H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = g^0(u) \triangleq -b(x + u - \sum_{t=\tau}^{\tau+\xi-1} d_t)$.

(ii) $u \in \left[\sum_{\ell=\tau}^{\tau+r-1} d_\ell - x, \sum_{\ell=\tau}^{\tau+r} d_\ell - x\right]$, for $1 \le r \le \xi - 1$:
$H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = g^r(u) \triangleq \sum_{t=\tau}^{\tau+r-1} h\big(x + u - \sum_{\ell=\tau}^{t} d_\ell\big) - b\big(x + u - \sum_{t=\tau}^{\tau+\xi-1} d_t\big)$.

(iii) $u \in \left[\sum_{\ell=\tau}^{\tau+\xi-1} d_\ell - x, +\infty\right)$: $H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = g^\xi(u) \triangleq \sum_{t=\tau}^{\tau+\xi-1} h(x + u - \sum_{\ell=\tau}^{t} d_\ell)$.

The continuity of H is easy to verify. Since $\frac{dg^j}{du} = jh - b$, for $0 \le j \le \xi - 1$, $\frac{dg^\xi}{du} = \xi h$, i.e., the slope of g^j is non-decreasing in the index j, it follows that H is convex in u. In particular, we have the following representation for H on $\mathbb{R}_+$: for all $u \ge 0$,

$$H(\tau, x, \xi, u, \mathbf{d}[\tau,\xi]) = \max\left\{g^0(u), g^1(u), \dots, g^\xi(u)\right\}. \tag{6}$$

□

Robust optimization approach for computing (u_τ, ξ) Let τ denote the first period of a new ordering cycle. Let x_τ denote the available inventory and $\mathbf{d}[1, \tau-1]$ denote the past demand realization. By Assumption 1, $\mathbf{d}[1, \tau-1]$ does not provide any new information about the future demand $(D_\tau, D_{\tau+1}, \dots, D_T)$. Following [14], we introduce the uncertainty set

$$\mathcal{D}[\tau] = \left\{\mathbf{d}[\tau] \middle| \bar{\omega}_t - \hat{\omega}_t \le d_t \le \bar{\omega}_t + \hat{\omega}_t, \quad \sum_{\ell=\tau}^{t} \frac{|d_\ell - \bar{\omega}_\ell|}{\hat{\omega}_\ell} \le \Gamma_{t-\tau+1}^{\tau}, \quad \forall\, \tau \le t \le T\right\}, \tag{7}$$

where $\bar{\omega}_t$ and $\hat{\omega}_t$ are, respectively, the *nominal value* and *deviation* of d_t, and are computed from $\{(\bar{d}_t, \sigma_t^2) : \tau \le t \le T\}$. In [14], $\bar{\omega}_t = \bar{d}_t$, $\hat{\omega}_t = \min\{2\sigma_t, \bar{\omega}_t\}$. The *budget of uncertainty* constraint $\sum_{\ell=\tau}^{t} \frac{|d_\ell - \bar{\omega}_\ell|}{\hat{\omega}_\ell} \le \Gamma_{t-\tau+1}^{\tau}$ was introduced in [12, 13]. It rules out large deviations in the cumulative demand; and, therefore, controls conservativeness of the robust constraints. Note that the uncertainty set $\mathcal{D}[\tau]$ defined in (7) is symmetric about the nominal demand $\bar{\omega}_t$. When the demand distribution is asymmetric about the mean, an asymmetric uncertainty set may be more appropriate. We consider such extensions in Section 6.2.2 and show that asymmetric uncertainty sets lead to significantly improved performance when the demand distribution is heavy tailed.

The worst-case average cost $F(u_\tau, \xi)$ over the cycle $(\tau, \tau+1, \ldots, \tau+\xi-1)$ is defined as follows:

$$F(u_\tau, \xi) = \max_{\mathbf{d}[\tau] \in \mathcal{D}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau > 0) + cu_\tau + H\Big(\tau, x_\tau, \xi, u_\tau, \mathbf{d}[\tau, \xi]\Big)}{\xi} \right\}, \tag{8}$$

where H denotes the total inventory holding and shortage costs over the cycle. The optimal decision (u_τ^*, ξ^*) is chosen by solving the optimization problem

$$\min_{\xi \in \left\{1,2,\ldots,\min\{T-\tau+1, U\}\right\}} \min_{0 \le u_\tau \le M - x_\tau} \Big\{F(u_\tau, \xi)\Big\}, \tag{9}$$

where U is a user-defined parameter. Setting $U = \infty$ allows the cycle length ξ to be set to any positive integer less than or equal to the number of remaining periods in the planning horizon. In our numerical experiments, we found that $U = 12$ was sufficient for very good performance. Note that in (8) and (9), τ and x_τ are known constants, i.e., part of the inputs.

It is clear that the traditional myopic policy is a special case of our policy with $\xi = 1$. The following example highlights the benefits of the flexibility of allowing longer cycle lengths. Consider a lost-sales inventory model with deterministic demand $d_t = \bar{d} = 100$, $T = 10$, $K = 1000$, $c = h = 0$, $b = 5$ and $x_1 = 0$. Since $K > b\bar{d}$, the myopic policy never places an order, and incurs a total cost of 5000. In contrast, our policy selects $\xi = 10$, $u_1 = 1000$, and incurs a total cost of 1000.

Solution algorithm To solve (9), we fix ξ, solve the inner minimization problem $\phi(\xi) = \min_{0 \le u_\tau \le M - x_\tau} F(u_\tau, \xi)$, and then enumerate all feasible value of ξ. Since the number of feasible ξ is at most U, and we show below that $\phi(\xi)$ can be solved efficiently, an optimal solution to (9) can be computed efficiently. The optimization problem $\phi(\xi)$ is equivalent to

$$\min_{0 \le u_\tau \le M - x_\tau} \max_{\mathbf{d}[\tau] \in \mathcal{D}[\tau]} \Big\{K\mathbf{1}(u_\tau > 0) + cu_\tau + H\big(\tau, x_\tau, \xi, u_\tau, \mathbf{d}[\tau, \xi]\big)\Big\}. \tag{10}$$

Since $\big(\tau, x_\tau, \xi\big)$ are known constants in (10), we abbreviate $H\big(\tau, x_\tau, \xi, u_\tau, \mathbf{d}[\tau, \xi]\big)$ as $H\big(u_\tau, \mathbf{d}[\tau, \xi]\big)$. Lemma 2 reformulates (10) into a more tractable form.

Lemma 2. The optimization problem (10) is equivalent to the following optimization problem:

$$\min \left\{ \begin{array}{c} \min_{0 \le u_\tau \le M - x_\tau} \max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} \Big\{ K + cu_\tau + H\big(u_\tau, \mathbf{d}[\tau,\xi]\big) \Big\}, \\ \max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} H\big(0, \mathbf{d}[\tau,\xi]\big) \end{array} \right\}, \tag{11}$$

where $\mathcal{D}[\tau,\xi]$ denotes the projection of the set $\mathcal{D}[\tau]$ to the space spanned by $\mathbf{d}[\tau,\xi]$ and is given by

$$\mathcal{D}[\tau,\xi] = \Big\{ \mathbf{d}[\tau,\xi] \Big| \bar{\omega}_t - \hat{\omega}_t \le d_t \le \bar{\omega}_t + \hat{\omega}_t, \quad \sum_{\ell=\tau}^{t} \frac{|d_\ell - \bar{\omega}_\ell|}{\hat{\omega}_\ell} \le \Gamma^{\tau}_{t-\tau+1}, \quad \forall\, \tau \le t \le \tau + \xi - 1 \Big\}. \tag{12}$$

Proof. Since $\mathbf{d}[\tau,\xi]$ is a sub-vector of $\mathbf{d}[\tau]$, it follows that

$$\max_{\mathbf{d}[\tau] \in \mathcal{D}[\tau]} \Big\{ K\mathbf{1}(u_\tau > 0) + cu_\tau + H\big(u_\tau, \mathbf{d}[\tau,\xi]\big) \Big\} = \max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} \Big\{ K\mathbf{1}(u_\tau > 0) + cu_\tau + H\big(u_\tau, \mathbf{d}[\tau,\xi]\big) \Big\}.$$

Thus, (10) is equivalent to

$$\min_{0 \le u_\tau \le M - x_\tau} \max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} \Big\{ K\mathbf{1}(u_\tau > 0) + cu_\tau + H\big(u_\tau, \mathbf{d}[\tau,\xi]\big) \Big\}. \tag{13}$$

To account for the discontinuity of the cost function at $u_\tau = 0$, we split (13) into two terms and take the minimum:

$$\min \left\{ \begin{array}{c} \min_{0 < u_\tau \le M - x_\tau} \max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} \Big\{ K + cu_\tau + H\big(u_\tau, \mathbf{d}[\tau,\xi]\big) \Big\}, \\ \max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} H\big(0, \mathbf{d}[\tau,\xi]\big) \end{array} \right\}.$$

(11) follows from observing that $\max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} K + H\big(0, \mathbf{d}[\tau,\xi]\big) \ge \max_{\mathbf{d}[\tau,\xi] \in \mathcal{D}[\tau,\xi]} H\big(0, \mathbf{d}[\tau,\xi]\big)$. □

Next, we show how to efficiently solve (11). We focus on the first program and show that the solution to the second program is a byproduct of solving the first. For notational simplicity, we drop all the subscripts appear in (11), and let $u = u_\tau$, $x = x_\tau$, $\mathbf{d} = \mathbf{d}[\tau,\xi]$, and $\mathcal{D} = \mathcal{D}[\tau,\xi]$. Then the first program in (11) can be written as

$$\min_{0 \le u \le M - x} \max_{\mathbf{d} \in \mathcal{D}} \Big\{ K + cu + H(u, \mathbf{d}) \Big\}. \tag{14}$$

1. **Initialize:** $\tilde{D} = \left\{ \tilde{\mathbf{d}} = (\bar{\omega}_{\tau}, \bar{\omega}_{\tau+1}, \ldots, \bar{\omega}_{\tau+\xi-1}) \right\}$, $LB = 0$ and $UB = +\infty$.

2. **Inventory manager's problem:** Solve the inventory manager's problem with $\mathbf{d} \in \tilde{D}$, i.e., the adversary is restricted to the set of demands in the working list $\tilde{D}$. Let

$$\tilde{u} \in \underset{0 \leq u \leq M-x}{\text{argmin}} \; \max_{\mathbf{d} \in \tilde{D}} K + cu + H(u, \mathbf{d}). \tag{15}$$

   Set $LB \leftarrow \max_{\mathbf{d} \in \tilde{D}} K + c\tilde{u} + H(\tilde{u}, \mathbf{d})$.

3. **Adversary's problem:** Solve the worst case demand sequence $\tilde{\mathbf{d}}$ corresponding to the policy $\tilde{u}$ computed in Step 2, i.e., compute

$$\tilde{\mathbf{d}} \in \underset{\mathbf{d} \in \mathcal{D}}{\text{argmax}} \, K + c\tilde{u} + H(\tilde{u}, \mathbf{d}). \tag{16}$$

   Note that $\tilde{\mathbf{d}}$ is selected from the full set $\mathcal{D}$ and not the current working list $\tilde{D}$. Set $UB \leftarrow \min\left(UB, K + c\tilde{u} + H(\tilde{u}, \tilde{\mathbf{d}})\right)$.

4. **Termination test and update:** If $UB - LB$ is small enough, then **EXIT**. Otherwise, set $\tilde{D} \leftarrow \tilde{D} \cup \{\tilde{\mathbf{d}}\}$ and return to Step 2.

Figure 1: Benders' decomposition

Since Lemma 1 shows that $cu + H(u, \mathbf{d})$ is convex in u, (14) is a convex optimization problem. We solve (14) using the Benders' decomposition approach [8, 15] displayed in Figure 1. In this procedure, one iteratively solves the inventory manager's problem that selects a candidate order quantity $\tilde{u}$ and the adversary's problem that computes the worst-case demand sequence $\tilde{\mathbf{d}} \in \mathcal{D}$ corresponding to $\tilde{u}$ that is added to the working set $\tilde{D}$ of candidate demand sequences for the inventory manager's problem. For this procedure to converge quickly, the inventory manager's problem and the adversary's problem should be efficiently solvable, and, in addition, the total number of iterations should be small. Since (14) is a one-dimensional convex optimization problem in u, and Benders' decomposition is a sub-gradient algorithm for solving (14), we expect the algorithm to converge quickly, and, indeed in our numerical experiments, we find that in nearly all cases the Benders' decomposition requires at most 5 iterations to guarantee that $\frac{UB-LB}{LB} < 10^{-5}$. In the rest of this section we show that both the inventory manager's problem and the adversary's problem can be solved efficiently.

