Inventory Management Based on Target-Oriented Robust Optimization

Yun Fong Lim • Chen Wang

Lee Kong Chian School of Business, Singapore Management University,

50 Stamford Road, Singapore 178899, Singapore

yflim@smu.edu.sg • cwang@smu.edu.sg

June 2, 2016

Abstract

We propose a target-oriented robust optimization approach to solve a multi-product, multi-period inventory management problem subject to ordering capacity constraints. We assume the demand for each product in each period is characterized by an uncertainty set, which depends only on a reference value and the bounds of the demand. Our goal is to find an ordering policy that maximizes the sizes of all the uncertainty sets such that all demand realizations from the sets will result in a total cost lower than a pre-specified cost target. We prove that a static decision rule is optimal for an approximate formulation of the problem, which significantly reduces the computation burden. By tuning the cost target, the resultant policy can achieve a balance between the expected cost and the associated cost variance. Numerical experiments suggest that, although only limited demand information is used, the proposed approach performs comparably to traditional methods based on dynamic programming and stochastic programming. More importantly, our approach significantly outperforms the traditional methods if the latter assume inaccurate demand distributions. We demonstrate the applicability of our approach through two case studies from different industries.

Key words: Inventory, Cost, Variability, Lead Time, Robust Optimization, Target

1 Introduction

Inventory represents a significant part of our economy. For example, the total investment in inventory in the United States for a quarter can be as high as $1.57 trillion (Nahmias, 2009). Inventory improves service level by buffering uncertainty, but carrying inventory incurs holding cost. As global competition becomes more intense, good inventory management becomes crucial to the success of many companies. In this paper we address a multi-product, multi-period inventory management problem with fixed ordering costs and lead times. This problem is notoriously hard to solve and has been studied since 1950s (Scarf, 1958).

Starting from 1960s, dynamic programming (DP) appears as a leading methodology in this area of research. Some crucial theoretical results have been developed, such as the optimality of base stock policies (Scarf, 1959, 1960; Veinott, 1966; Azoury, 1985; Miller, 1986; Zipkin, 2000). The idea of the DP approach is to decompose a large problem into many small problems that can be solved relatively easily. This leads to explosion of the number of small recursive optimization problems, which is well known as the "curse of dimensionality". This limits the applicability of DP especially for problems with a large number of variables.

Another common approach to solve inventory management problems is to use stochastic programming (SP) (Birge and Louveaux, 2011). However, SP is limited to problems with a very short planning horizon because the event tree grows exponentially with the length of the planning horizon.

The above methods also face other difficulties in practice. The DP approach requires the knowledge of the probability distribution functions of the underlying stochastic parameters (for example, product demands). On the other hand, the SP approach needs sufficient samples of these parameters to estimate the expected value of the cost function. Identifying the distribution functions of the stochastic parameters is a common challenge in practice. For example, for a new electronic product such as a mobile phone, it is difficult to identify the distribution of future demand precisely. Furthermore, the product life may not be long enough to collect adequate samples or to observe any demand pattern.

Since the DP and SP approaches depend heavily on the distributions of stochastic parameters, a solution optimized under a particular distribution may perform badly under another distribution, even with the same mean and variance (Bertsimas and Thiele, 2004, 2006). As a result, it is necessary to adopt a methodology that relies only on partial information of the uncertain parameters. A promising approach is robust optimization (RO) that was initially introduced to immunize uncertain mathematical optimization problems against infeasibility, while preserving the tractability of the models (Ben-Tal and Nemirovski, 1998, 1999, 2000; Bertsimas and Sim, 2004, 2006; Bertsimas et al., 2004; El-Ghaoui and Lebret, 1997; El-Ghaoui et al., 1998; Goh and Sim, 2010).

The research in RO experiences exponential growth in the last decade. Most RO methods share the following two merits: (1) Only limited knowledge of underlying uncertain parameters is used. Most of the early papers in the literature assume only the support sets of uncertain parameters. Some applications assume that means and support sets are known to obtain additional theoretical results (Wang et al., 2009). Chen and Sim (2009) derive a rather tight upper bound on the expected value of the positive part of a random variable. This result has been used to tackle problems in supply chain management and finance (See and Sim, 2010; Chen et al., 2009). (2) The tractability of an optimization problem can be well preserved by the RO approach: The robust counterpart of a linear programming (LP) problem remains as an LP problem if the uncertain parameters are characterized by linearly constrained support sets, or remains as a second-order cone optimization problem if the original optimization problem and the uncertain support sets can be described using second-order cones.

Thanks to the above mentioned properties, RO has already been applied to solve problems in supply chain and inventory management. For example, Bertsimas and Thiele (2004, 2006) introduce an RO approach for supply chain management based on a "budget of uncertainty" model proposed by Bertsimas and Sim (2004). Bertsimas and Thiele (2004, 2006) show that an optimal solution has a base-stock structure, and this result can be extended to a supply chain with capacity constraints. Their numerical results suggest that the robust solution performs comparably to a solution based on an exact distribution, but the former works well against distributional ambiguity. Papers along a similar direction include Adida and Perakis (2006), Bienstock and Özbay (2008), and Song et al. (2012).

Under the RO approach, decision variables can be parameterized as functions of uncertain factors (also called *decision rules*) to improve the performance of the solution. For example, Ben-Tal et al. (2004) introduce the concept of adjustable robust counterparts, also known as linear decision rules, to postpone decisions until uncertainties are realized. Ben-Tal et al. (2005) introduce this technique to handle a retailer-supplier flexible commitment contract problem. Their simulation results suggest that although only support sets of uncertainty are given, the RO solution based on a linear decision rule approximates the optimal policy well. Ang et al. (2012) apply linear decision rules to solve the storage-retrieval problem in a unit-load warehouse with time-varying arrivals and random departures of different products over multiple periods. See and Sim (2010) generalize the linear decision rule to a truncated linear decision rule for inventory management problems with ordering capacity constraints. The authors also demonstrate the advantages of the truncated linear decision rule by simulations.

A special case of decision rules is a *static rule*, which is a static function of the uncertain factors. Some papers discuss the optimality of a static rule under special conditions. Ben-Tal and Nemirovski (1999) consider a single-stage problem under uncertainty. They show that a static rule is optimal if each constraint is associated with an uncertainty set, and the uncertainty sets of different constraints are independent of each other. Bertsimas et al. (2013) show that a static rule is optimal if and only if a certain transformation of the uncertainty sets is convex. In case the transformation is non-convex, the authors give a tight bound on the performance of a static rule. In contrast, we show in this paper that a static rule is optimal if a worst-case scenario of uncertainty can be identified.

Besides requiring less information on uncertainty, RO also differs from DP and SP in the objective function. DP and SP usually minimize the expected value of a certain performance measure (such as cost), while RO minimizes the worst-case value of the measure. However, both objective functions may not be suitable for decision makers because sometimes their objective is to fulfill a pre-specified target for the performance measure (Chen and Sim, 2009). For example, if the decision makers' objective is to ensure that the cost is always lower than a pre-specified budget, then a policy that minimizes the expected cost may not work well because it may lead to large variance of cost. This reduces the probability of fulfilling a budget that is larger than the expected cost.

This paper fills the gap by proposing a *target-oriented robust optimization* formulation that aims to maximize the chance of fulfilling a pre-specified target. By tuning the target value, the resultant policy can achieve a balance between the aggressiveness on performance measure and the associated risk. Specifically, our contributions can be summarized as follows:

1. We propose a new method for inventory management based on target-oriented robust optimization. This approach has the following four characteristics: (i) It only makes use of some reference values (such as the means) and the bounds of products' demands, and is suitable for problems with demand distributional ambiguity. (ii) The proposed formulation can be solved in a reasonable amount of time for multi-product, multi-period problems with realistic sizes. We demonstrate this in two case studies from different industries. (iii) Ordering capacity and many other kinds of constraints can be incorporated easily into our formulation. (iv) By tuning the cost target, a balance between the expected cost and the cost variance is achievable, which cannot be done by other existing approaches.

2. We prove the optimality of a static rule for a set of problems where a worst-case scenario of uncertainty can be identified. Specifically, if the feasibility of a solution can be ensured by its feasibility under a worst-case uncertainty scenario, then the static rule is optimal and no complicated decision rules are needed. Fortunately, for the inventory management problem a worst-case uncertainty scenario

can be found by applying a few relaxations. This makes our approach computationally amiable.

3. We analyze the solution of our approach in detail for a single-product, single-period problem. We discover a piece-wise linear structure of the solution. If the inventory level is lower than a threshold value, the order quantity depends on the inventory level and the value of objective function. Otherwise, the ordering policy becomes an order-up-to policy.

This paper is organized as follows. Section 2 describes the inventory management problem and defines notation. Section 3 describes the traditional DP approach. Section 4 introduces our formulation and method. Section 5 compares our method with the DP and SP approaches through numerical experiments. Section 6 demonstrates through two case studies that our approach is implementable in practice. Section 7 discusses some generalizations and Section 8 concludes the paper.

2 Problem description and notation

Consider an inventory system with P products indexed as $i = 1, \ldots, P$ over a planning horizon with T periods indexed as $t = 1, \ldots, T$. Let $\mathcal{P} := \{1, 2, 3, \ldots, P\}$ and $\mathcal{T} := \{1, 2, 3, \ldots, T\}$. In each period t, the following procedure is repeated:

1. Let y_t^i denote the on-hand inventory level of product i at the start of period t. Based on this inventory level, we place an order (or production lot) of quantity x_t^i at the start of period t for product i. This incurs a fixed ordering (or setup) cost A_t^i and a variable ordering cost $c_t^i x_t^i$, where c_t^i represents the ordering cost per unit. In this problem, x_t^i are decision variables and y_t^i are dependent variables.

2. We consider a constant lead time of $l(i)$ periods for product i. Thus, an order placed at the start of period $t - l(i)$ will arrive at the start of period t.

3. Each product i in period t faces random demand $\tilde{d}_t^i$, which is realized as d_t^i at the end of the period. The inventory level at the end of period t becomes $y_t^i + x_{t-l(i)}^i - d_t^i$.

4. If $y_t^i + x_{t-l(i)}^i - d_t^i \geq 0$, the remaining inventory is carried to the next period $t+1$. This incurs a holding cost h_t^i per unit. On the other hand, if $y_t^i + x_{t-l(i)}^i - d_t^i < 0$, unsatisfied demand is backlogged to the next period. This incurs a backlog cost b_t^i per unit.

5. At the start of period $t+1$ the inventory level is $y_{t+1}^i = y_t^i + x_{t-l(i)}^i - d_t^i$. Steps 1–4 are repeated.

Table 1 defines the notation used in this paper. It is noteworthy that in practice the replenishment decision x_t^i is not necessarily determined at the beginning of the planning horizon (the start of period 1). Instead, to achieve good performance it can be postponed to period t after observing the realization $\mathbf{d}_{t-1}$ of $\tilde{\mathbf{d}}_{t-1}$. Therefore, x_t^i is a *non-anticipative* function, denoted by $x_t^i(\mathbf{d}_{t-1})$, which depends only on demand information up to period $t-1$.