Lemma 3. The inventory manager's problem (15) in Figure 1 can be reformulated as an LP with

2 decision variables and $(\xi+1)|\tilde{D}|+2$ constraints, where $|\tilde{D}|$ denotes the cardinality of the set $\tilde{D}$.

Proof. Let $\tilde{D} = \{\mathbf{d}^1, \mathbf{d}^2, \ldots, \mathbf{d}^{|\tilde{D}|}\}$, where $\mathbf{d}^j = (d^j_\tau, d^j_{\tau+1}, \ldots, d^j_{\tau+\xi-1})$. It is clear that (15) is equivalent to $\min\Big\{\ t \ \big| u \in [0, M-x], \quad t \geq K + cu + H(u, \mathbf{d}^j), \quad \forall\, j = 1, 2, \ldots, |\tilde{D}|\Big\}$, where (u, t) are the decision variables. Recall that Lemma 1 implies that $H(u, \mathbf{d}^j)$ is the maximum of $\xi+1$ linear functions of u; therefore, it follows that for each $1 \leq j \leq |\tilde{D}|$, constraint $t \geq K + cu + H(u, \mathbf{d}^j)$ can be reformulated as $\xi+1$ linear constraints in (u, t). □

The inventory manager's problem (15) can be reduced to an LP with no assumption on the uncertainty set $\mathcal{D}$. Next, we show that the adversary's problem (16) can be solved efficiently. The key insight in establishing the result is that the inventory level is non-increasing over a cycle.

Lemma 4. The adversary's problem (16) can be reduced to solving at most $\xi+1$ LPs.

Proof. Since the term $K + c\tilde{u}$ is fixed constant, the adversary's problem is equivalent to $\max_{\mathbf{d}\in\mathcal{D}} H(\tilde{u}, \mathbf{d})$. We first consider the case with backlogging dynamics. Consider the following two cases:

(i) $x + \tilde{u} \leq 0$: Since $\mathbf{d} \geq 0$, $x + \tilde{u} \leq 0$ implies $x + \tilde{u} - \sum_{\ell=\tau}^{t} d_\ell \leq 0$, for all $\tau \leq t \leq \tau + \xi - 1$. Thus, $H(\tilde{u}, \mathbf{d}) = \sum_{t=\tau}^{\tau+\xi-1} \max\Big\{h(x + \tilde{u} - \sum_{\ell=\tau}^{t} d_\ell), -b(x + \tilde{u} - \sum_{\ell=\tau}^{t} d_\ell)\Big\} = -\sum_{t=\tau}^{\tau+\xi-1} b(x + \tilde{u} - \sum_{\ell=\tau}^{t} d_\ell)$. Thus

$$\max_{\mathbf{d}\in\mathcal{D}} H(\tilde{u}, \mathbf{d}) = \max_{\mathbf{d}\in\mathcal{D}} \Big\{ -\sum_{t=\tau}^{\tau+\xi-1} b\big(x + \tilde{u} - \sum_{\ell=\tau}^{t} d_\ell\big)\Big\}. \tag{17}$$

Since $\mathcal{D}$ is a polyhedron, (17) is LP. In this case, we only have to solve a single LP.

(ii) $x + \tilde{u} > 0$: It is clear that

$$\max_{\mathbf{d}\in\mathcal{D}} H(\tilde{u}, \mathbf{d}) = \max_{r\in\{0,1,\ldots,\xi\}} \Big\{ \max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u}, \mathbf{d})\Big\}, \tag{18}$$

where

$$\mathcal{D}^r = \begin{cases} \Big\{\mathbf{d} \in \mathcal{D} \mid x + \tilde{u} - d_\tau \leq 0\Big\}, & r = 0, \\ \Big\{\mathbf{d} \in \mathcal{D} \mid x + \tilde{u} - \sum_{\ell=\tau}^{\tau+r-1} d_\ell \geq 0,\ x + \tilde{u} - \sum_{\ell=\tau}^{\tau+r} d_\ell \leq 0\Big\}, & 1 \leq r \leq \xi - 1, \\ \Big\{\mathbf{d} \in \mathcal{D} \mid x + \tilde{u} - \sum_{\ell=\tau}^{\tau+\xi-1} d_\ell \geq 0\Big\}, & r = \xi. \end{cases} \tag{19}$$

Fix r. Then

$$\begin{aligned}\max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u},\mathbf{d}) &= \max_{\mathbf{d}\in\mathcal{D}^r}\left\{\sum_{t=\tau}^{\tau+\xi-1}\max\left\{h(x+\tilde{u}-\sum_{\ell=\tau}^{t}d_\ell), -b(x+\tilde{u}-\sum_{\ell=\tau}^{t}d_\ell)\right\}\right\}\\ &= \max_{\mathbf{d}\in\mathcal{D}^r}\left\{\sum_{t=\tau}^{\tau+r-1}h(x+\tilde{u}-\sum_{\ell=\tau}^{t}d_\ell)-\sum_{t=\tau+r}^{\tau+\xi-1}b(x+\tilde{u}-\sum_{\ell=\tau}^{t}d_\ell)\right\}, \qquad (20)\end{aligned}$$

where the last equality follows from the structure of $\mathcal{D}^r$. Since $\mathcal{D}^r$ is polyhedron for all r, it follows that (20) is an LP, and the adversary's problem (16) reduces to solving $\xi+1$ LPs.

We now consider the case with lost-sales dynamics. In this case, $x+\tilde{u}$ is always non-negative. From Lemma 1, it follows that

$$\max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u},\mathbf{d}) = \begin{cases}\max_{\mathbf{d}\in\mathcal{D}^r}\left\{\sum_{t=\tau}^{\tau+r-1}h\Big(x+\tilde{u}-\sum_{\ell=\tau}^{t}d_\ell\Big)-b\Big(x+\tilde{u}-\sum_{t=\tau}^{\tau+\xi-1}d_t\Big)\right\}, & 0\le r\le \xi-1,\\ \max_{\mathbf{d}\in\mathcal{D}^r}\left\{\sum_{t=\tau}^{\tau+\xi-1}h\Big(x+\tilde{u}-\sum_{\ell=\tau}^{t}d_\ell\Big)\right\}, & r=\xi,\end{cases}$$

where $\mathcal{D}^r$ is given by (19). Since each $\mathcal{D}^r$ is a polyhedron, it follows that $\max_{\mathbf{d}\in\mathcal{D}^r} H(\tilde{u},\mathbf{d})$ is an LP for all r; thus, the adversary's problem (16) reduces to solving $\xi+1$ LPs. □

Unlike results in Lemma 3, the results in Lemma 4 exploit the special structure of the uncertainty set $\mathcal{D}$. That said, the piecewise linear structure of the inventory holding and shortage cost function H implies that the adversary's problem is efficiently solvable provided linear separation over the set $\mathcal{D}$ is efficient [22].

Lemma 3 and 4 imply that the Benders' decomposition approach should be efficient for solving (14), i.e., the first program in (11). Since the second program in (11) is simply the adversary's problem with $\tilde{u}=0$, we now have an efficient algorithm for solving (11).

4 Inventory model with non-zero lead time

In this section, we consider an inventory model which is identical to our basic model in all other respects except that the lead time is a positive constant L, i.e., the quantity ordered at the beginning of period k arrives at the installation at the beginning of period $k+L$. Thus, the inventory dynamics are as follows:

(a) When the excess demand is backlogged, $x_{k+1} = x_k + u_{k-L} - d_k$ for $k \in \{1, 2, \dots, T-1\}$.

(b) When the excess demand is lost, $x_{k+1} = \max\{x_k + u_{k-L} - d_k, 0\}$ for $k \in \{1, 2, \dots, T-1\}$.

Note that u_t, for $t \leq 0$, denotes orders placed before period 1, and are known constants at the beginning of the planning horizon. The cost C_k incurred during period k is given by

$$C_k = K\mathbf{1}(u_k > 0) + c \cdot u_k + \max\Big\{h \cdot (x_k + u_{k-L} - d_k), -b \cdot (x_k + u_{k-L} - d_k)\Big\}. \qquad (21)$$

In (21) we do not include the pipeline inventory costs; however, since the lead time is a constant L, this cost can be readily incorporated into the model by redefining the variable ordering cost as $c' = c + Lh'$, where h' denotes the per unit per period pipeline inventory cost. For simplicity, we assume that the installation capacity $M = +\infty$.

Positive lead times lead to a larger state space which makes computing the optimal policy for the inventory model much harder, especially when the excess demand is lost. Our understanding of the lost-sales inventory model with positive lead time is still very limited [44], and this model is still a very active area of research. See Karlin and Scarf [27], Morton [30, 31], Janakiraman and Roundy [25], Zipkin [44], Huh et al. [23] and references therein for the progression of our understanding of this model.

In order to generalize the basic policy in Section 3 to the inventory model with positive lead time L, we associate the *inventory cycle* $(\tau + L, \tau + L + 1, \dots, \tau + L + \xi - 1)$ with the ordering cycle $(\tau, \tau + 1, \dots, \tau + \xi - 1)$, i.e., the inventory cycle is obtained by shifting the corresponding ordering cycle forward by L periods. Our modified policy is defined as follows. As before orders are only placed at the beginning of each *ordering cycle*. However, unlike in the basic policy, (u_τ, ξ) are chosen to minimize the cost associated with placing the order u_τ and the inventory holding cost and shortage cost incurred over the corresponding *inventory cycle* averaged over cycle length ξ. The cost function is formally defined in (22). The motivation here is that the order decision u_τ, and the implicit decision $u_t = 0$ for $\tau + 1 \leq t \leq \tau + \xi - 1$, are matched with the inventory holding and shortage cost incurred over the corresponding inventory cycle. Note that when $L = 0$, the modified policy reduces to the basic policy.

Robust optimization approach for computing (u_τ, ξ) Let x_τ denote the available inventory and $\bar{\mathbf{u}} = (\bar{u}_{\tau-L}, \bar{u}_{\tau-L+1}, \ldots, \bar{u}_{\tau-1})$ the vector of outstanding orders at the beginning of the first period τ of a new ordering cycle. The inventory cycle corresponding to this ordering cycle is $(\tau+L, \tau+L+1, \ldots, \tau+L+\xi-1)$. Since $u_t = 0$ for all $t \in \{\tau+1, \ldots, \tau+\xi-1\}$, the order u_τ arrives at the beginning of the inventory cycle and the realized demand over the inventory cycle is $\mathbf{d}[\tau+L, \xi]$, the inventory holding cost and shortage cost incurred over the inventory cycle is given by $H\big(\tau+L, x_{\tau+L}(\mathbf{d}[\tau, L]), \xi, u_\tau, \mathbf{d}[\tau+L, \xi]\big)$, where the functional form of H is given by (3) (resp. (5)) when the unmet demand is backlogged (resp. lost), and we emphasize the fact that the inventory level $x_{\tau+L}$ is a function of $\mathbf{d}[\tau, L]$. The cost $G(u_\tau, \xi)$ associated with the decision (u_τ, ξ) is given by

$$G(u_\tau, \xi) \quad = \quad \max_{\mathbf{d}[\tau] \in \mathcal{D}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau > 0) + cu_\tau + H\Big(\tau+L, x_{\tau+L}\big(\mathbf{d}[\tau, L]\big), \xi, u_\tau, \mathbf{d}[\tau+L, \xi]\Big)}{\xi} \right\}. \tag{22}$$

The order quantity u_τ and cycle length ξ are computed by solving

$$\min_{\xi \in \big\{1,2,\ldots,\min\{T-\tau-L+1, U\}\big\}} \min_{u_\tau \geq 0} \Big\{G(u_\tau, \xi)\Big\}. \tag{23}$$

Solution algorithm We use the solution algorithm proposed in Section 3 to solve program (23). Specifically, we first fix ξ, solve the inner minimization problem, and then enumerate all feasible value of ξ. In (22) the adversary chooses a demand sequence $\mathbf{d}[\tau]$ that belongs to the uncertainty set $\mathcal{D}[\tau]$ which is defined in (7). A simple modification of the argument in Lemma 2 establishes that the adversary can be restricted to the set $\mathcal{D}[\tau, L+\xi]$ which is defined analogous to the set in (12). To simplify notation, let $u = u_\tau$, $\mathbf{d}^O = \mathbf{d}[\tau, L]$, $x(\mathbf{d}^O) = x_{\tau+L}(\mathbf{d}[\tau, L])$, $\mathbf{d}^I = \mathbf{d}[\tau+L, \xi]$, $\mathbf{d} = (\mathbf{d}^O, \mathbf{d}^I) = \mathbf{d}[\tau, L+\xi]$, and $\mathcal{D} = \mathcal{D}[\tau, L+\xi]$.

Since $(\tau+L, \xi)$ are constants in the inner minimization problem, we write $H\big(\tau+L, x(\mathbf{d}^O), \xi, u, \mathbf{d}^I\big)$

as $H\big(x(\mathbf{d}^O), u, \mathbf{d}^I\big)$. The inner minimization problem in u is equivalent to

$$\min_{u\geq 0}\max_{\mathbf{d}\in\mathcal{D}}\Big\{K\mathbf{1}(u>0)+cu+H\big(x(\mathbf{d}^O),u,\mathbf{d}^I\big)\Big\}. \tag{24}$$

Using results in Lemma 2, it is easy to show that (24) can be reformulated as

$$\min\left\{\begin{array}{c}\min_{u\geq 0}\max_{\mathbf{d}\in\mathcal{D}}\Big\{K+cu+H\big(x(\mathbf{d}^O),u,\mathbf{d}^I\big)\Big\},\\ \max_{\mathbf{d}\in\mathcal{D}}\Big\{H\big(x(\mathbf{d}^O),0,\mathbf{d}^I\big)\Big\}\end{array}\right\}. \tag{25}$$

We use Benders' decomposition to solve the first program in (25); the solution to the second program will be obtained as a special case.

Lemma 5. The inventory manager's problem in the Benders' decomposition approach to solve the first problem in (25) is given by $\min_{u\geq 0}\max_{\mathbf{d}\in\tilde{D}} K+cu+H\Big(x\big(\mathbf{d}^O\big),\ u,\ \mathbf{d}^I\Big)$,where $\tilde{D}$ is a given finite working list of demand sequences. This problem can be reformulated as an LP with 2 decision variables and $(\xi+1)|\tilde{D}|+1$ constraints, where $|\tilde{D}|$ denotes the cardinality of the set $\tilde{D}$.

The proof of this result is identical to that of Lemma 3. Next, we address the adversary's problem.

Lemma 6. The adversary's problem in the Benders' decomposition applied to the first problem in (25) is given by

$$\max_{\mathbf{d}=(\mathbf{d}^O,\mathbf{d}^I)\in\mathcal{D}} K+c\tilde{u}+H\Big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\Big), \tag{26}$$

where $\tilde{u}$ denotes the current working decision of the inventory manager. For backlogging dynamics, (26) is equivalent to solving at most $\xi+1$ LPs, and for lost-sales dynamics, (26) is equivalent to solving at most $(\xi+1)(L+1)$ LPs.

Proof. Since $K+c\tilde{u}$ is a constant, we can simply omit it and rewrite the adversary's problem as

$$\max_{\mathbf{d}=(\mathbf{d}^O,\mathbf{d}^I)\in\mathcal{D}} H\Big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\Big). \tag{27}$$

In the backlogging case $x(\mathbf{d}^O) = x_\tau + \sum_{t=\tau}^{\tau+L-1}(\bar{u}_{t-L} - d_t)$. Therefore, (27) is equivalent to

$$\max_{\mathbf{d}\in\mathcal{D}} \sum_{t=\tau+L}^{\tau+L+\xi-1} \max\left\{h\Big(x_\tau + \sum_{\ell=\tau-L}^{\tau-1}\bar{u}_\ell + \tilde{u} - \sum_{\ell=\tau}^{t} d_\ell\Big), -b\Big(x_\tau + \sum_{\ell=\tau-L}^{\tau-1}\bar{u}_\ell + \tilde{u} - \sum_{\ell=\tau}^{t} d_\ell\Big)\right\}. \quad (28)$$

Let $q = x_\tau + \sum_{\ell=\tau-L}^{\tau-1}\bar{u}_\ell + \tilde{u}$. We consider the following two cases.

(i) $q \le 0$: Since $q - \sum_{\ell=\tau}^{t} d_\ell \le 0$, for all $\tau + L \le t \le \tau + L + \xi - 1$, it follows that (28) reduces to the single LP, $\max_{\mathbf{d}\in\mathcal{D}}\Big\{ - \sum_{t=\tau+L}^{\tau+L+\xi-1} b\big(q - \sum_{\ell=\tau}^{t} d_\ell\big)\Big\}$.

(ii) $q > 0$: As in Lemma 4, (27) is equivalent to $\max_{r\in\{0,\ldots,\xi\}}\max_{\mathbf{d}\in\mathcal{D}^r}\Big\{H\Big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\Big)\Big\}$, where

$$\mathcal{D}^r = \begin{cases} \Big\{\mathbf{d}\in\mathcal{D} \mid q - \sum_{\ell=\tau}^{\tau+L} d_\ell \le 0\Big\}, & r = 0,\\ \Big\{\mathbf{d}\in\mathcal{D} \mid q - \sum_{\ell=\tau}^{\tau+L+r-1} d_\ell \ge 0,\ q - \sum_{\ell=\tau}^{\tau+L+r} d_\ell \le 0\Big\}, & 1 \le r \le \xi - 1,\\ \Big\{\mathbf{d}\in\mathcal{D} \mid q - \sum_{\ell=\tau}^{\tau+L+\xi-1} d_\ell \ge 0\Big\}, & r = \xi. \end{cases} \quad (29)$$

Since the definition of $\mathcal{D}^r$ implies that

$$\max_{\mathbf{d}\in\mathcal{D}^r}\Big\{H\Big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\Big)\Big\} = \max_{\mathbf{d}\in\mathcal{D}^r}\left\{\sum_{t=\tau+L}^{\tau+L+r-1} h\big(q - \sum_{\ell=\tau}^{t} d_\ell\big) - \sum_{t=\tau+L+r}^{\tau+L+\xi-1} b\big(q - \sum_{\ell=\tau}^{t} d_\ell\big)\right\},$$

and each $\mathcal{D}^r$ is polyhedron, $\max_{\mathbf{d}\in\mathcal{D}^r}\Big\{H\Big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\Big)\Big\}$ is an LP, and the adversary's problem (28) reduces to solving $\xi + 1$ LPs.

Next, we consider the case with lost-sales dynamics. Define

$$\theta^j(\mathbf{d}^O) = \begin{cases} x_\tau + \sum_{t=\tau}^{\tau+L-1}(\bar{u}_{t-L} - d_t), & j = 0,\\ \sum_{t=j+\tau}^{\tau+L-1}(\bar{u}_{t-L} - d_t), & 1 \le j \le L-1,\\ 0, & j = L. \end{cases} \quad (30)$$

Then the inventory level $x(\mathbf{d}^O) = \max\Big\{\theta^0(\mathbf{d}^O), \theta^1(\mathbf{d}^O), \ldots, \theta^L(\mathbf{d}^O)\Big\}$. Using this respresentation for $x(\mathbf{d}^O)$ we partition $\mathcal{D} = \cup_{j=0}^{L} \mathcal{E}^j$, where

$$\mathcal{E}^j \triangleq \Big\{\mathbf{d} \in \mathcal{D} \mid x(\mathbf{d}^O) = \theta^j(\mathbf{d}^O)\Big\} = \Big\{\mathbf{d} \in \mathcal{D} \mid \theta^j(\mathbf{d}^O) \geq \theta^\ell(\mathbf{d}^O), \forall\, \ell \neq j\Big\}. \tag{31}$$

We further partition $\mathcal{E}^j = \bigcup_{r=0}^{\xi} \mathcal{E}^{j,r}$, where

$$\mathcal{E}^{j,r} = \begin{cases} \Big\{\mathbf{d} \in \mathcal{E}^j \mid\ \theta^j(\mathbf{d}^O) + \tilde{u} - d_{\tau+L} \leq 0\Big\}, & r = 0, \\ \Big\{\mathbf{d} \in \mathcal{E}^j \mid\ \theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{\tau+L+r-1} d_\ell \geq 0, \theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{\tau+L+r} d_\ell \leq 0\Big\}, & 1 \leq r \leq \xi - 1, \\ \Big\{\mathbf{d} \in \mathcal{E}^j \mid\ \theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{\tau+L+\xi-1} d_\ell \geq 0\Big\}, & r = \xi. \end{cases}$$

Then $\max_{\mathbf{d}=(\mathbf{d}^O,\mathbf{d}^I)\in\mathcal{D}} H\big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\big) = \max_{j\in\{0,\dots,L\}} \max_{r\in\{0,\dots,\xi\}} \max_{\mathbf{d}\in\mathcal{E}^{j,r}} H\big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\big)$, and from the definition of $\mathcal{E}^{j,r}$ it follows that

$$\max_{\mathbf{d}\in\mathcal{E}^{j,r}} H\big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\big) = \begin{cases} \max_{\mathbf{d}\in\mathcal{E}^{j,r}} \Big\{ \sum_{t=\tau+L}^{\tau+L+r-1} h\Big(\theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{t} d_\ell\Big) \\ \qquad - b\Big(\theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{t=\tau+L}^{\tau+L+\xi-1} d_t\Big)\Big\}, & 0 \leq r \leq \xi - 1, \\ \max_{\mathbf{d}\in\mathcal{E}^{j,r}} \Big\{\sum_{t=\tau+L}^{\tau+L+\xi-1} h\Big(\theta^j(\mathbf{d}^O) + \tilde{u} - \sum_{\ell=\tau+L}^{t} d_\ell\Big)\Big\}, & r = \xi. \end{cases}$$

Since each $\theta^j(\mathbf{d}^O)$ is linear function of $\mathbf{d}^O$, and therefore of $\mathbf{d}$, $\mathcal{E}^{j,r}$ is polyhedron and $H\big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\big)$, when restricted to $\mathcal{E}^{j,r}$, is a linear function of $\mathbf{d}$, it follows that $\max_{\mathbf{d}\in\mathcal{E}^{j,r}} H\big(x\big(\mathbf{d}^O\big),\ \tilde{u},\ \mathbf{d}^I\big)$ is an LP. Thus, the adversary's problem (26) reduces to solving $(\xi + 1)(L + 1)$ LPs. □

5 Extension to joint pricing and inventory control

In this section, we consider a single installation joint pricing and inventory control model which is similar to that proposed in Federgruen and Heching [19] and Chen and Simchi-Levi [16]. See, also, Chen and Simchi-Levi [17] and Song et al. [40], for similar models. We show that a simple extension of our cycle-based control policy can efficiently compute a good solution for this model.

Dynamics We consider a single product single installation joint pricing and inventory control model over a finite discrete horizon of T periods. As in Section 2, we let x_k, u_k and d_k denote the inventory at the beginning of period k, the order quantity in period k, and the demand in period k, respectively. We denote by p_k the unit price of the product in period k. At the beginning of period k, the inventory manager chooses the order quantity $u_k \in [0,\ M - x_k]$ and the price $p_k \in \big[p^k_{\min},\ p^k_{\max}\big]$, where M denotes the capacity of the installation, and $\big[p^k_{\min},\ p^k_{\max}\big]$ denotes the feasible pricing interval. Following [19, 16], we assume that the excess demand in each period is backlogged, and the lead time is zero. Therefore, the system dynamics is given by $x_{k+1} = x_k + u_k - d_k$, for all $1 \le k \le T-1$.

Cost function The ordering cost and the inventory holding and shortage cost are assumed to be the same as in Section 2. In addition, the inventory manager receives revenue $p_k d_k$ in period k. Thus, the total cost incurred during period k is given by

$$C_k = K\mathbf{1}(u_k > 0) + c \cdot u_k + \max\Big\{h \cdot (x_k + u_k - d_k), -b \cdot (x_k + u_k - d_k)\Big\} - p_k d_k. \qquad (32)$$

Note that we assume that we receive full revenue even for the units that are backlogged. This assumption is common in the joint pricing and inventory control literature [19, 16].