Table 1: **Notation**

P: number of products
$\mathcal{P}$: $\{1, 2, 3, \ldots, \text{P}\}$
T: number of periods in the planning horizon
$\mathcal{T}$: $\{1, 2, 3, \ldots, \text{T}\}$
$l(i)$: lead time (in number of periods) of product i
h_t^i: inventory holding cost per unit of product i from period t to period $t+1$
b_t^i: backlog cost per unit of product i from period t to period $t+1$
A_t^i: fixed ordering cost per order for product i in period t
c_t^i: variable ordering cost per unit of product i in period t
$\tilde{d}_t^i$: demand for product i in period t with support set $\left[\underline{d}_t^i, \bar{d}_t^i\right]$
d_t^i: demand realization for product i in period t
$\tilde{\mathbf{d}}_t^i$: collection of demands for product i from period 1 to period t, $\tilde{\mathbf{d}}_t^i := \left(\tilde{d}_1^i, \ldots, \tilde{d}_t^i\right) \in \mathbb{R}^t$
$\tilde{\mathbf{d}}_t$: collection of demands for all products from period 1 to period t, $\tilde{\mathbf{d}}_t := \left(\tilde{\mathbf{d}}_t^1, \ldots, \tilde{\mathbf{d}}_t^P\right) \in \mathbb{R}^{t \times P}$
$\mathbf{d}_t^i$: collection of demand realizations for product i from period 1 to period t, $\mathbf{d}_t^i := \left(d_1^i, \ldots, d_t^i\right) \in \mathbb{R}^t$
$\mathbf{d}_t$: collection of demand realizations for all products from period 1 to period t, $\mathbf{d}_t := \left(\mathbf{d}_t^1, \ldots, \mathbf{d}_t^P\right) \in \mathbb{R}^{t \times P}$
y_t^i: on-hand inventory level of product i at the start of period t
x_t^i: order quantity placed for product i at the start of period t
$\mathbf{y}_t$: collection of y_t^i for all products in period t, $\mathbf{y}_t := \left(y_t^1, \ldots, y_t^P\right) \in \mathbb{R}^P$
$\mathbf{x}_t$: collection of x_t^i for all products in period t, $\mathbf{x}_t := \left(x_t^1, \ldots, x_t^P\right) \in \mathbb{R}^P$
$\bar{x}_t$: maximum total order quantity for all products in period t
$\mathcal{N}_t$: the set of non-anticipative functions mapping $\mathbb{R}^{(t-1) \times P}$ to $\mathbb{R}$

Given the initial inventory level y_1^i, replenishment quantities $x_{1-l(i)}^i, x_{2-l(i)}^i, \ldots, x_{t-l(i)}^i$, and demand realizations $\mathbf{d}_t^i$, the inventory level at the start of period $t+1$ is a function $y_{t+1}^i : \mathbb{R} \times \mathbb{R}^t \times \mathbb{R}^t \to \mathbb{R}$,

$$y_{t+1}^i = y_1^i + \sum_{k=1}^{t} x_{k-l(i)}^i - \sum_{k=1}^{t} d_k^i. \tag{1}$$

Note that in Equation (1) the order quantities $\left\{x_{1-l(i)}^i, \ldots, x_0^i\right\}$ are given before period $t=1$.

For any variable $y \in \mathbb{R}$, define $(y)^+ := \max\{0, y\}$ and $(y)^- := \max\{0, -y\}$. Thus, $\left(y_{t+1}^i\right)^+$ and $\left(y_{t+1}^i\right)^-$ represent inventory overage and underage, respectively, at the end of period t after demand $\tilde{d}_t^i$

is realized. Given the replenishment quantity x_t^i and the inventory level y_{t+1}^i, the total cost incurred for product i in period t is $A_t^i I\left(x_t^i\right) + c_t^i x_t^i + h_t^i \left(y_{t+1}^i\right)^+ + b_t^i \left(y_{t+1}^i\right)^-$, where $I\left(x_t^i\right)$ equals 1 if $x_t^i > 0$, and equals 0 otherwise.

3 A stochastic optimization model

Given the stochastic demands $\tilde{\mathbf{d}}_t$, our goal is to determine an ordering policy $x_t^i\left(\tilde{\mathbf{d}}_{t-1}\right)$, for all $i \in \mathcal{P}$, $t \in \mathcal{T}$, such that the expected total cost within the planning horizon is minimized. We formulate this as a multi-period stochastic optimization problem as follows:

$$
\begin{aligned}
\min \quad & \mathrm{E}\left[\sum_{i\in\mathcal{P}}\sum_{t\in\mathcal{T}} A_t^i I\left(x_t^i\left(\tilde{\mathbf{d}}_{t-1}\right)\right) + c_t^i x_t^i\left(\tilde{\mathbf{d}}_{t-1}\right) + h_t^i\left(y_{t+1}^i\right)^+ + b_t^i\left(y_{t+1}^i\right)^-\right] \\
s.t. \quad & \sum_{i\in\mathcal{P}} x_t^i\left(\tilde{\mathbf{d}}_{t-1}\right) \le \bar{x}_t, && t\in\mathcal{T}; && \text{(2a)}\\
& x_t^i\left(\tilde{\mathbf{d}}_{t-1}\right) \ge 0, && i\in\mathcal{P},\ t\in\mathcal{T}; && \text{(2b)}\\
& x_t^i \in \mathcal{N}_t, && i\in\mathcal{P},\ t\in\mathcal{T}; && \text{(2c)}
\end{aligned}
$$

where $\mathcal{N}_t$ is a set containing all non-anticipative functions mapping $\mathbb{R}^{(t-1)\times P}$ to $\mathbb{R}$. Constraint (2a) represents the ordering capacity constraint. We use the shorthand $x\left(\tilde{\mathbf{d}}_{t-1}\right) \le a$ to denote that the random variable $x\left(\tilde{\mathbf{d}}_{t-1}\right)$ is less than or equal to a almost surely. Hence, it is noteworthy that all the constraints in Problem (2) must be satisfied for all realizations of the demands $\tilde{\mathbf{d}}_{t-1}$. By solving Problem (2), a non-anticipative replenishment policy that minimizes the expected total cost can be determined. However, Problem (2) is generally a difficult optimization problem to solve in practice (Shapiro, 2006).

Theoretically, we can use dynamic programming to solve Problem (2). Define $\mathbf{q}_t^i := (x_{t-l(i)}^i, \dots, x_{t-1}^i)$ as outstanding order quantities for product i in period t. Recall that $\mathbf{q}_1^i = \left(x_{1-l(i)}^i, \dots, x_0^i\right)$ are given before period $t = 1$. Let $\mathbf{q}_t := (\mathbf{q}_t^1, \dots, \mathbf{q}_t^P)$. Given the order quantities $\mathbf{x}_t$, inventory levels $\mathbf{y}_t$, outstanding order quantities $\mathbf{q}_t$, and demand realizations $\mathbf{d}_{t-1}$, the expected cost incurred in period t is

$$
r_t\left(\mathbf{x}_t, \mathbf{y}_t, \mathbf{q}_t; \mathbf{d}_{t-1}\right) = \sum_{i\in\mathcal{P}}\left(A_t^i I\left(x_t^i\right) + c_t^i x_t^i\right) + \mathrm{E}_{\tilde{\mathbf{d}}_t}\left[\sum_{i\in\mathcal{P}}\left(h_t^i\left(y_{t+1}^i\right)^+ + b_t^i\left(y_{t+1}^i\right)^-\right)\middle|\, \tilde{\mathbf{d}}_{t-1} = \mathbf{d}_{t-1}\right]. \quad (3)
$$

Let $J_t\left(\mathbf{y}_t, \mathbf{q}_t; \mathbf{d}_{t-1}\right)$ denote the optimal expected cost from period t until the end of planning horizon. The optimality equations can be written as

$$
J_t\left(\mathbf{y}_t, \mathbf{q}_t; \mathbf{d}_{t-1}\right) = \min_{\mathbf{x}_t\in\mathcal{F}_t}\left\{r_t\left(\mathbf{x}_t, \mathbf{y}_t, \mathbf{q}_t; \mathbf{d}_{t-1}\right) + \mathrm{E}_{\tilde{\mathbf{d}}_t}\left[J_{t+1}\left(\mathbf{y}_{t+1}, \mathbf{q}_{t+1}; \tilde{\mathbf{d}}_t\right)\middle|\, \tilde{\mathbf{d}}_{t-1} = \mathbf{d}_{t-1}\right]\right\}, \quad (4)
$$

where $\mathcal{F}_t = \{\mathbf{x}_t | \, x_t^i \geq 0, i \in \mathcal{P}, \; \sum_{i \in \mathcal{P}} x_t^i \leq \bar{x}_t\}$. The boundary conditions are $J_{T+1}\left(\mathbf{y}_{T+1}, \mathbf{q}_{T+1}; \mathbf{d}_T\right) = 0$, for all $\mathbf{y}_{T+1}, \mathbf{q}_{T+1}$, and $\mathbf{d}_T$.

For a special case with $l(i) = 0, i \in \mathcal{P}$, and $\tilde{d}_t^i, i \in \mathcal{P}, t \in \mathcal{T}$, are independent of each other, the optimality equations reduce to the following familiar form:

$$J_t\left(\mathbf{y}_t\right) = \min_{\mathbf{x}_t \in \mathcal{F}_t} \left\{ r_t\left(\mathbf{x}_t, \mathbf{y}_t\right) + \mathrm{E}_{\left(\tilde{d}_t^1, \cdots, \tilde{d}_t^P\right)}\left[J_{t+1}\left(\mathbf{y}_{t+1}\right)\right] \right\}, \tag{5}$$

with boundary conditions $J_{T+1}\left(\mathbf{y}_{T+1}\right) = 0$, for all $\mathbf{y}_{T+1}$.

Let $\mathbf{x}_t^*$ denote the optimal decision in period t for Problem (4). An optimal replenishment policy for the entire planning horizon can be determined by the sequence $\{\mathbf{x}_1^*, \ldots, \mathbf{x}_T^*\}$, which we call the *DP policy*. However, Problem (4) is generally intractable due to its large state space in practice.

4 A target-oriented robust optimization approach

Solving the stochastic inventory management problem (Problem (2)) using the dynamic programming method in Section 3 requires demand distributions to compute the expected total cost. Unfortunately, the demand distributions are often not available. Even if the demand distributions are known, it is not always possible to obtain the optimal solution due to computational complexity.

To overcome these issues, we propose a method based on target-oriented robust optimization to solve the inventory management problem. Under our approach, only a reference value (for example, the mean) and the support set of each product's demand in each period are required. Furthermore, our approach only requires solving a moderate-size mixed-integer linear program. This significantly reduces the computational complexity.

4.1 Formulation

We find a replenishment policy based on robust optimization. For each $\tilde{d}_t^i$, define an *adjustable uncertainty set*

$$D_t^i(\gamma) := \left\{ d_t^i \mid \hat{d}_t^i - \gamma \underline{z}_t^i \leq d_t^i \leq \hat{d}_t^i + \gamma \bar{z}_t^i \right\}, \tag{6}$$

where $\hat{d}_t^i$ represents a reference value (for example, the mean or median) of the demand for product i in period t, $\underline{z}_t^i := \hat{d}_t^i - \underline{d}_t^i$, $\bar{z}_t^i := \bar{d}_t^i - \hat{d}_t^i$, and $\gamma \in [0,1]$ is called the *uncertainty set parameter*. Define $\mathbf{D}_t^i(\gamma) := D_1^i(\gamma) \times \cdots \times D_t^i(\gamma)$ and $\mathbf{D}_t(\gamma) := \mathbf{D}_t^1(\gamma) \times \cdots \times \mathbf{D}_t^P(\gamma)$.

We introduce a *cost target* τ, which is pre-specified by the decision maker. Our goal is to determine a replenishment policy that maximizes the sizes of all the adjustable uncertainty sets such that all demand realizations from the sets will result in a total cost no more than τ. We can achieve this by solving the following optimization problem. Given a cost target τ, we maximize the uncertainty set parameter γ subject to the cost target constraint:

$$\gamma^* := \max \quad \gamma \tag{7a}$$

$$s.t. \quad \sum_{i\in\mathcal{P}}\sum_{t\in\mathcal{T}}\left(A_t^i I\left(x_t^i\left(\mathbf{d}_{t-1}\right)\right) + c_t^i x_t^i\left(\mathbf{d}_{t-1}\right) + h_t^i\left(y_{t+1}^i\right)^+ + b_t^i\left(y_{t+1}^i\right)^-\right) \leq \tau, \quad \forall \mathbf{d}_T \in \mathbf{D}_T(\gamma); \tag{7b}$$

$$\sum_{i\in\mathcal{P}} x_t^i\left(\mathbf{d}_{t-1}\right) \leq \bar{x}_t, \quad \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), \qquad t \in \mathcal{T}; \tag{7c}$$

$$x_t^i\left(\mathbf{d}_{t-1}\right) \in \mathcal{N}_t, \quad \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), \qquad i \in \mathcal{P},\ t \in \mathcal{T}; \tag{7d}$$

$$x_t^i\left(\mathbf{d}_{t-1}\right) \geq 0, \quad \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), \qquad i \in \mathcal{P},\ t \in \mathcal{T}; \tag{7e}$$

$$0 \leq \gamma \leq 1. \tag{7f}$$

The above is called the Target-Oriented Robust Optimization (TRO) formulation for the inventory management problem. Under this formulation, the decision variables are $x_t^i\left(\mathbf{d}_{t-1}\right), i \in \mathcal{P}, t \in \mathcal{T}$, and γ. Each $x_t^i\left(\mathbf{d}_{t-1}\right)$ is expressed as a non-anticipative function of $\mathbf{d}_{t-1}$. For convenience, define $x_1^i\left(\mathbf{d}_0\right) = x_1^i, i \in \mathcal{P}$. The first constraint represents the cost target constraint. The remaining constraints correspond to constraints in Problem (2). Note that besides the ordering capacity constraints (7c), we can easily incorporate other constraints (for example, warehouse capacity constraint) into the above formulation. Problem (7) determines an ordering policy $\{\mathbf{x}_t\}_{t\in\mathcal{T}}$ that costs no more than τ under the largest attainable adjustable uncertainty sets.