Demand distribution In the joint pricing and inventory control model, the demand process $\mathfrak{D} = (D_1, D_2, \ldots, D_T)$ is no longer *exogenous*; the inventory manager can control the stochastic demand D_k by choosing the pricing decision p_k. We assume that the demand D_k is given by a multiplicative stochastic demand function $Q^{(m)}_k\big(p_k, z_k\big)$ defined as follows.

Assumption 3. The multiplicative stochastic demand function $Q^{(m)}_k\big(p_k, z_k\big) = \gamma_k(p_k) \cdot z_k$, where the functions $\{\gamma_k(\cdot)\}$ and the random variables $\{z_k\}$ satisfy the following two conditions.

(a) The function $\gamma_k(\cdot)$ is continuous and strictly decreasing on $\big[p^{(k)}_{\min},\ p^{(k)}_{\max}\big]$. Let $\rho_k(\cdot)$ denote the inverse of $\gamma_k(\cdot)$. We assume $\rho_k(\omega) \cdot \omega$ is concave in ω on $\big[\gamma_k(p^k_{\max}),\ \gamma_k(p^k_{\min})\big]$.

(b) The random variables $\{z_k\}_{k=1}^T$ are nonnegative and mutually independent, and only the mean and variance of z_k are known.

Assumption 3(a) is a common assumption on stochastic demand function (see [16]). Assumption 3(b) is similar to Assumption 1 and 2. We focus on the multiplicative demand function because it includes linear demand curve $\gamma(p) = \alpha - \beta p$ ($\alpha > 0, \beta > 0$), log-linear demand curve $\gamma(p) = e^{\alpha - \beta p}$ ($\alpha > 0, \beta > 0$) and iso-elastic demand curve $\gamma(p) = \alpha p^{-\beta}$ ($\alpha > 0, \beta > 1$) as special cases. Another motivation for studying multiplicative demand function is that the existing literature on mathematical programming based approaches to the joint pricing and inventory control problem is limited to the case with additive stochastic demand function with linear demand curve, see e.g., [1]. We show that our cycle-based policy is able to work with multiplicative demand functions in a mathematical programming framework. We will briefly discuss a broad family of "linear" stochastic demand functions at the end of this section.

Extension of the cycle-based policy Our cycle-based control policy extends to this setting as follows. Assumption 3(a) implies that, instead of selecting the price p_k, the inventory manager can equivalently select the *nominal demand* $\omega_k = \gamma_k(p_k) \in \left[\gamma_k(p^k_{\max}),\ \gamma_k(p^k_{\min})\right]$. As in Section 3, we break the planning horizon into ordering cycles. At the first period τ of a new ordering cycle, after observing the inventory level x_τ, the inventory manager chooses order quantity u_τ, cycle length ξ and nominal demand sequence $\omega[\tau, \xi]$ to minimize the worst-case average cost over the cycle $(\tau, \tau+1, \ldots, \tau+\xi-1)$, $F(u_\tau, \xi, \omega[\tau, \xi])$, which is defined as

$$F(u_\tau, \xi, \omega[\tau, \xi]) = \max_{z[\tau] \in \mathcal{Z}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau > 0) + cu_\tau + W(u_\tau, \omega[\tau, \xi], z[\tau, \xi])}{\xi} \right\}, \tag{33}$$

where $z[\tau]$ denotes $(z_\tau, z_{\tau+1}, \ldots, z_T)$, $z[\tau, \xi]$ denotes $(z_\tau, z_{\tau+1}, \ldots, z_{\tau+\xi-1})$, $\mathcal{Z}[\tau]$ is defined similarly as $\mathcal{D}[\tau]$ in (7), i.e.,

$$\mathcal{Z}[\tau] = \left\{ z[\tau] \middle| \bar{\delta}_t - \hat{\delta}_t \le z_t \le \bar{\delta}_t + \hat{\delta}_t, \quad \sum_{\ell=\tau}^{t} \frac{|z_\ell - \bar{\delta}_\ell|}{\hat{\delta}_\ell} \le \Gamma^\tau_{t-\tau+1}, \quad \forall\, \tau \le t \le T \right\}, \tag{34}$$

where $\bar{\delta}_t$ and $\hat{\delta}_t$ are, respectively, the *nominal value* and *deviation* of z_t, and

$$\begin{aligned} W(u_\tau, \omega[\tau,\xi], z[\tau,\xi]) \quad &= \quad \sum_{t=\tau}^{\tau+\xi-1} \max\Big\{h(x_\tau + u_\tau - \sum_{\ell=\tau}^{t} \omega_\ell z_\ell), -b(x_\tau + u_\tau - \sum_{\ell=\tau}^{t} \omega_\ell z_\ell)\Big\} \\ &\quad - \sum_{t=\tau}^{\tau+\xi-1} \rho_t(\omega_t)\omega_t z_t. \end{aligned}$$

The optimal decision $(u^*_\tau, \xi^*, \omega^*[\tau,\xi])$ is chosen by solving the optimization problem

$$\min_{\xi\in\Big\{1,2,\ldots,\min\{T-\tau+1,U\}\Big\}} \quad \min_{\Big\{0\le u_\tau\le M-x_\tau, \omega[\tau,\xi]\in\Omega[\tau,\xi]\Big\}} F(u_\tau, \xi, \omega[\tau,\xi]), \tag{35}$$

where

$$\Omega[\tau,\xi] = \Big\{\omega[\tau,\xi] \;\Big|\; \omega_t \in [\gamma_t(p^t_{\max}),\ \gamma_t(p^t_{\min})],\ \text{for } \tau \le t \le \tau+\xi-1\Big\}. \tag{36}$$

Solution algorithm To solve (35), we fix ξ, solve the inner minimization problem, and then enumerate all feasible value of ξ. A simple modification of Lemma 2 establishes that the adversary can be restricted to the set $\mathcal{Z}[\tau,\xi]$ which is defined analogously to the set in (12). To simplify notation, let $u = u_\tau$, $x = x_\tau$, $\omega = \omega[\tau,\xi]$, $z = z[\tau,\xi]$, $\Omega = \Omega[\tau,\xi]$ and $\mathcal{Z} = \mathcal{Z}[\tau,\xi]$. Following the arguments in Lemma 2, it is clear that solving the inner minimization problem of (35) is equivalent to solving the following optimization problem:

$$\min \left\{ \begin{array}{c} \min_{\Big\{0\le u\le M-x, \omega\in\Omega\Big\}} \max_{z\in\mathcal{Z}} \Big\{K + cu + W(u,\omega,z)\Big\}, \\ \min_{\omega\in\Omega} \max_{z\in\mathcal{Z}} \Big\{W(0,\omega,z)\Big\} \end{array} \right\}. \tag{37}$$

We use Benders' decomposition to solve each of the optimization problems in (37). We focus on the first program without loss of generality, i.e., we show how to apply Benders' decomposition to efficiently solve

$$\min_{\Big\{0\le u\le M-x, \omega\in\Omega\Big\}} \max_{z\in\mathcal{Z}} \Big\{K + cu + W(u,\omega,z)\Big\}. \tag{38}$$

Assumption 3(a) and the fact that $z \ge 0$ ensure that $W(u,\omega,z)$ is convex in (u,ω); therefore, (38) is a convex optimization problem. Next, we show that the inventory manager's problem and the adversary's problem of (38) can be solved efficiently.

Lemma 7. The inventory manager's problem of (38) is $\min_{\left\{0\le u\le M-x,\omega\in\Omega\right\}} \max_{z\in\tilde{Z}} \Big\{K + cu + W(u,\omega,z)\Big\}$, where $\tilde{Z}$ is a given finite working list of random terms. The inventory manager's problem is a convex optimization problem.

By noting that $W(u,\omega,z)$ is convex in (u,ω), the proof of Lemma 7 follows immediately. Next, we consider the adversary's problem.

Lemma 8. The adversary's problem of (38) is $\max_{z\in\mathcal{Z}} \Big\{K + c\tilde{u} + W(\tilde{u},\tilde{\omega},z)\Big\}$, where $(\tilde{u},\tilde{\omega})$ denotes the current working decision of the inventory manager. The adversary's problem can be reduced to solving at most $\xi + 1$ LPs.

The proof of this lemma is similar to that of Lemma 4. Thus we have completed our algorithm.

Linear stochastic demand function Our cycle-based policy extends to the "linear" stochastic demand functions $Q_k^{(l)}(p_k, z_k)$ defined as follows.

Assumption 4. The linear stochastic demand functions $\{Q_k^{(l)}(p_k, z_k)\}$ satisfy the following three conditions.

(a) $Q_k^{(l)}(p_k, z_k)$ is separately linear in p_k and z_k, i.e., linear in p_k for fixed z_k, and vice versa.

(b) $Q_k^{(l)}(p_k, z_k)$ is non-increasing in p_k for fixed z_k.

(c) $\{z_k\}_{k=1}^T$ are mutually independent, and only the mean and variance of z_k are known.

Assumption 4(c) is essentially identical to Assumption 3(b) except that we do not require z_k to be nonnegative. This class of demand functions includes the widely-used additive demand function with linear demand curve $Q(p,z) = \alpha - \beta p + z$ $(\alpha > 0, \beta \ge 0)$ as a special case. Although the demand functions in [19] are allowed to be more general, the convexity related regularity conditions essentially restrict the demand functions to the class of "linear" demand functions defined here.

Our policy can be extended to this case by making the following simple modification. Since $Q_k^{(l)}$ is linear in p_k, we directly work with the prices $\{p_k\}$ as decision variables. In particular, at the first period τ of a new ordering cycle, the inventory manager chooses order quantity u_τ, cycle length ξ and price sequence $p[\tau,\xi]$ to minimize the worst-case average cost over the cycle $(\tau,\tau+1,\ldots,\tau+\xi-1)$, $F(u_\tau,\xi,p[\tau,\xi])$,

$$F(u_\tau,\xi,p[\tau,\xi]) = \max_{z[\tau]\in\mathcal{Z}[\tau]} \left\{ \frac{K\mathbf{1}(u_\tau>0)+cu_\tau+W(u_\tau,p[\tau,\xi],z[\tau,\xi])}{\xi} \right\}, \tag{39}$$

where $\mathcal{Z}[\tau]$ is defined in (34), and

$$\begin{aligned} W(u_\tau,p[\tau,\xi],z[\tau,\xi]) \quad = \quad & \sum_{t=\tau}^{\tau+\xi-1} \max\left\{h(x_\tau+u_\tau-\sum_{\ell=\tau}^{t} Q_\ell^{(l)}(p_\ell,z_\ell)), -b(x_\tau+u_\tau-\sum_{\ell=\tau}^{t} Q_\ell^{(l)}(p_\ell,z_\ell))\right\} \\ & -\sum_{t=\tau}^{\tau+\xi-1} p_t Q_t^{(l)}(p_t,z_t). \end{aligned}$$

Note that for fixed z-sequence $z[\tau,\xi]$, $W(u_\tau,p[\tau,\xi],z[\tau,\xi])$ is a convex piecewise quadratic function in $(u_\tau,p[\tau,\xi])$. The optimal decision $(u_\tau^*,\xi^*,p^*[\tau,\xi])$ is chosen by solving the optimization problem

$$\min_{\xi\in\left\{1,2,\ldots,\min\{T-\tau+1,U\}\right\}} \; \min_{\left\{0\le u_\tau\le M-x_\tau,\, p[\tau,\xi]\in\mathcal{P}[\tau,\xi]\right\}} F(u_\tau,\xi,p[\tau,\xi]), \tag{40}$$

where

$$\mathcal{P}[\tau,\xi] = \left\{p[\tau,\xi] \mid p_{\min}^t \le p_t \le p_{\max}^t, \text{ for } \tau\le t\le \tau+\xi-1\right\}. \tag{41}$$

We use solution algorithm which is identical to that of (35) to solve (40). In particular, we apply Benders' decomposition to solve the counterpart of program (38).

Lemma 9. The inventory manager's problem $\min_{\left\{0\le u\le M-x,p\in\mathcal{P}\right\}} \max_{z\in\tilde{Z}} \left\{K+cu+W(u,p,z)\right\}$ is a convex quadratically constrained quadratic program (QCQP). When the linear demand function $Q_k^{(l)}(p,z)$ is additive, i.e., $Q_k^{(l)}(p,z) = \alpha_k - \beta_k p + z$ ($\alpha_k>0, \beta_k\ge 0$), the inventory manager's problem reduces to a convex quadratic program (QP).