It is noteworthy that there may exist multiple optimal ordering policies for Problem (7). How do we choose from these multiple policies? If the cost target constraint (7b) is binding under an optimal solution for Problem (7), each optimal policy $\{\mathbf{x}_t\}_{t\in\mathcal{T}}$ gives the same worst-case total cost τ (with the extreme values of $\mathbf{d}_{t-1}$). We can choose any one of these optimal policies arbitrarily. If the cost target constraint (7b) is not binding (for example, when we are given a big budget such that τ equals a large value), different optimal policies for Problem (7) may give different worst-case total costs (although they are all no more than τ). In this situation, we choose an optimal policy $\{\mathbf{x}_t\}_{t\in\mathcal{T}}$ that minimizes the worst-case total cost.

4.2 A static rule

In Problem (7) each decision variable $x_t^i(\mathbf{d}_{t-1})$ belongs to the set $\mathcal{N}_t$ of non-anticipative functions of $\mathbf{d}_{t-1}$. In practice, it is intractable to consider all possible functions in $\mathcal{N}_t$. Instead, to overcome this issue the optimal policy $x_t^i(\mathbf{d}_{t-1})$ can be approximated as a static function, as an affine function (Ben-Tal et al., 2004), or as a piece-wise affine function (Chen et al., 2008; Wang et al., 2010) of the uncertainty variables $\mathbf{d}_{t-1}$. In this section, we approximate the optimal policy using a static function as $x_t^i(\mathbf{d}_{t-1}) = x_t^i$, for all $i \in \mathcal{P}, t \in \mathcal{T}$. We call this policy a static rule. We study an approximation of Problem (7) in which the static rule is optimal. It is worth noting that under the static rule although the ordering decisions cannot be adjusted as the information of the uncertainty variables is unveiled, the order quantity x_t^i for product i can be different for different t.

Another challenge of solving Problem (7) is caused by the $(\cdot)^+$ and $(\cdot)^-$ functions in Constraint (7b). There are various ways to represent these functions. Note that both $(y_{t+1}^i)^+$ and $(y_{t+1}^i)^-$ are scalars, and for any period t one of them must be zero. Problem (7) can be approximated by the following formulation.

$$
\begin{aligned}
\gamma'' := \max \quad & \gamma \\
s.t. \quad & \sum_{i\in\mathcal{P}}\sum_{t\in\mathcal{T}} \left(A_t^i I\left(x_t^i(\mathbf{d}_{t-1})\right) + c_t^i x_t^i(\mathbf{d}_{t-1}) + \theta_t^i\right) \le \tau, && \forall \mathbf{d}_{T-1} \in \mathbf{D}_{T-1}(\gamma); \\
& \theta_t^i \ge h_t^i \left(y_1^i + \sum_{k=1}^{t} \left(x_{k-l(i)}^i\left(\mathbf{d}_{k-l(i)-1}\right) - d_k^i \right) \right), \ \forall \mathbf{d}_t \in \mathbf{D}_t(\gamma), && i \in \mathcal{P},\ t \in \mathcal{T}; \quad (8a) \\
& \theta_t^i \ge -b_t^i \left(y_1^i + \sum_{k=1}^{t} \left(x_{k-l(i)}^i\left(\mathbf{d}_{k-l(i)-1}\right) - d_k^i \right) \right), \ \forall \mathbf{d}_t \in \mathbf{D}_t(\gamma), && i \in \mathcal{P},\ t \in \mathcal{T}; \quad (8b) \\
& \sum_{i\in\mathcal{P}} x_t^i(\mathbf{d}_{t-1}) \le \bar{x}_t,\ \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), && t \in \mathcal{T}; \\
& x_t^i(\mathbf{d}_{t-1}) \in \mathcal{N}_t,\ \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), && i \in \mathcal{P},\ t \in \mathcal{T}; \\
& x_t^i(\mathbf{d}_{t-1}) \ge 0,\ \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), && i \in \mathcal{P},\ t \in \mathcal{T}; \\
& 0 \le \gamma \le 1.
\end{aligned}
$$

It is worth noting that given any γ, the worst-case total cost of Problem (7) is no larger than the worst-case total cost of Problem (8). Thus, the cost target constraint is tighter in Problem (8), which leads to $\gamma'' \le \gamma^*$. This is because given any γ the worst-case total cost of Problem (7) is induced by a full vector of worst-case demand realizations $\mathbf{d}_{T-1}^* = \left(\mathbf{d}_{T-1}^{1*}, \ldots, \mathbf{d}_{T-1}^{P*}\right)$. In contrast, given any γ the worst-case total cost of Problem (8) comprises many different θ_t^i, each corresponds to a sub-vector of worst-case demand realizations $\breve{\mathbf{d}}_t^i$, for $i \in \mathcal{P},\ t \in \mathcal{T}$. Note that for each i each element of $\breve{\mathbf{d}}_t^i$ may vary with t.

We can tighten Constraint (8a) by replacing d_k^i with $\hat{d}_k^i - u_k^i$, and tighten Constraint (8b) by replacing d_k^i with $\hat{d}_k^i + v_k^i$, where u_k^i and v_k^i are uncertainty variables falling in $U_k^i(\gamma) := \left\{u_k^i | 0 \le u_k^i \le \gamma \underline{z}_k^i\right\}$ and $V_k^i(\gamma) := \left\{v_k^i | 0 \le v_k^i \le \gamma \bar{z}_k^i\right\}$ respectively. As a result, each d_t^i and $D_t^i(\gamma)$ in the above formulation can be represented as $d_t^i = d_t^i\left(u_t^i, v_t^i\right) = \hat{d}_t^i - u_t^i + v_t^i$ and $D_t^i(\gamma) = \left\{\hat{d}_t^i - u_t^i + v_t^i \mid u_t^i \in U_t^i(\gamma), v_t^i \in V_t^i(\gamma)\right\}$ respectively. Thus, each decision variable $x_t^i\left(\mathbf{d}_{t-1}\right)$ can be expressed as a function of u_k^i and v_k^i, $i \in \mathcal{P}$, $k = 1, \ldots, t-1$. For notational convenience, define $\mathbf{u}_t^i(\gamma) := \left(u_1^i(\gamma), \ldots, u_t^i(\gamma)\right)$, $\mathbf{u}_t(\gamma) := \left(\mathbf{u}_t^1(\gamma), \ldots, \mathbf{u}_t^P(\gamma)\right)$, $\mathbf{v}_t^i(\gamma) := \left(v_1^i(\gamma), \ldots, v_t^i(\gamma)\right)$, and $\mathbf{v}_t(\gamma) := \left(\mathbf{v}_t^1(\gamma), \ldots, \mathbf{v}_t^P(\gamma)\right)$. Define $\mathbf{U}_t^i(\gamma) := U_1^i(\gamma) \times \cdots \times U_t^i(\gamma)$ and $\mathbf{U}_t(\gamma) := \mathbf{U}_t^1(\gamma) \times \cdots \times \mathbf{U}_t^P(\gamma)$. Similarly, define $\mathbf{V}_t^i(\gamma) := V_1^i(\gamma) \times \cdots \times V_t^i(\gamma)$ and $\mathbf{V}_t(\gamma) := \mathbf{V}_t^1(\gamma) \times \cdots \times \mathbf{V}_t^P(\gamma)$. Problem (8) can be approximated as

$$
\begin{aligned}
\gamma' := {} & \max \quad \gamma && (9)\\
\text{s.t.} \quad & \sum_{i\in\mathcal{P}}\sum_{t\in\mathcal{T}} \left(A_t^i I\left(x_t^i\left(\mathbf{d}_{t-1}\right)\right) + c_t^i x_t^i\left(\mathbf{d}_{t-1}\right) + \theta_t^i\right) \le \tau,\ \forall \mathbf{u}_{T-1} \in \mathbf{U}_{T-1}(\gamma), \forall \mathbf{v}_{T-1} \in \mathbf{V}_{T-1}(\gamma); && \\
& \theta_t^i \ge h_t^i\left(y_1^i - \sum_{k=1}^{t} \hat{d}_k^i + \sum_{k=1}^{t}\left(x_{k-l(i)}^i\left(\mathbf{d}_{k-l(i)-1}\right) + u_k^i\right)\right),\ \forall \mathbf{u}_t \in \mathbf{U}_t(\gamma), \forall \mathbf{v}_t \in \mathbf{V}_t(\gamma), && i \in \mathcal{P}, t \in \mathcal{T};\\
& \theta_t^i \ge -b_t^i\left(y_1^i - \sum_{k=1}^{t} \hat{d}_k^i + \sum_{k=1}^{t}\left(x_{k-l(i)}^i\left(\mathbf{d}_{k-l(i)-1}\right) - v_k^i\right)\right),\ \forall \mathbf{u}_t \in \mathbf{U}_t(\gamma), \forall \mathbf{v}_t \in \mathbf{V}_t(\gamma), && i \in \mathcal{P}, t \in \mathcal{T};\\
& \sum_{i\in\mathcal{P}} x_t^i\left(\mathbf{d}_{t-1}\right) \le \bar{x}_t,\ \forall \mathbf{u}_{t-1} \in \mathbf{U}_{t-1}(\gamma), \forall \mathbf{v}_{t-1} \in \mathbf{V}_{t-1}(\gamma), && t \in \mathcal{T};\\
& x_t^i\left(\mathbf{d}_{t-1}\right) \in \mathcal{N}_t,\ \forall \mathbf{u}_{t-1} \in \mathbf{U}_{t-1}(\gamma), \forall \mathbf{v}_{t-1} \in \mathbf{V}_{t-1}(\gamma), && i \in \mathcal{P}, t \in \mathcal{T};\\
& x_t^i\left(\mathbf{d}_{t-1}\right) \ge 0,\ \forall \mathbf{u}_{t-1} \in \mathbf{U}_{t-1}(\gamma), \forall \mathbf{v}_{t-1} \in \mathbf{V}_{t-1}(\gamma), && i \in \mathcal{P}, t \in \mathcal{T};\\
& 0 \le \gamma \le 1.
\end{aligned}
$$

Recall that $d_t^i = d_t^i\left(u_t^i, v_t^i\right) = \hat{d}_t^i - u_t^i + v_t^i$ in the above formulation. Since the second and third constraints of Problem (9) are more restrictive than Constraints (8a) and (8b), we have $\gamma' \le \gamma'' \le \gamma^*$.