Proof. The inventory manager's problem can be reformulated as

$$\begin{aligned}
\min \quad & K + cu + t, \\
\text{s.t.} \quad & W(u,p,z^j) \le t, \quad \forall z^j \in \tilde{Z}, \\
& 0 \le u \le M - x, \; p \in \mathcal{P}.
\end{aligned}$$

Recall that $W(u,p,z)$ is a convex piecewise quadratic function of (u,p) for a fixed z. Therefore, the constraint $W(u,p,z^j) \le t$ can be reformulated as a collection of convex quadratic constraints.

When the linear demand function is additive, the quadratic term in p in

$$\sum_{t=\tau}^{\tau+\xi-1} p_t Q_t^{(l)}(p_t, z_t) = \sum_{t=\tau}^{\tau+\xi-1} p_t\big(\alpha_t - \beta_t p_t + z_t\big) = \sum_{t=\tau}^{\tau+\xi-1} p_t z_t + \sum_{t=\tau}^{\tau+\xi-1} p_t\big(\alpha_t - \beta_t p_t\big)$$

is independent of z. Consequently, the inventory manager's problem can be reformulated as a convex quadratic program, i.e., with a convex quadratic objective but only linear constraints. □

At the same time, we can still compute the optimal solution to the adversary's problem by solving at most $\xi + 1$ LPs.

6 Numerical experiments

In this section, we report the results of a set of numerical experiments that investigate the empirical performance of our cycle-based inventory control policy (CI).

6.1 Backlogging dynamics

We compared CI with the rolling horizon version of the open-loop policy (RHBT) proposed in [14]. RHBT efficiently computes the ordering decisions for a backlogging inventory control problem by solving LPs or MIPs. In [14] it was shown that the performance of RHBT is comparable or superior to that of the DP-based policies, and in [15] it was shown that the performance of RHBT is superior to that of the robust base-stock policy.

We set the parameters $\bar{\omega}_t$, $\hat{\omega}_t$ and Γ_j^t defining the uncertainty sets (7) to $\bar{\omega}_t = \bar{d}_t$, $\hat{\omega}_t =$

Parameter	Value	Parameter	Value
Time horizon T	48	Fixed ordering cost K	500
Initial inventory x_1	0	Variable ordering cost/unit c	1
Demand distribution D_t	$N(\bar{d}_t, \sigma_t^2)$	Inventory holding cost/unit h	4
Mean of demand $\bar{d}_t$	$100 + 40\sin(\frac{2\pi}{12}t)$	Shortage cost/unit b	6
Std. dev. of demand σ_t	$0.25\bar{d}_t$	Lead time L	0

Table 1: Parameters for base backlogging model

$\min\{2\sigma_t, \bar{\omega}_t\}$, $\Gamma_j^t = \sqrt{j}$, for $1 \leq t \leq T$ and $1 \leq j \leq T - t + 1$ (see Section 5 in [14]). To ensure a fair comparison, we use the same uncertainty sets for both CI and RHBT. We set $U = 12$ for CI. All the LPs and MIPs in the implementation of CI and RHBT are solved using Gurobi Optimizer v4.5. We call Gurobi Optimizer from MATLAB using Gurobi Mex [42].

Our goal in the numerical experiments was to understand the impact of problem parameters on the performance of CI and RHBT. In particular, we investigated the impact of different choices of the realized demand distribution, the mean demand trajectory, the standard deviation of demand, the cost parameters (K, c, h, b) and the lead time L. The parameters for the base model are displayed in Table 1, where the choice of c, h and b is suggested in [14]. To study the impact of the parameters, we change one parameter at a time while keeping all other parameters at their respective base values. For each configuration, we randomly generated $N = 100$ demand sequences and tested the performance of CI and RHBT on the same demand realizations. In these numerical experiments, we set $c = 0$ when *computing* the CI policy decisions. This is motivated by the fact that variable ordering costs incurred over a fixed finite horizon can be viewed as a sunk cost when c is smaller than backlogging cost b. We found this modification lead to better performance for the CI policy. Note we still used the true value of c when *evaluating* the performance of the CI policy. We report the mean and standard deviation of the total cost incurred by CI and RHBT on the $N = 100$ sample paths. At the end of this section, we also report the running time of CI and RHBT.

Impact of the realized demand distribution. We considered the following 5 different families of realized demand distributions: normal, Student's-$t(4)$, gamma, uniform and lognormal. In each

Distribution	RHBT Mean: std dev	CI Mean(%): std dev
Normal	$29,330.6 : 646.1$	$28,894.3(-1.49\%) : 754.3$
T(4)	$28,990.6 : 654.4$	$28,562.1(-1.48\%) : 693.1$
Gamma	$29,422.5 : 561.1$	$28,782.3(-2.17\%) : 727.8$
Uniform	$29,649.1 : 573.8$	$28,996.7(-2.20\%) : 674.9$
Lognormal	$29,446.3 : 704.1$	$28,823.5(-2.11\%) : 723.1$

Table 2: Impact of realized demand distribution

β	RHBT Mean: std dev	CI Mean(%): std dev
0	$29,919.9 : 516.7$	$28,684.7(-4.13\%) : 648.7$
20	$29,546.2 : 574.4$	$28,867.0(-2.30\%) : 654.2$
40	$29,330.6 : 646.1$	$28,894.3(-1.49\%) : 754.3$
60	$28,596.9 : 632.2$	$28,984.6(+1.36\%) : 788.2$
80	$27,844.3 : 654.6$	$29,609.1(+6.34\%) : 1,087.9$

Table 3: Impact of β

configuration, $\{D_t\}_{t=1}^T$ were drawn from the same family, and with pre-specified mean $\{\bar{d}_t\}_{t=1}^T$ and variance $\{\sigma_t^2\}_{t=1}^T$ in the base model. Table 2 summaries the performance measures. Since CI and RHBT both only use the first two moments of demand distributions, their performance is fairly robust across various realized demand distributions. The performance of the two policies is comparable – the cost of the CI policy is 1.5% to 2.2% less on average, but volatility of the cost is higher.

Impact of the mean demand trajectory. We considered mean demand trajectory $\bar{d}_t = 100 + \beta \sin(\frac{2\pi}{12}t)$, with $\beta = 0, 20, 40, 60, 80$. In the base model, $\beta = 40$. Note that a larger β implies greater variability in $\{\bar{d}_t\}_{t=1}^T$. Table 3 summarizes the performance measures. When $\beta = 0$, i.e., $\{D_t\}_{t=1}^T$ is an IID demand, the CI cost is approximately 4% lower than RHBT. As β increases, the cost of RHBT decreases while the cost of CI increases. When $\beta = 80$, i.e., the demand trajectory is highly variable, CI costs 6% more on average and incurs a higher standard deviation. Thus, CI outperforms RHBT when the demand is changing smoothly, and the reverse is true when the demand is changing sharply.

γ	RHBT Mean: std dev	CI Mean(%): std dev
0.15	$28,303.1: 351.8$	$28,236.7(-0.23\%): 439.8$
0.20	$28,830.5: 489.7$	$28,714.6(-0.40\%): 596.9$
0.25	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
0.30	$29,709.3: 728.6$	$29,203.1(-1.70\%): 790.1$
0.35	$30,195.1: 937.6$	$29,629.2(-1.87\%): 888.5$

Table 4: Impact of γ

K	RHBT Mean: std dev	CI Mean(%): std dev
0	$9,492.5: 536.6$	$9,492.5(-0.00\%): 536.6$
250	$21,277.1: 479.6$	$21,378.1(+0.47\%): 575.7$
500	$29,330.6: 646.1$	$28,894.3(-1.49\%): 754.3$
750	$34,101.2: 813.3$	$34,301.4(+0.59\%): 876.1$
1000	$38,070.8: 847.4$	$38,619.8(+1.44\%): 871.1$

Table 5: Impact of K

Impact of the standard deviation of demand. We fixed $\gamma = \frac{\sigma_t}{d_t}$ for $1 \leq t \leq T$, therefore the ratio γ determines the relative magnitude of standard deviation. In the base model, $\gamma = 0.25$. Table 4 summarizes the performance measures. As γ increases, the performance of CI improves relative to RHBT. When $\gamma = 0.35$, i.e., the standard deviation of demand is high, *both* the average cost and volatility of the cost of the CI policy are lower than those of the RHBT policy.

Impact of the fixed ordering cost. Table 5 summarizes the performance measures. CI has a slightly higher average cost in 3 cases ($K = 250, 750, 1000$). When $K = 0$, we found that CI and RHBT incurred the same cost on every sample path.

Impact of the variable ordering cost. Table 6 summarizes the performance measures. CI has a lower average cost across all the 5 configurations, but the magnitudes of savings are decreasing as c increases. When c increases to 3, there are jumps in the volatility of the cost of both CI and RHBT.

c	RHBT Mean: std dev	CI Mean(%): std dev
0	$24,545.7$: 519.1	$23,994.9(-2.24\%)$: 615.4
1	$29,330.6$: 646.1	$28,894.3(-1.49\%)$: 754.3
2	$33,923.3$: 658.5	$33,537.2(-1.14\%)$: 678.5
3	$38,702.7$: 947.1	$38,323.1(-0.98\%)$: $1,085.1$
4	$43,556.1$: 849.5	$43,339.9(-0.50\%)$: $1,075.6$

Table 6: Impact of c

h	RHBT Mean: std dev	CI Mean(%): std dev
2	$24,254.5$: 489.5	$23,883.2(-1.53\%)$: 555.3
4	$29,330.6$: 646.1	$28,894.3(-1.49\%)$: 754.3
6	$31,637.7$: 755.9	$31,530.0(-0.34\%)$: 900.9
8	$33,106.6$: 875.7	$33,100.8(-0.02\%)$: 838.7
10	$34,162.2$: 952.8	$34,299.6(+0.40\%)$: $1,024.7$

Table 7: Impact of h

Impact of the inventory holding cost. Table 7 summarizes the performance measures. CI has a lower average cost in 4 out of the 5 cases, but the relative performance of CI over RHBT deteriorates as h increases.

Impact of the shortage cost. Table 8 summarizes the performance measures. CI has a lower average cost across all the 5 cases, but also incurs significantly higher volatility of cost when $b = 6, 8$.

Impact of the lead time. When lead time is L, we assumed that the outstanding orders at the beginning of the planning horizon, $(u_{1-L}, \ldots, u_0) = (\bar{d_1}, \bar{d_2}, \ldots, \bar{d_L})$. Table 9 summarizes the performance measures. It is clear from the table, that the average CI cost is significantly lower when compared to the RHBT cost across all 5 configurations. The volatility of the cost of the two policies is comparable; it is also clear that the volatility of the cost of both policies increases relatively quickly as lead time L increases.

Next, we compared the running time of CI and RHBT. In particular, we investigated how the

b	RHBT Mean: std dev	CI Mean(%): std dev
2	$22,430.9$: 664.4	$21,917.1(-2.29\%)$: 539.8
4	$26,639.4$: 625.1	$26,273.5(-1.37\%)$: 623.4
6	$29,330.6$: 646.1	$28,894.3(-1.49\%)$: 754.3
8	$30,870.1$: 612.9	$30,813.5(-0.18\%)$: 799.6
10	$31,881.5$: 685.4	$31,492.1(-1.22\%)$: 664.6

Table 8: Impact of b

L	RHBT Mean: std dev	CI Mean(%): std dev
0	$29,330.6$: 646.1	$28,894.3(-1.49\%)$: 754.3
1	$29,951.6$: 638.5	$28,899.3(-3.51\%)$: 773.4
2	$30,519.4$: $1,154.7$	$29,214.4(-4.28\%)$: $1,231.2$
4	$31,255.0$: $1,788.9$	$29,957.6(-4.15\%)$: $1,747.9$
6	$31,812.8$: $2,227.6$	$31,057.7(-2.37\%)$: $2,142.3$

Table 9: Impact of L

time horizon, T, and the fixed ordering cost, K, impact the running time of the two policies. For each configuration below, we randomly generated $N = 10$ demand sequences, and implemented CI and RHBT on the same demand realizations. We report the mean and standard deviation of the running time of CI and RHBT on the $N = 10$ sample paths.