It is more economic to solve Problem (9) because the static rule $x_t^i\left(\mathbf{d}_{t-1}\right) = x_t^i$ is optimal under such an approximation. Thus, no complicated decision rules need to be considered. To show this result, define a vector $w_t^i := \left(u_t^i,\ v_t^i\right) \in U_t^i(\gamma) \times V_t^i(\gamma)$ and let $\mathbf{w} = \left(w_1^1, \ldots, w_T^1, \ldots, w_1^P, \ldots, w_T^P\right)$. Note that $\mathbf{w}$ represents a collection of uncertainty variables u_t^i and v_t^i, for $i \in \mathcal{P}$ and $t \in \mathcal{T}$. Let $\mathbf{W}(\gamma)$ denote the support set of $\mathbf{w}$. Let $\boldsymbol{\pi}(\mathbf{w})$ denote a vector that contains all decision variables of Problem (9), and let

$\mathcal{F}$ denote the feasible ranges of these variables. We can rewrite Problem (9) in the following general form:

$$
\begin{aligned}
\gamma' = \max \ & \gamma && (10)\\
s.t. \ & \mathbf{A}(\mathbf{w})\boldsymbol{\pi}(\mathbf{w}) \le \mathbf{b}(\mathbf{w}), \ \forall \mathbf{w} \in \mathbf{W}(\gamma);\\
& \boldsymbol{\pi}(\mathbf{w}) \in \mathcal{F}, \ \forall \mathbf{w} \in \mathbf{W}(\gamma);
\end{aligned}
$$

where $\mathbf{A}(\mathbf{w})$ and $\mathbf{b}(\mathbf{w})$ represent all deterministic and uncertain coefficients of Problem (9). It is noteworthy that $\boldsymbol{\pi}(\mathbf{w})$ is a decision rule. We consider the static rule $\boldsymbol{\pi}(\mathbf{w}) = \boldsymbol{\pi}$, which can be determined by solving the following problem:

$$
\begin{aligned}
\gamma^s := \max \ & \gamma && (11)\\
s.t. \ & \mathbf{A}(\mathbf{w})\boldsymbol{\pi} \le \mathbf{b}(\mathbf{w}), \ \forall \mathbf{w} \in \mathbf{W}(\gamma);\\
& \boldsymbol{\pi} \in \mathcal{F}.
\end{aligned}
$$

We will show that an optimal solution of Problem (11) is also optimal for Problem (10). In other words, the static rule is an optimal decision rule for Problem (10). We first give the following definition.

Definition 1 (Worst-case scenario of uncertainty). *Given* $\mathbf{A}(\mathbf{w})$ *and* $\mathbf{b}(\mathbf{w})$*, an element* $\breve{\mathbf{w}}(\gamma)$ *in the set* $\mathbf{W}(\gamma)$ *is called a worst-case scenario of uncertainty if for each* $\boldsymbol{\pi} \in \mathcal{F}$ *that satisfies* $\mathbf{A}(\breve{\mathbf{w}}(\gamma))\boldsymbol{\pi} \le \mathbf{b}(\breve{\mathbf{w}}(\gamma))$*, it also satisfies* $\mathbf{A}(\mathbf{w})\boldsymbol{\pi} \le \mathbf{b}(\mathbf{w}), \ \forall \mathbf{w} \in \mathbf{W}(\gamma)$*.*

According to Definition 1, for any $\gamma \in [0,1]$, once we identify a feasible solution for Problem (11) under a worst-case scenario of uncertainty, then the solution is also feasible for any other realizations of the uncertainty variables. We use this property to show that the static rule is optimal for Problem (10).

Theorem 1 (Optimality of the static rule). *Suppose for any* $\gamma \in [0,1]$ *Problem (11) has a worst-case scenario of uncertainty, denoted by* $\breve{\mathbf{w}}(\gamma) \in \mathbf{W}(\gamma)$*. Let* $\boldsymbol{\pi}^\dagger$ *denote a static rule that represents the solution of the following deterministic optimization problem:*

$$
\begin{aligned}
\gamma^\dagger := \max \ & \gamma && (12)\\
s.t. \ & \mathbf{A}(\breve{\mathbf{w}}(\gamma))\boldsymbol{\pi} \le \mathbf{b}(\breve{\mathbf{w}}(\gamma));\\
& \boldsymbol{\pi} \in \mathcal{F}.
\end{aligned}
$$

The static rule $\boldsymbol{\pi}^\dagger$ *is also optimal for Problem (10) and* $\gamma^\dagger = \gamma^s = \gamma'$*.*

Proof. Clearly, we have $\gamma^s \leq \gamma'$. We can also observe that

$$\begin{aligned}
\gamma^s &= \max\{\gamma : \mathbf{A}(\mathbf{w})\boldsymbol{\pi} \leq \mathbf{b}(\mathbf{w}),\ \forall \mathbf{w} \in \mathbf{W}(\gamma),\ \boldsymbol{\pi} \in \mathcal{F}\} \\
&\geq \max\{\gamma : \mathbf{A}(\breve{\mathbf{w}}(\gamma))\boldsymbol{\pi} \leq \mathbf{b}(\breve{\mathbf{w}}(\gamma)),\ \boldsymbol{\pi} \in \mathcal{F}\} \\
&= \gamma^\dagger \\
&= \max\{\gamma : \mathbf{A}(\breve{\mathbf{w}}(\gamma))\boldsymbol{\pi}(\breve{\mathbf{w}}(\gamma)) \leq \mathbf{b}(\breve{\mathbf{w}}(\gamma)),\ \boldsymbol{\pi}(\breve{\mathbf{w}}(\gamma)) \in \mathcal{F}\} \\
&\geq \max\{\gamma : \mathbf{A}(\mathbf{w})\boldsymbol{\pi}(\mathbf{w}) \leq \mathbf{b}(\mathbf{w}),\ \boldsymbol{\pi}(\mathbf{w}) \in \mathcal{F},\ \forall\ \mathbf{w} \in \mathbf{W}(\gamma)\} \\
&= \gamma',
\end{aligned}$$

where the first inequality is due to the definition of the worst-case scenario of uncertainty and the last inequality is due to the fact that $\breve{\mathbf{w}}(\gamma) \in \mathbf{W}(\gamma)$. Note that the static rule $\boldsymbol{\pi}^\dagger$ is optimal for Problem (12) under the worst-case scenario of uncertainty $\breve{\mathbf{w}}\left(\gamma^\dagger\right)$. Thus, $\boldsymbol{\pi}^\dagger$ is feasible for Problem (10). Since $\boldsymbol{\pi}^\dagger$ achieves the optimal objective γ', the static rule is also optimal for Problem (10). $\square$

Theorem 1 implies that if Problem (10) has a worst-case scenario of uncertainty for any γ, then it can be solved by just considering the static rule and the worst-case scenario of uncertainty for any γ. This significantly reduces the computation burden because no complicated decision rules are required. Since Problem (10) has a general form, this result is not limited to the inventory management problem.

Thus, to solve Problem (9) we only need to consider the static rule $x_t^i\left(\mathbf{d}_{t-1}\right) = x_t^i$ and the worst-case scenario of uncertainty for any γ. The solution of Problem (9) can be determined by solving the following mixed-integer program:

$$\begin{aligned}
\gamma^\dagger = \max\ & \gamma && (14) \\
s.t.\ & \sum_{i \in \mathcal{P}} \sum_{t \in \mathcal{T}} \left(A_t^i I\left(x_t^i\right) + c_t^i x_t^i + \theta_t^i\right) \leq \tau, \\
& \theta_t^i \geq h_t^i \left(y_1^i - \sum_{k=1}^{t} \hat{d}_k^i + \sum_{k=1}^{t} \left(x_{k-l(i)}^i + \gamma \underline{z}_k^i\right)\right), && \forall\ i \in \mathcal{P},\ t \in \mathcal{T}; \\
& \theta_t^i \geq -b_t^i \left(y_1^i - \sum_{k=1}^{t} \hat{d}_k^i + \sum_{k=1}^{t} \left(x_{k-l(i)}^i - \gamma \bar{z}_k^i\right)\right), && \forall\ i \in \mathcal{P},\ t \in \mathcal{T}; \\
& \sum_{i \in \mathcal{P}} x_t^i \leq \bar{x}_t, && \forall\ t \in \mathcal{T}; \\
& x_t^i \geq 0, && \forall\ i \in \mathcal{P},\ t \in \mathcal{T}; \\
& 0 \leq \gamma \leq 1.
\end{aligned}$$

In the following sections, we use Problem (14) to approximate Problem (7). The former is a mixed-integer program with an optimal objective value $\gamma^\dagger$ and the latter is a robust optimization problem with an optimal objective value $\gamma^* \geq \gamma^\dagger$. Problem (14) can be solved by using any commercial or open-source solvers. Furthermore, if the fixed ordering cost $A_t^i = 0$ for all $i \in \mathcal{P}$ and $t \in \mathcal{T}$, then Problem (14) reduces to a linear program. We call the solution to Problem (14) the *TRO policy*.

4.3 Target coefficient

The pre-specified cost target τ is an important parameter in our approach. If τ is too small, it may cause Problem (14) to be infeasible or may lead to a replenishment policy with weak robustness (because $\gamma^\dagger$ is too small and thus the resultant uncertainty sets $D_t^i\left(\gamma^\dagger\right)$ are too small). On the other hand, an overly large τ may result in a replenishment policy that is too conservative ($D_t^i\left(\gamma^\dagger\right)$ are too large). In practice, it can be tedious to choose a proper value for τ. This involves some estimation of the total cost for the entire planning horizon. Furthermore, the total cost depends on the initial inventory levels, which change over time. To overcome this issue, we define *target coefficient* as

$$\alpha := \frac{\rho(1) - \tau}{\rho(1) - \rho(0)}, \tag{15}$$

where $\rho(\gamma)$, $\gamma = 0, 1$, can be determined by solving the following mixed-integer program:

$$\begin{aligned}
\rho(\gamma) := \min \quad & \sum_{i\in\mathcal{P}} \sum_{t\in\mathcal{T}} \left(A_t^i I\left(x_t^i\right) + c_t^i x_t^i + \theta_t^i\right) \\
s.t. \quad & \theta_t^i \geq h_t^i \left(y_1^i - \sum_{k=1}^{t} \hat{d}_k^i + \sum_{k=1}^{t} \left(x_{k-l(i)}^i + \gamma \underline{z}_k^i \right) \right), && \forall\, i \in \mathcal{P},\ t \in \mathcal{T}; \\
& \theta_t^i \geq -b_t^i \left(y_1^i - \sum_{k=1}^{t} \hat{d}_k^i + \sum_{k=1}^{t} \left(x_{k-l(i)}^i - \gamma \bar{z}_k^i \right) \right), && \forall\, i \in \mathcal{P},\ t \in \mathcal{T}; \\
& \sum_{i\in\mathcal{P}} x_t^i \leq \bar{x}_t, && \forall\, t \in \mathcal{T}; \\
& x_t^i \geq 0, && \forall\, i \in \mathcal{P},\ t \in \mathcal{T}.
\end{aligned}$$

Given a target coefficient $\alpha \in [0, 1]$, the corresponding cost target is $\tau(\alpha) = (1-\alpha)\rho(1) + \alpha\rho(0)$, which falls in the interval $[\rho(0), \rho(1)]$.

It is useful to define target coefficient because Problem (14) is often solved in a rolling horizon manner. Since the initial inventory levels change from one planning horizon to the next planning horizon, the range of τ, $[\rho(0), \rho(1)]$ may vary accordingly. Thus, by defining target coefficient we can conveniently stick to a constant choice for α (say, $\alpha = 0.7$) without knowing $\rho(0)$ and $\rho(1)$.

4.4 A special case

We consider a special case with one product over a single period, and drop all the superscripts and subscripts for notational simplicity in this section. The TRO policy in this case is given as follows.

Theorem 2 (The TRO policy for the single-product, single-period problem). *Consider a single-product single-period inventory management problem with zero lead time, zero fixed ordering cost, variable ordering cost c, unit holding cost h, unit backlog cost b ($b > c$), stochastic demand with a reference value $\hat{d}$ and support set $[\hat{d} - \underline{z}, \hat{d} + \bar{z}]$, and no ordering capacity constraint. Under this setting, Problem (14) is equivalent to Problem (7) and their optimal objective value is*

$$\gamma^* = \begin{cases} \left(\tau + cy - c\hat{d}\right)(b+h)/\left((c+h)b\bar{z} + (b-c)h\underline{z}\right), & \underline{y} \le y < y_a; \\ 1, & y_a \le y \le y_b; \\ \left(\tau - hy + h\hat{d}\right)/h\underline{z}, & y_b < y \le \bar{y}; \end{cases} \tag{16}$$

where $\underline{y} = \hat{d} - \tau/c$, $y_a = \hat{d} - \tau/c + \left((c+h)b\bar{z} + (b-c)h\underline{z}\right)/(cb+ch)$, $y_b = \hat{d} + \tau/h - \underline{z}$, and $\bar{y} = \hat{d} + \tau/h$. The corresponding replenishment policy is a state-dependent order-up-to policy:

$$x_{opt} = \max\left\{Y(\gamma^*) - y, 0\right\}, \tag{17}$$

where $Y(\gamma^*) = \hat{d} + \gamma^*(b\bar{z} - h\underline{z})/(h+b)$ represents a state-dependent order-up-to level.

Proof. See Appendix A. □

Figure 1 shows the objective γ^*, the order quantity x_{opt}, and the inventory level $y + x_{opt}$ for the single-product, single-period problem. We set $\tau = 5.3$, $c = 2$, $h = 1.2$, $b = 4$, and $\tilde{d}$ follows a beta distribution $Beta(2,5)$ on $[0,5]$. The horizontal axis of Figure 1 can be divided into four intervals. Note that $\gamma^* = 1$ in the second and third intervals (corresponding to Cases (a.1) and (a.2), respectively, in the proof in Appendix A). In contrast, $\gamma^* < 1$ in the first and fourth intervals (corresponding to Cases (b.1) and (b.2), respectively, in the proof in Appendix A).

As the initial inventory level y increases, γ^* first increases until it hits its maximum value 1. It remains at this maximum value until it begins to drop when $y \ge y_b \approx 4.4$ as we enter the fourth interval (see Equation (31) in Appendix A for an explanation). On the other hand, the order quantity x_{opt} decreases as y increases. Note that x_{opt} decreases with a larger rate when γ^* hits the value 1, and the order quantity finally becomes 0.

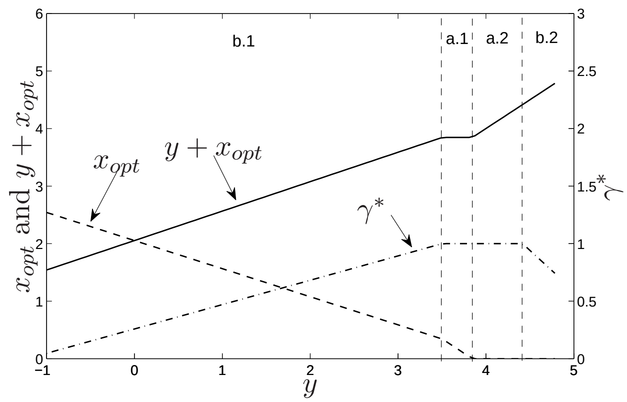

Figure 1: **Solution structure of the single-product, single-period problem.**

5 Numerical studies

We use numerical simulations to compare the performance of the TRO policy in Section 4.2 with that of the DP policy in Section 3 and of a *myopic policy* described as follows. Given the inventory levels $\mathbf{y}_t$ and outstanding order quantities $\mathbf{q}_t$ in each period t, the myopic policy adopts the order quantities $x_1^i, i \in \mathcal{P}$, which are obtained by solving Problem (2) with $T = \max_{1 \le i \le P} l(i) + 1$ and $x_2^i = x_3^i = \cdots = x_T^i = 0, i \in \mathcal{P}$:

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{P}} \left(A_1^i I\left(x_1^i\right) + c_1^i x_1^i \right) + \mathrm{E}_{\tilde{\mathbf{d}}_T} \left[\sum_{i \in \mathcal{P}} \sum_{t=1}^{l(i)+1} \left(h_t^i \left(y_{t+1}^i\right)^+ + b_t^i \left(y_{t+1}^i\right)^- \right) \right] \\ s.t. \quad & \sum_{i \in \mathcal{P}} x_1^i \le \bar{x}_1; \\ & x_1^i \ge 0, \;\; i \in \mathcal{P}. \end{aligned} \tag{18}$$

The expected value of the objective function of Problem (18) can be determined based on demand distributions if they are available. Otherwise, it can be computed using sample average approximation.

Throughout this section, let N denote the number of simulated periods in each simulation, which could be different from the length of planning horizon T in each approach. In Sections 5 and 6, we assume $\hat{d}_t^i$ represents the mean of the demand for product i in period t in Equation (6).

5.1 The multi-product, single-period problem

We compare the TRO policy (based on Problem (14)) with the DP policy (based on Equations (4)) and the myopic policy (based on Problem (18)) for a multi-product, single-period problem. We set $P = 10$, $T = 1$, $l(i) = 0$, $\bar{x}_t = 15$, $A_t^i = 1$, $c_t^i = 2$, $h_t^i = 1 + 0.2i$, and $b_t^i = 3i$, for all $i \in \mathcal{P}$, $t \in \mathcal{T}$. We assume

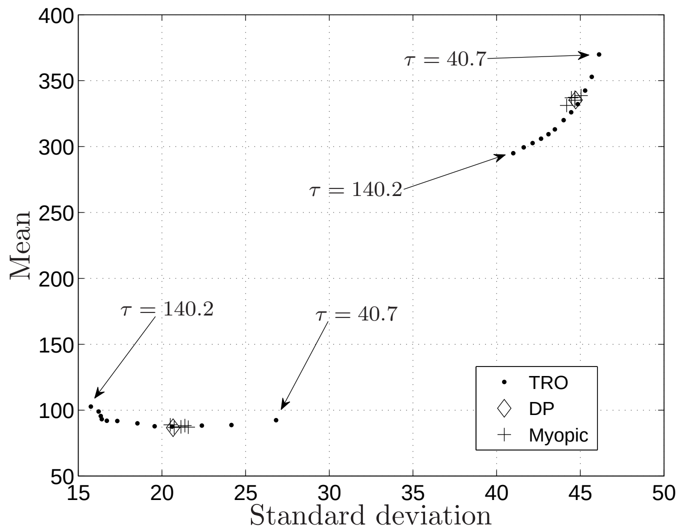

Figure 2: **Performance of different policies for the multi-product, single-period problem.**

the demand for each product in each period falls in $[0, 5]$, and it follows a beta distribution $Beta(2, 5)$. After obtaining the solution of each approach, we do simulations to evaluate them.

Figure 2 shows the results of simulations with $N = 1$. Each data point in the graph represents the mean and standard deviation of the costs of 1,000 simulations. The dots represent the results of the TRO policy. Each dot is generated using a distinct value of τ. The diamonds and the crosses represent the results of the DP policy and the myopic policy respectively. To see the effects of limited demand information, the myopic policy is determined based on sample average approximation using a sample set with cardinality 50. Each cross corresponds to a distinct sample set.

The bottom-left corner of Figure 2 shows the simulation results under different approaches when we use the same demand distribution $Beta(2, 5)$ in the simulations. The results suggest that, for similar standard deviations, the mean costs of the TRO policy are very close to that of the DP policy and of the myopic policy. This is surprising because only limited demand information ($\hat{d}_t^i$, $\bar{z}_t^i$, and $\underline{z}_t^i$, for all i and t), excluding the demand distribution, is used to obtain the TRO policy. In contrast, the other two policies require the information of the demand distribution $Beta(2, 5)$.

More importantly, the TRO policy provides additional flexibility that allows us to tune the cost target τ. This flexibility is promising because one can balance the cost and its associated risk by properly choosing τ. For example, when $\tau = 140.2$ the mean cost of the TRO policy is roughly 10% more than that of the DP policy and of the myopic policy, but its standard deviation of costs is 25% lower. This suggests that we can substantially reduce the standard deviation of costs of the TRO policy

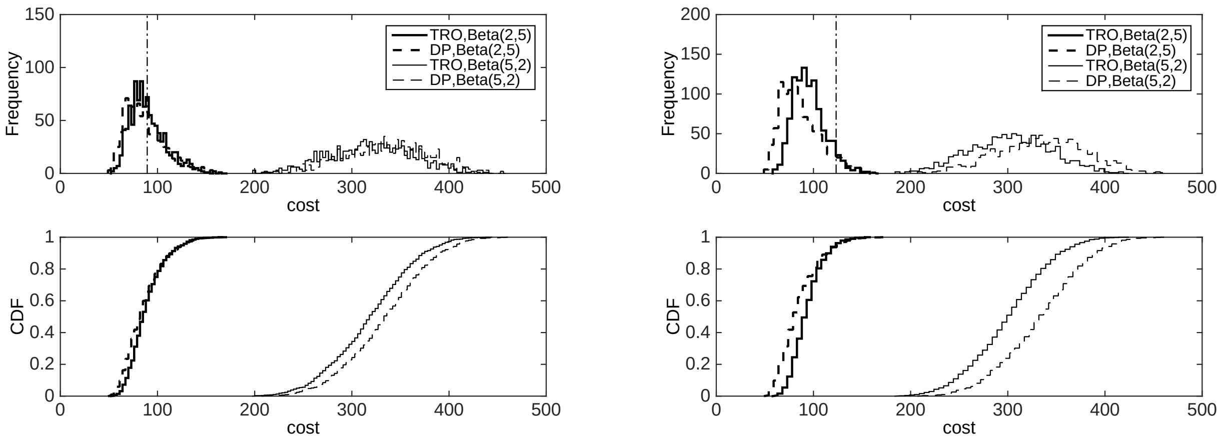

(a) (b)

Figure 3: **Histograms and cumulative distributions of simulation costs.**

by tuning τ without increasing its mean cost too much. In contrast, neither the DP policy nor the myopic policy provides this degree of freedom.

To show the robustness of different approaches against distributional ambiguity, we repeat the simulations using a different demand distribution $Beta(5,2)$. The results are shown in the top-right corner of Figure 2. These results suggest that if a wrong demand distribution is used for solving Equations (4) and Problem (18), the DP policy and the myopic policy, respectively, perform much worse than the TRO policy. In most cases, both mean and standard deviation of costs provided by the TRO policy are significantly lower than that of the other two policies.

To compare the TRO policy and the DP policy in more detail, we plot the distribution of simulation costs under each policy. We set $\tau = 89.4$ for the TRO policy. The top graph of Figure 3(a) shows the histograms of costs under the two policies. The bold solid line and the bold dashed line correspond to the TRO policy and the DP policy, respectively, when we use the demand distribution $Beta(2,5)$ for the simulations. The distribution of costs under the DP policy is skewed toward the left compared to the distribution under the TRO policy. However, the result is reversed when we use a different demand distribution for the simulations. The right side of the same graph shows the histograms under the two policies using the demand distribution $Beta(5,2)$ for the simulations. The solid line and the dashed line correspond to the TRO policy and the DP policy respectively. Apparently, the costs under the TRO policy are generally lower compared to the DP policy, which assumes an inaccurate demand distribution. This demonstrates the robustness of the TRO policy against distributional ambiguity.

The bottom graph of Figure 3(a) shows the cumulative distribution functions of costs under the two policies. The left side of the graph shows the results when the DP policy assumes an accurate demand distribution. Note that as the cost increases although the DP policy dominates initially, the two policies give approximately the same cumulative probability at the cost target $\tau = 89.4$ (corresponding to the vertical dashed line). This suggests that the percentage of simulations with costs less than τ is similar under both policies. On the other hand, the right side of the same graph suggests that if the DP policy assumes an inaccurate demand distribution, the simulation cost under the TRO policy is first-order stochastically dominated by that of the DP policy.

Figures 3(b) shows the histograms and the cumulative distribution functions of costs under both approaches when we set $\tau = 123.3$ for the TRO policy. If the DP policy assumes an inaccurate demand distribution, the TRO policy significantly outperforms the DP policy for this value of τ.

5.2 The single-product, multi-period problem

To evaluate the TRO policy in more detail, we compare it with the DP policy with an accurate demand distribution (which gives the optimal expected cost). We focus on the single-product, multi-period problem so that we can compute the DP policy within a reasonable amount of time. We assume demand for the product in each period follows a beta distribution $Beta(2, 2)$ in the interval $[0, 5]$, and we use the same distribution in the simulations. We set $T = 10$, $l(1) = 1$, and $\bar{x}_t = 5$, for $t \in \mathcal{T}$. We adopt a rolling-horizon principle to determine the TRO policy in each period: After demand is realized in each period of a simulation, we use the new inventory level to resolve Problem (14) to find the order quantity for the next period. Table 2 shows the simulation results.

Each row of Table 2 corresponds to a setting of cost parameters. Each entry in the table shows the average value and standard deviation (in brackets) of the costs of 1,000 simulations. The TRO policy can be optimized by fine tuning the target coefficient α. The lowest cost of the TRO policy for each setting of cost parameters is marked with an asterisk, and the highest cost is marked with a dagger sign. The relative gap between the lowest cost and the optimal cost by the DP policy is shown in the second-to-last column of Table 2. The last column of the table shows the relative gap between the highest cost of the TRO policy and the optimal cost. Our results suggest that the cost of the TRO policy is within 120% of the optimal cost.