Impact of the time horizon on running time. We considered $T = 24, 36, 48, 60, 72$, with all other parameters were chosen from the base model in Table 1. Table 10 reports the running time. We found that the running time of CI grows linearly with T, whereas the running time of RHBT

T	RHBT Mean: std dev (in seconds)	CI Mean: std dev (in seconds)
24	0.72 : 0.02	4.74 : 0.21
36	5.16 : 0.24	7.77 : 0.38
48	28.25 : 0.88	10.38 : 0.25
60	72.99 : 1.04	13.16 : 0.34
72	141.77 : 1.87	16.24 : 0.26

Table 10: Running time: Impact of T

increases exponentially as T increases. This result is not surprising – CI only relies on solving LPs, whereas RHBT solves MIPs with increasing size as T increases.

Impact of the fixed ordering cost on running time. We fixed $T = 48$ and let $K = 0, 250, 500, 750, 1000$ in each of the 5 scenarios. Table 11 reports the running time. Table 11 shows

K	RHBT Mean: std dev (in seconds)	CI Mean: std dev (in seconds)
0	1.24 : 0.01	26.73 : 1.76
250	1.65 : 0.02	21.49 : 0.50
500	28.25 : 0.88	10.38 : 0.25
750	48.68 : 0.89	10.26 : 0.41
1000	46.05 : 1.08	9.12 : 0.55

Table 11: Running time: Impact of K

some interesting results. When $K = 0$ or 250, RHBT costs less than 2 seconds – approximately 5% of that of CI. However, as K increases, the running time of RHBT increases significantly, indicating that the MIPs in RHBT are becoming harder to solve as K increases. When $K = 750$ or 1000, it takes RHBT almost 50 seconds to solve the problem. On the other hand, the running time of CI in fact decreases as K increases, since a larger K implies a lower ordering frequency, and therefore fewer decisions to make. It is clear that, when $K = 1000$, the running time of RHBT is more than 5 times of that of CI.

6.2 Lost-sales dynamics

6.2.1 Inventory control when lead time is zero

In the lost-sales model with zero lead time, we compared CI with the DP solution and the optimal base-stock policy [43] in *hindsight* (BH). Following [14], we compute the DP policy *assuming* $D_t \sim N(\bar{d}_t, \sigma_t^2)$ for $1 \leq t \leq T$, and approximate $N(\bar{d}_t, \sigma_t^2)$ by a probability mass function supported on the five points $\{\bar{d}_t - 2\sigma_t, \bar{d}_t - \sigma_t, \bar{d}_t, \bar{d}_t + \sigma_t, \bar{d}_t + 2\sigma_t\}$. The BH policy is base-stock policy but the order-up-to level S is computed in hindsight, i.e., *after* observing the entire demand realization $\mathbf{d}$. We compute the optimal base-stock level for a particular $\mathbf{d}$ by an exhaustive search. It is clear

Parameter	Value	Parameter	Value
Time horizon T	48	Fixed ordering cost K	0
Initial inventory x_1	0	Variable ordering cost/unit c	1
Demand distribution D_t	$N(\bar{d}_t, \sigma_t^2)$	Inventory holding cost/unit h	4
Mean of demand $\bar{d}_t$	100	Shortage cost/unit b	12
Std. dev. of demand σ_t	$0.25\bar{d}_t$	Lead time L	0

Table 12: Parameters for base lost-sales model

that the cost of BH is a lower bound for the cost of any non-anticipatory base-stock policy. The choice of BH as a comparison policy is motivated by the fact that base-stock policies with constant order-up-to level, are commonly used in practice in lost-sales setting [25].

We set the parameters $\bar{\omega}_t$, $\hat{\omega}_t$, Γ_j^t and U in the CI policy to those used in Section 6.1. In the DP policy, the state at each time instance is the inventory at hand. We restricted the state to the interval $[0, 500]$, i.e., set $M = 500$, and further restricted both the state and action values to multiples of 0.1.

We investigated the impact of the same set of parameters as in Section 6.1 on the performance of CI, DP and BH, except that we did not consider positive lead time since DP does not scale with lead time L. The parameters for the base model are displayed in Table 12. To study the impact of the parameters, we vary one parameter at a time while keeping all other parameters at their respective base values. For each configuration, we randomly generated $N = 100$ demand sequences and tested the performance of CI, DP and BH on the same demand realizations. In this section, we still set $c = 0$ when *computing* the CI policy decisions and used the true value of c when *evaluating* the performance of the CI policy. We report the mean and standard deviation of the total cost incurred by CI, DP and BH on the $N = 100$ sample paths. We report the running time for CI and DP at the end of this section.

Impact of the realized demand distribution. As in Section 6.1, we considered the following 5 different families of realized demand distributions: normal, Student's-$t(4)$, gamma, uniform and lognormal. Table 13 summaries the performance measures. DP and CI have comparable performance across the 5 scenarios; they are both robust with respect to the changes in realized demand

Distribution	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
Normal	$11,090.3 : 545.5$	$11,105.8(+0.14\%) : 520.1$	$10,611.4(-4.32\%) : 619.2$
T(4)	$10,845.0 : 774.7$	$10,871.7(+0.25\%) : 750.2$	$9,954.9(-8.21\%) : 925.9$
Gamma	$11,414.9 : 717.9$	$11,408.8(-0.05\%) : 668.3$	$10,812.9(-5.27\%) : 834.0$
Uniform	$11,009.3 : 421.6$	$10,998.1(-0.10\%) : 407.8$	$10,780.0(-2.08\%) : 452.3$
Lognormal	$11,418.7 : 670.9$	$11,405.9(-0.11\%) : 673.6$	$10,691.6(-6.37\%) : 953.2$

Table 13: Impact of realized demand distribution

β	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0	$11,090.3 : 545.5$	$11,105.8(+0.14\%) : 520.1$	$10,611.4(-\ 4.32\%) : 619.2$
20	$11,229.2 : 532.5$	$11,215.6(-0.12\%) : 517.7$	$11,763.6(+\ 4.76\%) : 799.7$
40	$11,247.0 : 555.7$	$11,209.8(-0.33\%) : 529.9$	$14,641.3(+30.18\%) : 989.3$
60	$11,144.8 : 603.9$	$11,145.8(+0.01\%) : 587.7$	$17,674.9(+58.59\%) : 1,345.3$
80	$11,206.2 : 609.4$	$11,218.7(+0.11\%) : 643.5$	$21,180.5(+89.01\%) : 1,415.5$

Table 14: Impact of β

distributions. BH incurs significantly less cost, since the mean demand trajectory is flat and there is no fixed ordering cost in the base model.

Impact of the mean demand trajectory. We considered mean demand trajectory $\bar{d}_t = 100 + \beta \sin(\frac{2\pi}{12}t)$, with $\beta = 0, 20, 40, 60, 80$. In the base model, $\beta = 0$, i.e., demands are IID random variables. Table 14 summarizes the performance measures. The costs of DP and CI are almost equal; however, the cost of BH increases significantly as β increases. This result is not surprising – the performance of base-stock policies with constant base-stock level is very sensitive to the variability of the mean demand trajectory. In cases where the mean demands vary relatively sharply, base-stock policies deteriorate quickly.

Impact of the standard deviation of demand. We fixed $\gamma = \frac{\sigma_t}{\bar{d}_t}$ for $1 \leq t \leq T$. In the base model, $\gamma = 0.25$. Table 15 summarizes the performance measures. It is clear from the table, that the costs of DP and CI are still very close.

γ	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0.15	$8,568.6$: 307.9	$8,576.3(+0.09\%)$: 297.0	$8,293.9(-3.21\%)$: 338.2
0.20	$9,800.2$: 493.2	$9,790.2(-0.10\%)$: 470.9	$9,342.9(-4.67\%)$: 601.1
0.25	$11,090.3$: 545.5	$11,105.8(+0.14\%)$: 520.1	$10,611.4(-4.32\%)$: 619.2
0.30	$12,326.4$: 752.1	$12,337.4(+0.09\%)$: 732.2	$11,728.2(-4.85\%)$: 881.5
0.35	$13,699.2$: 854.7	$13,707.5(+0.06\%)$: 836.4	$13,064.4(-4.63\%)$: $1,036.1$

Table 15: Impact of γ

K	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0	$11,090.3$: 545.5	$11,105.8(+0.14\%)$: 520.1	$10,611.4(-4.32\%)$: 619.2
250	$22,935.7$: 574.0	$23,138.0(+0.88\%)$: 522.6	$22,613.9(-1.40\%)$: 668.4
500	$31,234.3$: 913.5	$32,170.9(+3.00\%)$: 792.4	$34,630.4(+10.87\%)$: 733.8
750	$37,061.8$: 951.0	$37,754.4(+1.87\%)$: 711.0	$46,562.9(+25.64\%)$: 737.5
1000	$42,124.8$: $1,043.6$	$43,339.0(+2.88\%)$: 953.7	$58,557.4(+39.01\%)$: 735.6

Table 16: Impact of K

Impact of the fixed ordering cost. Table 16 summarizes the performance measures. CI incurs 1% to 3% higher average cost compared to DP when $K \neq 0$, but saves approximately 10% to 20% in standard deviation. On the other hand, the performance of BH deteriorates significantly as K grows, since base-stock policies place positive orders in every period as long as demand is positive, regardless of the value of K. This suggests that base-stock policies are in general not good candidates when there are non-zero fixed ordering costs.

Impact of the variable ordering cost. Table 17 summarizes the performance measures. The performance of DP and CI is comparable when $c = 0, 1, 2, 3$ – DP incurs a slightly lower average cost while CI incurs a slightly lower standard deviation. When $c = 4$, CI costs 1.5% more on average but is more stable in that the standard deviation of CI is only half of that of DP.

Impact of the holding cost. Table 18 summarizes the performance measures. CI incurs 3% more in average when $h = 2$ and 2.5% less when $h = 6$ compared to DP; in other scenarios the differences are small. On the other hand, CI has a smaller standard deviation across all 5 cases. In

c	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
0	$6,411.9$: 681.1	$6,408.7(-0.05\%)$: 681.3	$5,950.7(-7.19\%)$: 732.4
1	$11,090.3$: 545.5	$11,105.8(+0.14\%)$: 520.1	$10,611.4(-4.32\%)$: 619.2
2	$15,785.7$: 549.7	$15,829.4(+0.28\%)$: 539.0	$15,238.5(-3.47\%)$: 715.9
3	$20,421.0$: 470.9	$20,468.4(+0.23\%)$: 445.1	$19,710.3(-3.48\%)$: 840.0
4	$24,886.2$: $1,171.6$	$25,268.1(+1.53\%)$: 553.0	$24,330.1(-2.23\%)$: 927.8

Table 17: Impact of c

h	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
2	$8,508.8$: 507.5	$8,778.9(+3.17\%)$: 325.6	$8,411.0(-1.15\%)$: 501.2
4	$11,090.3$: 545.5	$11,105.8(+0.14\%)$: 520.1	$10,611.4(-4.32\%)$: 619.2
6	$13,059.5$: $1,110.3$	$12,732.3(-2.51\%)$: 665.2	$12,329.4(-5.59\%)$: 763.5
8	$14,154.8$: $1,023.1$	$14,095.3(-0.42\%)$: 854.7	$13,780.5(-2.64\%)$: 877.8
10	$14,918.8$: $1,079.7$	$14,920.1(+0.01\%)$: $1,007.9$	$14,681.5(-1.59\%)$: $1,004.0$

Table 18: Impact of h

particular, when $h = 6$, the standard deviation of CI cost is approximately 40% less than that of DP cost.

Impact of the shortage cost. Table 19 summarizes the performance measures. In all cases

b	DP Mean: std dev	CI Mean(%): std dev	BH Mean(%): std dev
8	$9,943.5$: 682.4	$9,987.1(+0.44\%)$: 483.2	$9,610.7(-3.35\%)$: 575.9
10	$10,951.5$: 485.9	$10,652.9(-2.73\%)$: 519.1	$10,234.4(-6.55\%)$: 650.7
12	$11,090.3$: 545.5	$11,105.8(+0.14\%)$: 520.1	$10,611.4(-4.32\%)$: 619.2
14	$11,236.6$: 610.4	$11,428.5(+1.71\%)$: 581.3	$10,917.3(-2.84\%)$: 701.4
16	$11,520.0$: 684.2	$11,863.4(+2.98\%)$: 591.4	$11,233.5(-2.49\%)$: 818.2

Table 19: Impact of b

except when $b = 10$, CI incurs a higher average cost and a lower standard deviation compared to DP. When $b = 10$, CI costs 2.7% less but incurs a higher standard deviation of cost.

Impact of the fixed ordering cost on running time. Next, we report the running time for CI and DP. In this set of experiments we were interested in investigating the impact of the fixed ordering cost K on the running time of these two policies. We set $K = 0, 250, 500, 750, 1000$, and set all other parameters to the values in the base model displayed in Table 12. Table 20 reports the running time. As in the backlogging case, the running time for CI typically decreases as K increases, and CI goes from being approximately 10 times faster to approximately 20 times faster as we increase K from 0 to 1000.

K	DP (in seconds)	CI Mean: std dev (in seconds)
0	154.32	16.88 : 1.03
250	155.94	15.76 : 0.28
500	156.39	7.85 : 0.13
750	156.09	7.96 : 0.21
1000	157.13	8.34 : 0.20

Table 20: Running time: Impact of K

6.2.2 Inventory control with positive lead time

In this section, we consider the lost-sales model with positive lead time L. In this case, the DP optimal policy is very hard to compute because of the "curse of dimensionality". Consequently, comparing CI with DP is impractical. Zipkin [44] compared the performance of the optimal policy and eight plausible heuristics – two different myopic policies, the dual-balancing policy of Levi et al. [28], four different base-stock policies and the constant-order policy in Reiman [34] – for an infinite horizon lost-sales inventory model with stationary demand and the long-run average cost criterion. The parameters for this model were as follows: fixed cost $K = 0$, variable ordering cost $c = 0$, inventory holding cost $h = 1$, lost-sales penalty cost $b \in \{4, 9, 19, 39\}$, lead time $L \in \{1, 2, 3, 4\}$, and demand distribution D_t was either Poisson with mean 5 or geometric with mean 5. For details, see Section 6 of [44].

In this section, we adopt the simulation design in [44]. Since [44] considered the long-run average cost criterion, and our cycle-based policy CI is designed for finite horizon inventory problems, we used the batch means method to approximate the long-run average cost. The specific simulation strategy was as follows: for each combination of parameter values, we simulated *one* sample path under the CI policy for a total time horizon of $T = 50,000$ periods. We discarded the first $T_b = 10,000$ periods as the "burn-in" periods. We equally divided the rest $T_s = 40,000$ periods into $N = 10$ batches, each consists of $T_s/N = 4,000$ consecutive periods. We computed the average cost over each batch and then the mean and standard deviation across the $N = 10$ batches. We report the mean as an estimator for the long-run average cost of CI.

Since in all the cases we considered, the mean demand was 5, we set $\bar{\omega}_t = 5$. We set $\Gamma_j^t = \sqrt{j}$, for $1 \leq t \leq T$ and $1 \leq j \leq T - t + 1$, and set $U = 5$. The deviation $\hat{\omega}_t$ was selected as a function of the penalty cost b and the demand distribution. The specific choice for $\hat{\omega}_t$ are displayed in Table 21. Note that the choice of $\hat{\omega}_t$ is independent of the lead time – this was done for simplicity

	Demand Distribution	
b	Poisson	geometric
4	2	3
9	2	5
19	3	5
39	4	5

Table 21: Choice of $\hat{\omega}_t$

and one can further improve performance by choosing a lead time dependent $\hat{\omega}_t$. Our choice of $\hat{\omega}_t$ is non-decreasing in b, and for a given b, we chose a larger $\hat{\omega}_t$ for the geometric case. This choice is motivated by the fact that as the lost-sales penalty b increases the decision maker should prefer a "fatter" uncertainty set to protect against the very expensive stock-out events more aggressively. And since the geometric distribution has heavier tails, the uncertainty set is chosen to be larger in this case.

Poisson demand. We summarize the performance of CI when demand is Poisson in Table 22. We report the optimal average cost taken from [44], the long-run average cost for CI, and the amount by which the CI cost exceeds the optimal cost as a percentage of the optimal cost. We also report the standard deviation of the batch averages in parenthesis. For the sake of brevity, we only report the rank of CI among the nine heuristics, i.e., the original eight heuristics in [44] and CI, for each set of parameter values. We refer the interested readers to [44] for the cost for each of the other eight heuristics.

	Lead time							
	1		2		3		4	
(a) $b = 4$								
Optimal	4.04		4.40		4.60		4.73	
CI	4.07	+0.74%	4.49	+2.00%	4.86	+5.66%	5.13	+8.42%
	(0.05)	2nd	(0.07)	4th	(0.06)	5th	(0.12)	5th
(b) $b = 9$								
Optimal	5.44		6.09		6.53		6.84	
CI	5.83	+7.12%	6.23	+2.25%	6.70	+2.62%	7.00	+2.37%
	(0.10)	8th	(0.11)	4th	(0.14)	3rd	(0.10)	4th
(c) $b = 19$								
Optimal	6.68		7.66		8.36		8.89	
CI	7.19	+7.62%	7.75	+1.22%	8.61	+2.93%	9.08	+2.14%
	(0.17)	7th	(0.25)	2nd	(0.25)	5th	(0.18)	3rd
(d) $b = 39$								
Optimal	7.84		9.11		10.04		10.79	
CI	8.21	+4.66%	9.42	+3.39%	10.37	+3.24%	10.93	+1.27%
	(0.16)	7th	(0.29)	7th	(0.30)	5th	(0.25)	2nd

Table 22: Poisson case

Recall that the optimal policy and the eight heuristics in [44] utilize the full distribution information of demand and specifically optimize the long-run average cost. In contrast, CI only uses partial demand information and adopts a min-max approach to compute the order decisions. Consequently, the simulations favor of the polices in [44] over CI. The results in Table 22 indicate that the performance of CI is still quite good. In 12 out of the 16 parameter combinations, the CI cost exceeds the optimal cost by no more than 5%. Note that when $b = 4$, i.e., the lost-sales penalty is small, the relative performance of CI deteriorates as lead time L increases; whereas for $b \in \{9, 19, 39\}$, the opposite is true, i.e., the relative performance improves with lead time. This behavior can be explained by considering the dynamics of the CI policy for positive lead times. Recall that the CI policy selects u_τ to balance the ordering cost with the holding cost and lost-sales penalty over the *inventory cycle* $(\tau + L, \ldots, \tau + L + \xi - 1)$. In order to counteract the possibility of stock-out due to the uncertain demand over the intervening interval $(\tau, \ldots, \tau + L - 1)$, the CI policy tends to choose a larger value of u_τ when the lead time L is large. When the lost-sales penalty b is small, this higher inventory position becomes a "liability" because it forces the inventory manager to pay more inventory holding cost; however, when b is large, higher inventory position becomes an "asset" since it helps the inventory manager to avoid or alleviate the expensive stock out.

Geometric demand. Table 23 summarizes the performance of CI when demand is geometric. CI performs quite well when $b \in \{4, 9\}$, i.e., when the lost-sales penalty is modest. However, as b becomes large, the performance of CI deteriorates quickly. When $b = 39$, CI cost is significantly larger than the optimal cost and is beaten by most other heuristics.

	Lead time							
	1		2		3		4	
(a) $b = 4$								
Optimal	9.82		10.24		10.47		10.61	
CI	9.95	+1.28%	10.57	+3.26%	11.19	+6.84%	11.75	+10.74%
	(0.19)	3rd	(0.25)	4th	(0.23)	6th	(0.23)	6th
(b) $b = 9$								
Optimal	14.51		15.50		16.14		16.58	
CI	14.82	+2.13%	15.56	+0.37%	16.63	+3.06%	17.38	+4.85%
	(0.32)	5th	(0.38)	1st	(0.26)	3rd	(0.39)	3rd
(c) $b = 19$								
Optimal	19.22		20.89		22.06		22.95	
CI	22.70	+18.13%	22.36	+7.02%	23.25	+5.39%	23.95	+4.34%
	(1.16)	7th	(0.99)	6th	(0.68)	5th	(0.61)	4th
(d) $b = 39$								
Optimal	23.87		26.21		27.96		29.36	
CI	37.57	+57.37%	35.57	+35.69%	36.25	+29.64%	35.32	+20.31%
	(1.36)	9th	(1.00)	8th	(1.70)	7th	(1.49)	7th

Table 23: Geometric case

The poor performance of CI for large b is because we use a symmetric uncertainty set while the true demand distribution is highly asymmetric and has a considerable probability mass outside of the uncertainty set which, in turn results, in a high lost-sales penalty. When $b = 39$, we found that the relatively infrequent stock-out events, which occur on average once every 7–9 periods, account for 70–80% of the total cost. In order to correct this we considered the following asymmetric uncertainty set:

$$\mathcal{D}[\tau] = \left\{ \mathbf{d}[\tau] \;\middle|\; \bar{\omega}_t - \hat{\omega}_t^l z_t \le d_t \le \bar{\omega}_t + \hat{\omega}_t^r z_t, \quad 0 \le z_t \le 1, \quad \sum_{\ell=\tau}^{t} z_\ell \le \Gamma_{t-\tau+1}^{\tau}, \quad \forall\, \tau \le t \le T \right\}, \tag{42}$$

where $\hat{\omega}_t^l$ (resp. $\hat{\omega}_t^r$) denotes the *left* (resp. *right*) deviation. It is clear that we recover the symmetric uncertainty set (7) by setting $\hat{\omega}_t^l = \hat{\omega}_t^r = \hat{\omega}_t$. We re-did the cases with $b = 19$ and 39, where CI with symmetric uncertainty sets performs poorly. To ensure a fair comparison, we tested CI with asymmetric uncertainty sets on the same sample paths that we tested CI with symmetric uncertainty sets. In our numerical experiments, we set $\hat{\omega}_t^l = 5$ and $\hat{\omega}_t^r = 10$. We report the results in Table 24. The new results are quite encouraging. When $b = 39$, the cost of CI with asymmetric uncertainty sets exceeds the optimal cost by no more than 5%.

	Lead time							
	1		2		3		4	
(c) $b = 19$								
Optimal	19.22		20.89		22.06		22.95	
CI(Asy)	19.43	+1.08%	21.41	+2.47%	23.58	+6.87%	25.58	+11.44%
	(0.63)	5th	(0.55)	5th	(0.25)	5th	(0.48)	5th
(d) $b = 39$								
Optimal	23.87		26.21		27.96		29.36	
CI(Asy)	25.00	+4.75%	26.43	+0.82%	28.77	+2.90%	30.20	+2.87%
	(0.75)	7th	(0.67)	3rd	(1.03)	5th	(0.94)	4th

Table 24: Geometric case continued

These experiments indicate that tailoring the uncertainty sets by using additional information of the demand distribution, e.g., whether the demand is asymmetric about the mean, the tail decays slowly, etc., and problem parameters, e.g., lost-sales penalty cost b, results in a close-to-optimal performance for a very large set of scenarios.

6.3 Joint pricing and inventory control

In this section, we report the empirical performance of the cycle-based joint pricing and inventory control policy in Section 5. We continue to refer to our proposed cycle-based policy as CI.

We considered two scenarios which are adapted from the two scenarios considered in [19]. In both scenarios, we assumed that the stochastic demand functions are of the form $Q_k(p_k, z_k) = \alpha_k - \beta_k p_k + z_k$, i.e., additive demand functions with linear demand curve. Thus, Lemma 9 implies that the inventory manager's problem can be reformulated as a convex quadratic program. We use Gurobi Optimizer v4.5 to solve the QPs in the implementation of CI. We set the parameters $\bar{\delta}_t$, $\hat{\delta}_t$ and Γ_j^t defining the uncertainty sets (34) to $\bar{\delta}_t = \bar{z}_t$, $\hat{\delta}_t = 2\sigma_t$, $\Gamma_j^t = \frac{1}{4}\sqrt{j}$, for $1 \leq t \leq T$ and $1 \leq j \leq T - t + 1$. We also set $U = 15$.

We compared the performance of CI with the stochastic DP solution computed by assuming that the random term z_t was distributed according to a 5-point approximation to the $N(\bar{z}_t, \sigma_t^2)$ distribution supported on the set $\{\bar{z}_t - 2\sigma_t, \bar{z}_t - \sigma_t, \bar{z}_t, \bar{z}_t + \sigma_t, \bar{z}_t + 2\sigma_t\}$. The state at each time instance is, as before, the inventory at hand. We restricted the state $x \in [-200, 500]$ and discretized the state to take values in the set of integers. The action space in the DP was two dimensional – u and p – and we restricted each of the two dimensions to take values in the set of integers.