Table 2 shows that the lowest costs of the TRO policy occur when α is between 0.5 and 0.8. This suggests that we should set the target coefficient α in the interval $[0.5, 0.8]$ for this problem instance. The lowest cost of the TRO policy for each parameter setting is within 104% of the optimal cost.

Table 2: **Performance of the TRO and DP policies for the single-product, multi-period problem**

$(A_t^1, c_t^1, h_t^1, b_t^1)$	**DP policy**	**TRO policy**									**Relative gaps**	
		$\alpha = 0.1$	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	min	max
$(1, 2, 0.1, 3)$	55.37	66.16†	65.16	63.24	61.74	60.47	59.30	58.26	57.59*	58.04	4.00%	19.49%
	(7.84)	(6.24)	(6.34)	(6.62)	(6.74)	(6.87)	(7.12)	(7.75)	(8.24)	(9.31)		
$(1, 1, 0.1, 5)$	34.66	40.54†	39.65	38.24	37.36	36.76	36.07	35.73*	36.26	37.88	3.08%	16.96%
	(5.40)	(4.11)	(4.28)	(4.46)	(4.57)	(4.73)	(4.91)	(5.55)	(7.12)	(9.04)		
$(1, 2, 1, 3)$	71.42	80.54†	78.72	77.08	75.74	74.78	74.20	74.10*	74.55	75.38	3.75%	12.77%
	(8.64)	(4.11)	(4.62)	(5.34)	(6.18)	(7.15)	(8.31)	(9.43)	(10.85)	(12.05)		
$(1, 1, 1, 5)$	54.50	63.71†	61.20	58.77	57.10	56.03	55.74*	56.48	58.30	61.95	2.27%	16.90%
	(7.24)	(4.30)	(4.36)	(4.66)	(5.37)	(6.42)	(7.96)	(9.96)	(12.20)	(14.99)		
$(1, 1, 2, 5)$	67.87	75.07†	72.94	71.17	69.78	68.87	68.44*	68.82	69.83	71.46	0.84%	10.61%
	(9.83)	(8.50)	(8.42)	(8.51)	(8.81)	(9.36)	(10.18)	(11.61)	(13.16)	(14.75)		
$(1, 1, 3, 2)$	59.56	61.45†	61.11	60.86	60.71	60.66*	60.71	60.87	61.12	61.42	1.84%	3.18%
	(8.07)	(9.31)	(9.02)	(8.75)	(8.52)	(8.33)	(8.19)	(8.09)	(8.04)	(8.01)		
$(1, 1, 5, 2)$	66.89	71.34	69.92	69.30	68.39	67.92*	67.95	68.73	70.13	72.11†	1.54%	7.80%
	(9.79)	(12.24)	(11.61)	(10.96)	(10.46)	(10.20)	(10.21)	(10.57)	(11.33)	(12.32)		
$(10, 2, 1, 3)$	110.41	120.87†	119.24	119.65	116.12	114.39	113.48	112.94	112.54*	112.77	1.93%	9.47%
	(13.78)	(9.67)	(10.11)	(10.25)	(10.80)	(11.11)	(11.57)	(12.01)	(12.85)	(13.56)		

†: The highest mean cost. $*$: The lowest mean cost.

5.3 The multi-product, multi-period problem

We also evaluate the TRO policy in a multi-product, multi-period setting with $P = 3$ and $T = 5$. To ensure tractability, we assume that demand for each product in each period has only three possible values 0, 2.5, and 5 with probabilities 0.6, 0.3, and 0.1 respectively. We consider products $i = 1, 2, 3$ of the problem in Section 5.1 and assume all other parameters remain unchanged. For the TRO policy, we set $\alpha = 0.2, 0.3, \ldots, 0.7$. For each value of α, we adopt the same rolling-horizon principle described in Section 5.2 to determine the TRO policy in each period of every simulation. Figure 4(a) shows the mean and standard deviation of the costs of 1,000 simulations under the TRO and the DP policies.

The left side of the graph shows the results when the presumed demand distribution: $Pr(0) = 0.6$, $Pr(2.5) = 0.3$, and $Pr(5) = 0.1$ is used for the simulations. The TRO policy gives lower standard deviations but slightly higher means compared to the DP policy. On the right side of the graph, we plot the results of the simulations with a different demand distribution: $Pr(0) = 0.1$, $Pr(2.5) = 0.3$, and $Pr(5) = 0.6$. The results suggest that if the DP policy assumes an inaccurate demand distribution, the TRO policy always gives a lower mean cost and sometimes also a lower standard deviation.

It is worth noting that we control the target coefficient α (instead of cost target τ) in Figure 4(a). This is because in each simulation we use the rolling-horizon principle to determine the TRO policy for each period. The planning horizon of Problem (14) reduces as we progress in time. Thus, given a specific α, the values of $\rho(0)$ and $\rho(1)$ in Equation (15) may change from period to period. This

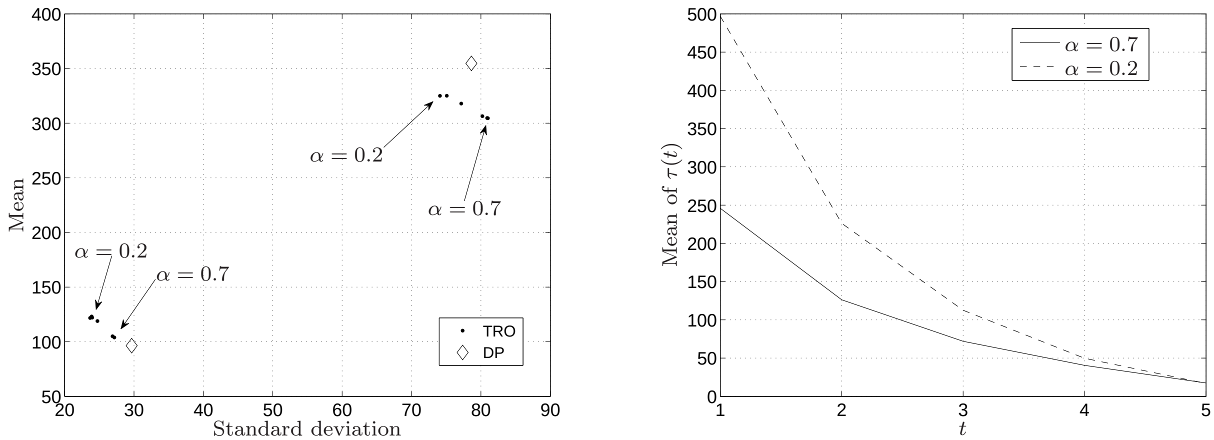

(a) (b)

Figure 4: **Comparison of the TRO and the DP policies in the multi-product, multi-period problem.**

may result in different values of τ in different periods. Let $\tau(t)$ denote the cost target in period t. In addition, the initial inventory levels for period t may vary across different simulations. This will also lead to different values of $\tau(t)$ in different simulations for the same period t. Figure 4(b) shows that the average value of $\tau(t)$ over all simulations decreases with t. This is consistent with our intuition: As t increases, the planning horizon of Problem (14) reduces, which results in a lower cost target.

5.4 A robustness comparison

In order to test the robustness of the TRO policy, we compare it with the DP policy over a family of demand distributions using the same problem instance in Section 5.1. When we compute the DP policy we assume demand for each product in each period follows the beta distribution $Beta(4, 4)$ in the interval $[0, 5]$. For the simulations, we use the demand distribution $Beta(\eta, \eta)$ in $[0, 5]$ where $0.1 \leq \eta \leq 6$. Under this distribution, the demand for each product in each period has mean equal to 2.5 for all η. Based on the recommendations of Table 2, we set $\alpha = 0.7$ for the TRO policy.

Figure 5(a) shows the mean and standard deviation of the costs of 1,000 simulations under each policy for each η. The results of the two policies corresponding to the same value of η are connected by a dotted line. Figure 5(a) suggests that if the demand distributions used in the simulations are close to the one assumed by the DP policy (that is, for distributions that are similar to the one shown in Figure 5(b) with $\eta = 4$), the DP policy outperforms the TRO policy in mean cost. However, the latter gives a lower standard deviation of costs for most of the values of η. If the demand distributions used

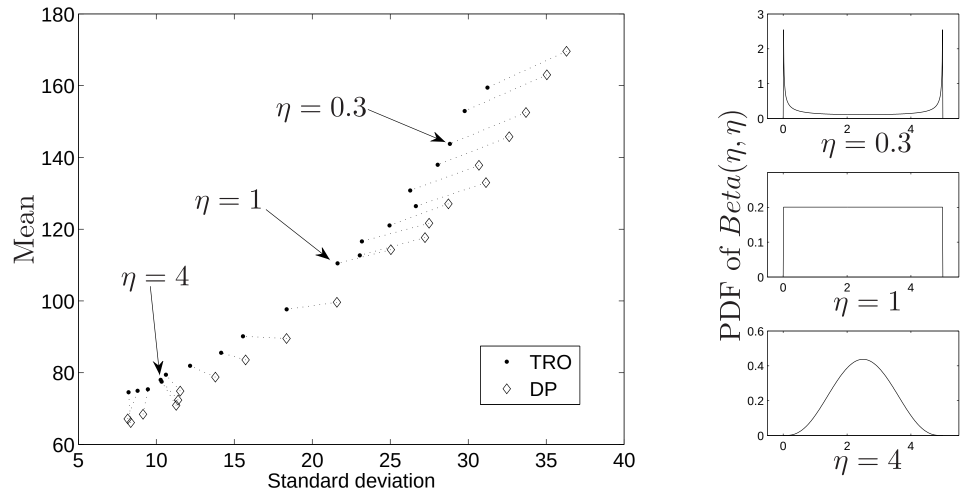

(a) (b)

Figure 5: **Comparison of the TRO and DP policies under a family of demand distributions.**

in the simulations are substantially different from the one assumed by the DP policy (for example, the distributions shown in Figure 5(b) with $\eta = 0.3$ or 1), then the TRO policy is superior to the DP policy in both mean and standard deviation of costs. This suggests that the TRO policy is robust over a broad family of demand distributions.

We also investigate the performance of the TRO and DP policies when the demands for different products follow different distributions. We consider $P = 5$, $T = 1$, $\bar{x}_t = 15$, $A_t^i = 1$, $c_t^i = 1$, $h_t^i = 1+0.2i$, $b_t^i = 2.5i$, for all $i \in \mathcal{P}$, $t \in \mathcal{T}$. To compute the DP policy, we assume the demands follow a $Beta(4, 4)$ distribution. However, in the simulations, the demand for each product i follows a $Beta(\eta_1^i, \eta_2^i)$ distribution, where η_1^i and η_2^i are randomly chosen from the set $\{0.2, 0.6, 0.8, 1, 2, 4, 6\}$. Figure 6(a) shows the mean and the standard deviation of costs under each policy based on 20,000 simulation runs. The results suggest that the TRO policy outperforms the DP policy in both mean and standard deviation of costs for a wide range of α (from 0.476 to 0.938). Figure 6(b) shows the distribution of costs under each policy (with $\alpha = 0.476$ for the TRO policy). The histograms in the top graph suggest that the TRO policy yields a significantly narrower distribution. The bottom graph shows the cumulative distribution of costs under each policy. As the cost increases, the TRO policy outperforms the DP policy by yielding a significantly larger cumulative probability at the cost target $\tau = 42$ ($\alpha = 0.476$, corresponding to the vertical dashed line in the graph). The above results suggest that the TRO policy is robust against distributional ambiguity even if the demands for different products follow different distributions.

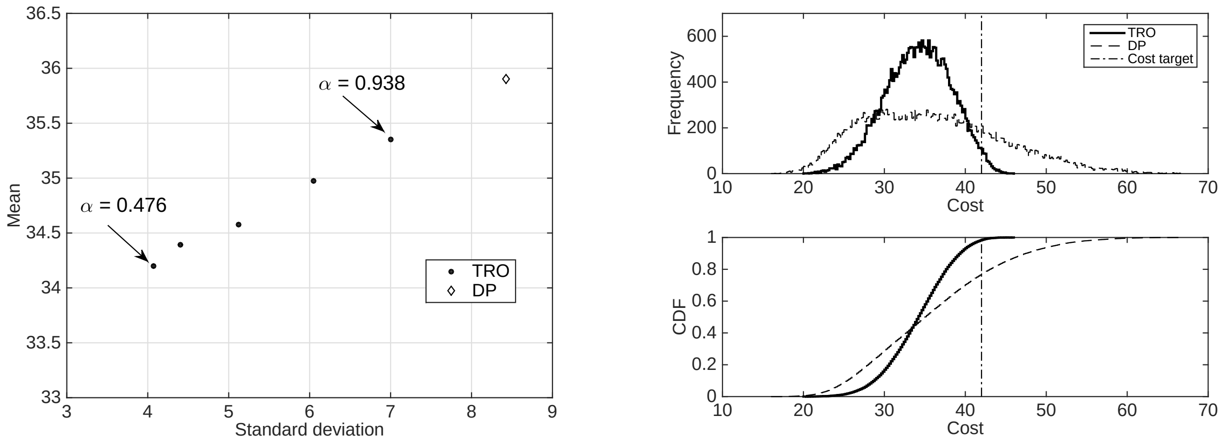

(a) (b)

Figure 6: **Comparison of the TRO and DP policies when demand distributions are random.**

5.5 Quality of the approximation

To evaluate the approximation by Problem (14), we compare its solution with that of the original formulation (Problem (7)). We focus on the case with two products over two periods so that Problem (7) can be solved by enumerating all possible extreme values of demands $\tilde{d}_t^i$, $i, t = 1, 2$.

We consider four different sets of cost parameters. For each set of the parameters, we begin from three different initial inventory settings. This leads to twelve different problem instances. Figure 7 shows the optimal order quantities x_1^i, $i = 1, 2$, for period 1 for both Problems (7) and (14) across various problem instances. The order quantity x_1^i for product i under the approximate formulation (Problem (14)) is very close to that of the original formulation (Problem (7)). Table 3 compares the optimal value of the uncertainty set parameter γ^* under the original formulation with its approximation $\gamma^\dagger$ determined by solving Problem (14). The results suggest that although $\gamma^\dagger$ is generally smaller than γ^*, they are very close for some problem instances.

Table 3: **Optimal uncertainty set parameter γ^* versus its approximation $\gamma^\dagger$**

$\gamma^*/\gamma^\dagger$	**Parameter set 1**	**Parameter set 2**	**Parameter set 3**	**Parameter set 4**
Inventory setting 1	0.5641 / 0.5641	0.5454 / 0.4444	0.4667 / 0.3975	0.6250 / 0.5652
Inventory setting 2	0.3566 / 0.3501	0.3211 / 0.2815	0.3044 / 0.2593	0.5250 / 0.4462
Inventory setting 3	0.4728 / 0.4421	0.3943 / 0.2585	0.2745 / 0.1795	0.5143 / 0.3437

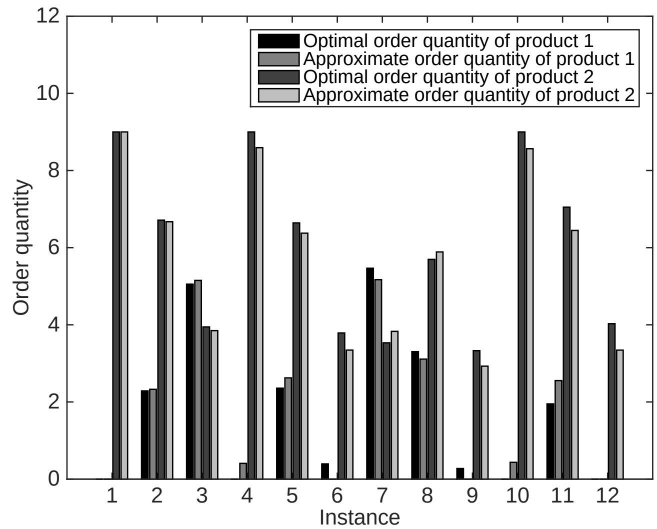

Figure 7: **Optimal order quantities for Problems (7) and (14).**

6 Case studies

To demonstrate the applicability of the TRO policy, we perform case studies with two companies in Singapore. In the first case, we study the inventory management problem of a distributor of petrochemical products. The second case study is conducted with a distributor of hardware parts.

6.1 Case study 1: A distributor of petrochemical products

We evaluate the TRO policy using data from a distributor of petrochemical products in Singapore. Each period in our formulation corresponds to a month for this distributor. All products are from a single supplier. Since the distributor usually receives the products from the supplier within a few days after she places an order, we assume negligible lead times for all the products.

The autocorrelation of the total monthly demand for all 260 products of the distributor does not indicate any obvious seasonality or pattern. We use a moving average method to forecast the mean demand $\hat{d}_t^i$ for product i in period t using its actual demands in periods $t-3$ to $t-1$. Let σ_t^i denote the standard deviation of the demand for product i in period t. Similarly, we use the actual demands for product i in periods $t-3$ to $t-1$ to estimate σ_t^i. We assume the demand for product i in period t falls in the interval $\left[\max\left\{0, \hat{d}_t^i - 3\sigma_t^i\right\}, \hat{d}_t^i + 3\sigma_t^i\right]$. Figure 8(a) shows the forecast mean and bounds of the total demand for all the products.

We use $\alpha = 0.7$ for the TRO policy and set $T = 2$, which gives the best results among other possible

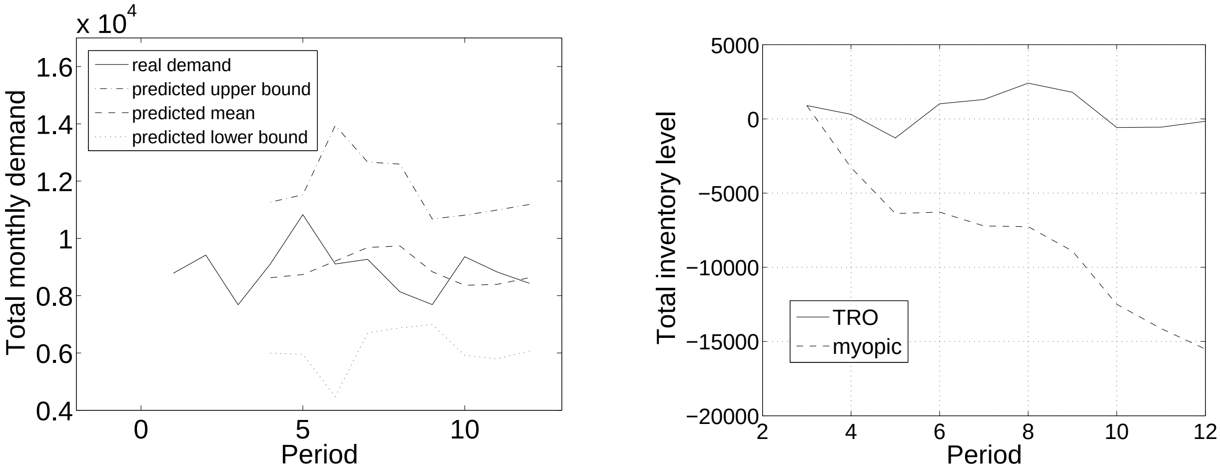

(a) (b)

Figure 8: **Demand forecast and total inventory level.**

values. The values of A_t^i, c_t^i, h_t^i, and b_t^i are known and fall in the ranges $[0.2, 1]$, $[0.06, 0.3]$, $[0.02, 0.1]$, and $[0.1, 0.5]$ respectively. As there are 260 products, the dynamic program in Section 3 becomes intractable. Thus, we compare the TRO policy with the myopic policy. Based on the realized demands in the previous three periods, we use sample average approximation to compute the expected value of the objective function of Problem (18) for the myopic policy. Following a rolling-horizon principle, we compute the resultant cumulative costs of the TRO and the myopic policies using the actual demands from period 4 to period 12.

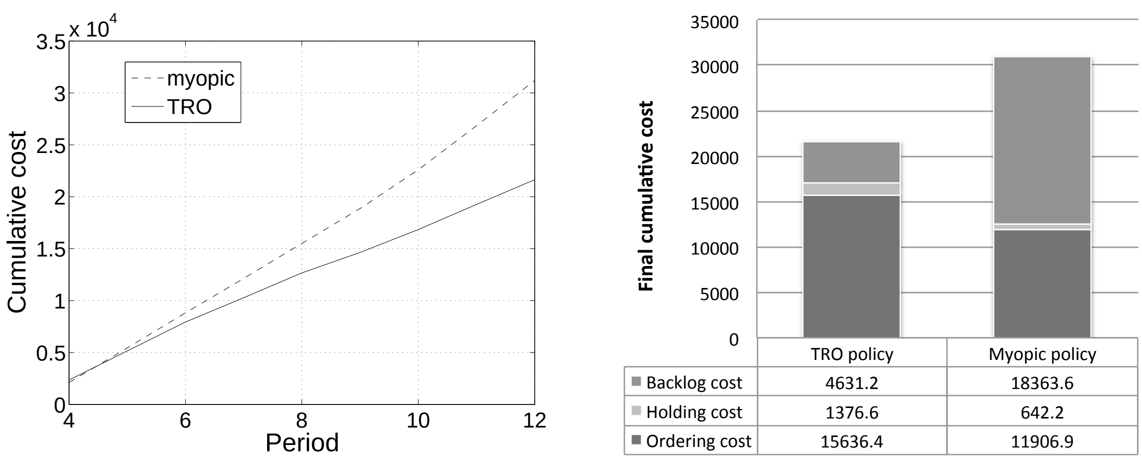

(a) (b)

Figure 9: **Cumulative costs and their components.**

Figure 9(a) shows the cumulative cost of each policy over time. The TRO policy significantly outperforms the myopic policy. The cumulative cost from period 4 to period 12 under the myopic policy is 30,910.4, whereas the cumulative cost under the TRO policy is 21,644.2. Figure 9(b) shows the components of these cumulative costs. The figure suggests that the system has a large backlog cost under the myopic policy. In contrast, the TRO policy achieves a better balance between the backlog and holding costs. This is consistent with Figure 8(b), which shows the total inventory level of all the products under each policy. The TRO policy maintains a total inventory level above or near zero, whereas the total inventory level under the myopic policy falls far below zero after some time.

6.2 Case study 2: A distributor of hardware parts

The second case study is conducted with a distributor of hardware parts in Singapore, which manages a wide range of industrial products such as relay, switch, rotary encoder, programmable logic controller, etc. Demands for these products are driven by both walk-in and project-based customers. Demands of walk-in customers are quite steady in terms of frequency and quantity, whereas demands of project-based customers only occur occasionally but each order requests a large quantity. Under the distributor's current experience-based replenishment policy, it constantly faces overstock and understock.

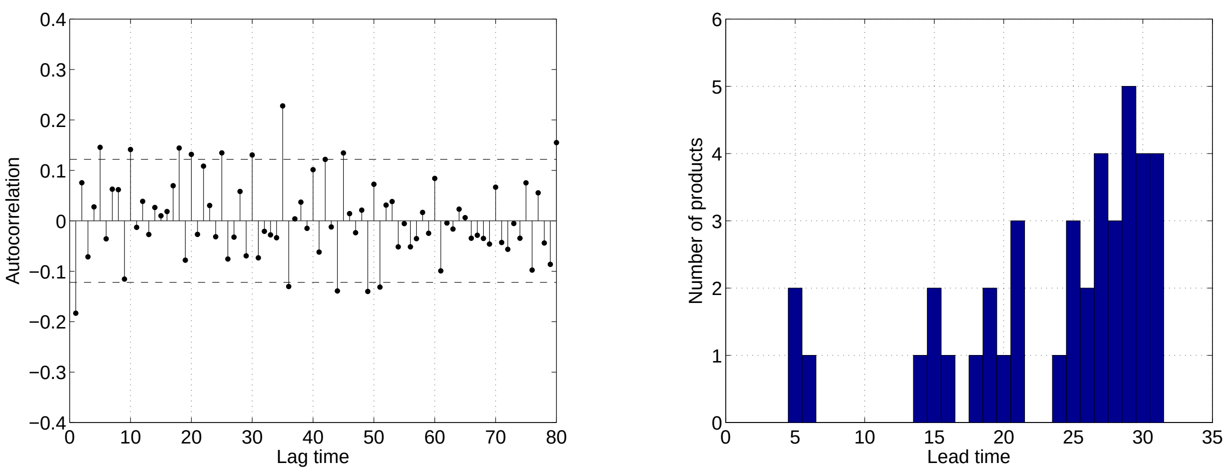

(a) (b)

Figure 10: **Autocorrelation of total daily demand and distribution of lead times.**

Since the demands of project-based customers are handled separately, they are removed from this case study. As a result, we focus on 40 products that are mainly demanded by walk-in customers. Figure 10(a) shows the autocorrelation of the total daily demand of the 40 products over time. Peaks at every 5 time lags suggest a weekly seasonality pattern (due to 5 working days per week). Lead times of the 40 products vary between 5 to 31 days and the distribution is shown in Figure 10(b). Each period corresponds to a day and we choose $T = 33$, which is two days more than the longest lead time. The values of A_t^i, c_t^i, h_t^i, and b_t^i are given and fall in the ranges $[3.2, 40.18]$, $[0.6, 12.05]$, $[0.2, 4.02]$, and $[2.2, 20.09]$ respectively.