Table 25 displays the parameters of the first scenario – the "Dress Scenario" in [19]. We considered the following 5 different families of realized random term distributions: normal, Student's-$t(4)$, gamma, uniform and lognormal. In each configuration, $\{z_t\}_{t=1}^T$ were drawn from the same family, and with pre-specified mean $\{\bar{z}_t\}_{t=1}^T$ and variance $\{\sigma_t^2\}_{t=1}^T$. For each configuration, we randomly generated $N = 100$ random term sequences and tested the performance of CI and DP on the same random term realizations. Since in the two scenarios of this section the variable ordering cost c is close to or larger than the backlogging cost b and the CI policy is myopic, the optimal decision of the policy tends to order less over a given cycle and "transfers" the variable ordering cost to the subsequent cycle. In order to correctly account for the externality of the ordering decision over a given cycle, in the numerical experiments we added another penalty term $c \cdot \max\{-(x_\tau + u_\tau - \sum_{\ell=\tau}^{\tau+\xi-1} Q_\ell^{(l)}(p_\ell, z_\ell)), 0\}$ to $W(u_\tau, p[\tau, \xi], z[\tau, \xi])$, i.e., we penalized the end of the cycle backlogged demand by $b + c$ per unit. It is easy to see that after the modification our

Parameter	Value	Parameter	Value
Time horizon T	48	Std. dev. of random term σ_t	15
Fixed ordering cost K	500	Intercept of demand curve α_t	174
Variable ordering cost/unit c	22.15	Slope of demand curve β_t	3
Inventory holding cost/unit h	0.22	$p^t_{\max}$	44
Shortage cost/unit b	21.78	$p^t_{\min}$	25
Mean of random term $\bar{z}_t$	0	Capacity M	500

Table 25: Parameters of Scenario 1 – "Dress Scenario"

Distribution	DP Mean: std dev	CI Mean(%): std dev
Normal	$40,668.4 : \ 2,100.3$	$39,652.3(-2.50\%) : \ 2,236.9$
T(4)	$41,123.2 : \ 1,646.4$	$40,224.8(-2.18\%) : \ 1,746.6$
Gamma	$41,003.5 : \ 1,615.3$	$40,028.7(-2.38\%) : \ 1,681.2$
Uniform	$40,782.6 : \ 1,878.0$	$39,792.9(-2.43\%) : \ 1,930.2$
Lognormal	$40,937.1 : \ 1,804.9$	$40,013.7(-2.26\%) : \ 1,838.7$

Table 26: Performance summary of Scenario 1

algorithm can be implemented in an identical fashion. We report the mean and standard deviation of the total *profit* generated by CI and DP on the $N = 100$ sample paths. Table 26 summarizes the performance measures. Table 27 displays the parameters of the second scenario – the "Skirt Scenario" in [19]. We repeated the procedures as in the first scenario. Table 28 summarizes the performance measures. In all cases considered, the average profit of CI policy is approximately 2-2.5% below that of the DP-based policy, and the difference in the average profit of CI policy and the DP-based policy is around half of the confidence interval for the average profit of the DP-based policy.

The main difference between the two policies is the running time. We display the running time of both policies in Table 29. From the table, it is clear that for the two scenarios we considered, CI is approximately 10 times faster than DP. Moreover, DP costs around 50% more time in Scenario 2 than in Scenario 1 because the price range $[15, 44]$ contains 50% more grid points than the price range $[25, 44]$, whereas the running times of CI in both scenarios are almost identical. Thus, the running time of the DP-based policy is very sensitive to the size of the parameter ranges, but that

Parameter	Value	Parameter	Value
Time horizon T	48	Std. dev. of random term σ_t	6.5
Fixed ordering cost K	500	Intercept of demand curve α_t	57
Variable ordering cost/unit c	14.05	Slope of demand curve β_t	1
Inventory holding cost/unit h	0.17	$p^t_{\max}$	44
Shortage cost/unit b	16.83	$p^t_{\min}$	15
Mean of random term $\bar{z}_t$	0	Capacity M	500

Table 27: Parameters of Scenario 2 – "Skirt Scenario"

Distribution	DP Mean: std dev	CI Mean(%): std dev
Normal	$19,307.4$: 874.1	$18,852.5(-2.36\%)$: 834.2
T(4)	$19,382.3$: $1,038.2$	$18,930.1(-2.33\%)$: $1,041.8$
Gamma	$19,200.2$: 941.2	$18,765.2(-2.27\%)$: 936.8
Uniform	$19,086.8$: 997.7	$18,652.2(-2.28\%)$: 966.1
Lognormal	$19,179.8$: 847.1	$18,746.4(-2.26\%)$: 850.3

Table 28: Performance summary of Scenario 2

the running time of CI is completely insensitive to the size of the parameter sets.

Scenario	DP (in seconds)	CI Mean: std dev (in seconds)
1	72.21	8.15 : 1.20
2	106.45	8.15 : 1.29

Table 29: Running time comparison

7 Conclusions

In this paper we have proposed a new robust cycle-based control policy for a single installation, finite horizon inventory model with non-stationary uncertain demand. Our policy is an extension of the EOQ and (R,T)-type policies, and, like these policies, it is structurally simple and the decisions in the policy have an intuitive appeal. Our policy provides a unified framework for efficiently solving inventory control problems under various model assumptions. In particular, our policy can be efficiently implemented both when the excess demand is backlogged or lost, when the fixed ordering cost is non-zero, and when the lead time is non-zero. The policy can also be extended to the joint pricing and inventory control problem with a simple generalization.

The optimal decisions in our policy are computed by a Benders decomposition approach where all the sub-problems are linear programs of modest size, even when the problem contains positive fixed ordering costs. Since our policy only requires LP solvers, it is easily implementable and has a higher likelihood of being adopted in practice. Our numerical experiments showed that the performance of our policy is very close to that of the competing polices while requiring significantly lower computing resources.

In this paper, we have restricted our attention to the single product, single installation inventory models. Just as in the (R, T) policy, the cycle length, or equivalently, the next reorder epoch, is an explicit decision variable in our policy. Consequently, coordinating orders in more complex inventory networks is relatively easy in the context of our policy [33]. We are currently investigating extensions of our robust cycle-based policy to more general supply chain settings, e.g., serial system [18] and one-warehouse-multi-retailer system (OWMR) [38].

References

[1] E. Adida and G. Perakis. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming*, 107(1):97–129, 2006.

[2] D.R. Atkins and P.O. Iyogun. Periodic versus 'can-order' policies for coordinated multi-item inventory systems. *Management Science*, pages 791–796, 1988.

[3] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization.* Princeton Series in Applied Mathematics. Princeton University Press, 2009.

[4] A. Ben-Tal, B. Golany, A. Nemirovski, and J.P. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *M&SOM*, 7:248–271, 2005.

[5] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Math. OR.*, 23:769–805, 1998.

[6] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Oper. Research Lett.*, 25:1–13, 1999.

[7] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming., Ser. A*, 88:411–424, 2000.

[8] J. Benders. Partitioning procedures for solving mixed variables programming problems. *Numerische Mathematik*, 4:238–252, 1962.

[9] D. Bertsekas. *Dynamic Programming And Optimal Control*, volume 1. Athena Scientific, 1995.

[10] D. Bertsimas, D. Iancu, and P. Parrilo. A hierarchy of near-optimal policies for multi-stage adaptive optimization. *Submitted for publication*, 2009.

[11] D. Bertsimas, D. Iancu, and P. Parrilo. Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.*, 35(2):363–394, 2010.

[12] D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Math. Programming, Ser. B*, 98:49–71, 2003.

[13] D. Bertsimas and M. Sim. The price of robustness. *Oper. Res.*, 52:35–53, 2004.

[14] D. Bertsimas and A. Thiele. A robust optimization approach to inventory theory. *Oper. Res.*, 54:150–168, 2006.

[15] D. Bienstock and N. Ozbay. Computing robust basestock levels. *Discrete Optimization*, 5:389–414, 2008.

[16] X. Chen and D. Simchi-Levi. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Operations Research*, pages 887–896, 2004.

[17] X. Chen and D. Simchi-Levi. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. *Mathematics of Operations Research*, pages 698–723, 2004.

[18] A. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management Sci.*, 6(4):475–490, 1960.

[19] A. Federgruen and A. Heching. Combined pricing and inventory control under uncertainty. *Oper. Res.*, 47(3):454–475, 1999.

[20] G. Gallego and I. Moon. The distribution free newsboy problem: review and extensions. *Journal of the Operational Research Society*, pages 825–834, 1993.

[21] G. Gallego, J. Ryan, and D. Simchi-Levi. Minimax analysis for finite horizon inventory models. *IIE Trans.*, 33:861–874, 2001.

[22] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer Verlag, 1988.

[23] W.T. Huh, G. Janakiraman, J.A. Muckstadt, and P. Rusmevichientong. An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems with censored demand. *Mathematics of Operations Research*, 34(2):397–416, 2009.

[24] G. N. Iyengar. Robust dynamic programming. *Math. Oper. Res.*, 30(2):257–280, 2005.

[25] G. Janakiraman and R. Roundy. Lost-sales problems with stochastic lead times: Convexity results for base-stock policies. *Oper. Res.*, 52(5):795–803, 2004.

[26] S. Karlin. Dynamic inventory policy with varying stochastic demands. *Management Science*, pages 231–258, 1960.

[27] S. Karlin and H. Scarf. Inventory models of the Arrow-Harris-Marschak type with time lag. *Studies in the Mathematical Theory of Inventory and Production*, pages 155–178, 1958.

[28] R. Levi, G. Janakiraman, and M. Nagarajan. Provably near-optimal balancing policies for stochastic inventory control models with lost sales. *Mathematics of Operations Research*, 33(2):351–374, 2008.

[29] I. Moon and G. Gallego. Distribution free procedures for some inventory models. *Journal of the Operational research Society*, pages 651–658, 1994.

[30] T.E. Morton. Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM review*, 11(4):572–596, 1969.

[31] T.E. Morton. The near-myopic nature of the lagged-proportional-cost inventory problem with lost sales. *Operations Research*, pages 1708–1716, 1971.

[32] A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Oper. Res.*, 53(5):780–798, 2005.

[33] U.S. Rao. Properties of the periodic review (R, T) inventory control policy for stationary, stochastic demand. *Manufacturing & Service Operations Management*, 5(1):37–53, 2003.

[34] M.I. Reiman. A new and simple policy for the continuous review lost sales inventory model. *Unpublished manuscript*, 2004.

[35] H. Scarf. A min-max solution of an inventory problem. *Studies in the Mathematical Theory of Inventory and Production*, pages 201–209, 1958.

[36] H. Scarf. The optimality of (s, S) policies in the dynamic inventory problem. In K. Arrow, S. Karlin, and P. Suppes, editors, *Mathematical Methods in the Social Sciences*, pages 196–202. Stanford University Press, Stanford, 1960.

[37] S. Sethi and F. Cheng. Optimality of (s, S) policies in inventory models with markovian demand. *Operations Research*, 45(6):931–939, 1997.

[38] K.H. Shang and S.X. Zhou. Optimizing replenishment intervals for two-echelon distribution systems with stochastic demand. *Working paper*, 2011.

[39] J.S. Song and P. Zipkin. Inventory control in a fluctuating demand environment. *Operations Research*, pages 351–370, 1993.

[40] Y. Song, S. Ray, and T. Boyaci. Optimal dynamic joint inventory-pricing control for multiplicative demand with fixed order costs and lost sales. *Operations Research*, 57(1):245–250, 2009.

[41] A.F. Veinott Jr. On the optimality of (s, S) inventory policies: New conditions and a new proof. *SIAM Journal on Applied Mathematics*, pages 1067–1083, 1966.

[42] W. Yin. *Gurobi Mex: A MATLAB interface for Gurobi.* URL: http://convexoptimization.com/wikimization/index.php/gurobi_mex, 2009-2011.

[43] P. Zipkin. *Foundations of Inventory Management.* McGraw-Hill Higher Education, 2000.

[44] P. Zipkin. Old and new methods for lost-sales inventory systems. *Operations Research*, 56(5):1256–1263, 2008.