Following the observation on Figure 10(a), we use historical data of the latest 9 weeks to predict the products' demands in each day of the following week. For example, the mean demand for product i on Monday $\hat{d}_M^i$ of the following week is the average of actual demands for product i on Mondays in the latest 9 weeks. We assume the demand for product i on Monday of the following week falls in the range $\left[\max\left\{0, \hat{d}_M^i - 3\sigma_M^i\right\}, \hat{d}_M^i + 3\sigma_M^i\right]$, where σ_M^i represents the sample standard deviation of actual demands for product i on Mondays in the latest 9 weeks. Figure 11 compares the actual total demand and the predicted total mean demand on Monday of each week.

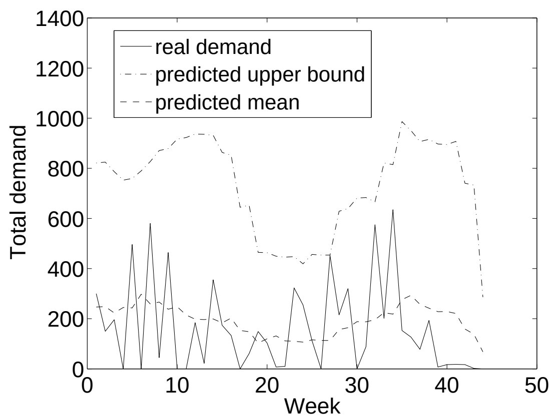

Figure 11: **Demand forecast for Monday of each week.**

To compare the TRO and the myopic policies, we use the first 9 weeks of historical data to estimate the means and the bounds of demands for week 10. Following a rolling-horizon principle, we recompute the policies every period and evaluate them using the actual demands. The cumulative cost of each policy over 100 periods is shown in Figure 12(a). Again, the TRO policy significantly outperforms the myopic policy. Figure 12(b) shows the components of their final cumulative costs. The two policies result in similar ordering and backlog costs. Due to non-zero lead times of the products, both policies maintain a substantial inventory level to avoid backlogs. However, the TRO policy results in a lower inventory holding cost.

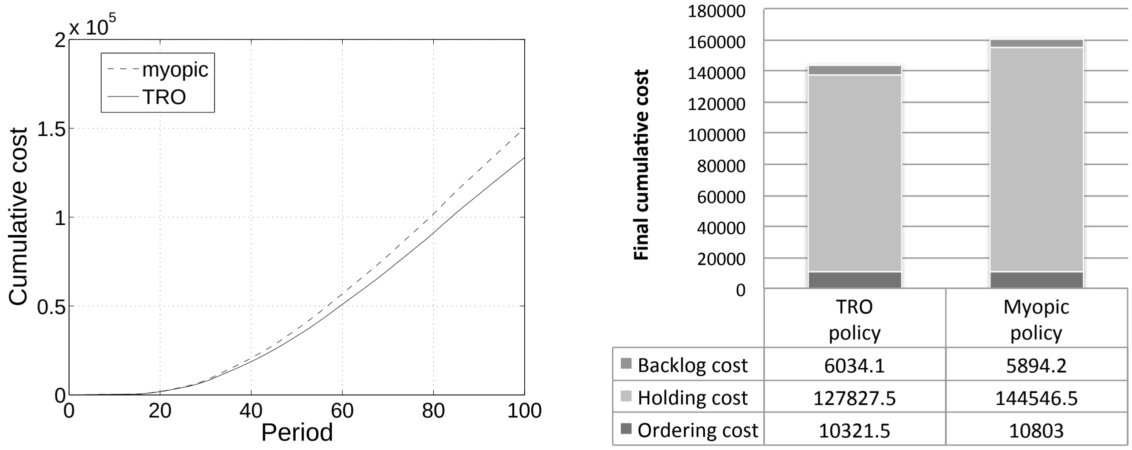

(a) (b)

Figure 12: **Cumulative costs and their components.**

7 Generalizations

Problem (7) can be generalized to handle more complex situations. We discuss two possibilities below.

7.1 Generalizing the cost target

In some situations, a company may have different budget plans for different periods. Furthermore, the demands may vary significantly over time and the resultant cost may change substantially from one period to another. Therefore, a single cost target for the entire planning horizon may not be sensible in this kind of situations. We can generalize our formulation in Problem (7) by assuming that we have a cost target τ_t in each period $t \in \mathcal{T}$. We maximize the sizes of all the adjustable uncertainty sets by maximizing γ, subject to the cost target constraint in each period t as follows:

$$\gamma^* := \max \quad \gamma \tag{19a}$$

$$s.t. \quad \sum_{i \in \mathcal{P}} \left(A_t^i I\left(x_t^i\left(\mathbf{d}_{t-1}\right)\right) + c_t^i x_t^i\left(\mathbf{d}_{t-1}\right) + h_t^i\left(y_{t+1}^i\right)^+ + b_t^i\left(y_{t+1}^i\right)^- \right) \leq \tau_t, \ \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), \quad t \in \mathcal{T}; \tag{19b}$$

$$\sum_{i \in \mathcal{P}} x_t^i\left(\mathbf{d}_{t-1}\right) \leq \bar{x}_t, \quad \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), \qquad t \in \mathcal{T}; \tag{19c}$$

$$x_t^i\left(\mathbf{d}_{t-1}\right) \in \mathcal{N}_t, \quad \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), \qquad i \in \mathcal{P}, \ t \in \mathcal{T}; \tag{19d}$$

$$x_t^i\left(\mathbf{d}_{t-1}\right) \geq 0, \quad \forall \mathbf{d}_{t-1} \in \mathbf{D}_{t-1}(\gamma), \qquad i \in \mathcal{P}, \ t \in \mathcal{T}; \tag{19e}$$

$$0 \leq \gamma \leq 1. \tag{19f}$$

Note that there are T budget constraints in (19b). Instead of having a single budget constraint for the entire planning horizon such as Constraint (7b), we require that the total cost in each period t should not exceed the cost target τ_t in the above formulation. Problem (19) can be solved using a similar approximation procedure described in Section 4.2, and Theorem 1 continues to hold. Thus, Problem (19) can be approximated by Problem (14) with its first constraint replaced by $\sum_{i \in \mathcal{P}} \left(A_t^i I \left(x_t^i \right) + c_t^i x_t^i + \theta_t^i \right) \leq \tau_t$, for all $t \in \mathcal{T}$. As long as the length of planning horizon T is not too large, the computational time to find the ordering policy x_t^i for all $i \in \mathcal{P}, t \in \mathcal{T}$, will not increase significantly.

7.2 Generalizing the uncertainty set parameter

If the demand distributions vary significantly across products and periods, the same value of γ may represent different levels of robustness. For example, if $\tilde{d}_t^i$ follows a normal distribution with low variance, then the adjustable uncertainty set $D_t^i(\gamma)$ captures most demand realizations even with a low value of γ. In contrast, if $\tilde{d}_t^i$ follows a distribution with low probability around the median but high probability near the ends of its support set (such as $Beta(0.3, 0.3)$ in Figure 5(b)), we may need a high value of γ to ensure the robustness.

To make the level of robustness more consistent across different products and periods, we denote an uncertainty set parameter as $\gamma_t^i \in [0, 1]$ and define the corresponding adjustable uncertainty set as

$$D_t^i \left(\gamma_t^i \right) := \left\{ d_t^i \;\middle|\; \hat{d}_t^i - \gamma_t^i \; \underline{z}_t^i \leq d_t^i \leq \hat{d}_t^i + \gamma_t^i \; \bar{z}_t^i \right\}, \tag{20}$$

for all $i \in \mathcal{P}, t \in \mathcal{T}$. We change the objective function of Problem (7) to $\sum_{i \in \mathcal{P}, t \in \mathcal{T}} \epsilon_t^i \gamma_t^i$, where $\epsilon_t^i \in (0, 1)$ represents the weight of γ_t^i and $\sum_{i \in \mathcal{P}, t \in \mathcal{T}} \epsilon_t^i = 1$. We set ϵ_t^i to a large value if $\tilde{d}_t^i$ follows a distribution with high probability around the median (such as a normal distribution). We set ϵ_t^i to a small value if $\tilde{d}_t^i$ follows a distribution with low probability around the median, but high probability near the ends of its support set (such as $Beta(0.3, 0.3)$). Likewise, the resultant problem can be solved using a similar approximation procedure described in Section 4.2, and Theorem 1 remains valid.

8 Conclusion

Traditional methods based on dynamic programming or stochastic programming for inventory management usually require full information of demand distributions. Unfortunately, precise demand distributions are often not available in practice. As a result, the traditional methods may perform poorly in practice when they are given inaccurate demand distributions. Furthermore, these methods often become intractable for a problem with multiple products over multiple periods considering fixed ordering costs and lead times.

To address the above issues, we propose an inventory management approach based on target-oriented robust optimization (TRO). This approach determines an ordering policy that meets a pre-specified cost target, while it accommodates demand variability as much as possible. The proposed approach has three main merits. First, it only makes use of a reference value (such as the mean or median) and the bounds of each product's demand in each period, making it suitable to handle distributional ambiguity of demands. Second, the proposed approach is well scalable and the multi-product, multi-period problem can be solved in a reasonable amount of time. Third, by tuning the cost target, a balance between the expected cost of the problem and the associated risk (cost variance) is achievable.

To reduce computation burden, we focus on a static rule under the TRO approach. This static rule (also called the TRO policy) is computationally more amiable compared to other complicated decision rules. We also prove that the static rule is optimal for an approximation of the original TRO formulation for the problem.

We first consider a special case with a single product over one period to gain some insights. We study the TRO policy in detail and discover that the order quantity is a piece-wise linear function of the inventory level (see Figure 1). If the inventory level is lower than a certain threshold value, the order quantity depends on both the inventory level and the optimal objective of the TRO formulation (γ^*). If the inventory level exceeds the threshold value, the ordering policy becomes an order-up-to policy (see Equations (16)-(17)).

The performance of the TRO policy is then benchmarked against dynamic programming (the DP policy) in a multi-product, single-period problem. Our numerical results suggest that the costs of the TRO policy are very close to that of the DP policy. This is surprising because the TRO policy only makes use of limited demand information (means and bounds). In contrast, the optimal DP policy requires full information of demand distributions. Furthermore, the TRO policy achieves a significantly smaller variance of costs by having only a slightly higher expected cost (see the bottom-left corner of Figure 2). More importantly, if the DP policy assumes inaccurate demand distributions, the TRO policy may significantly outperform the DP policy in both expected value and variance of costs (see the top-right corner of Figure 2). We observe similar results in a multi-product, multi-period problem.

A critical parameter of the TRO policy is the target coefficient, which determines the pre-specified cost target. To find an optimal value of this parameter, we conduct an experiment based on a single-product, multi-period problem. Under various cost parameter settings, our numerical results suggest that the best target coefficient falls in the range $[0.5, 0.8]$.

One advantage of the TRO policy is its robustness against distributional ambiguity. We compare the performance of the TRO and DP policies over a family of beta distributions with the mean and bounds fixed. Our numerical experiments suggest that even if the demand distributions are close to that assumed by the DP policy, the TRO policy generally leads to a smaller variance of costs than the DP policy (see the bottom-left corner of Figure 5(a)). If the demand distributions are very different from that assumed by the DP policy, the TRO policy outperforms the DP policy in both expected cost and variance (see the top-right corner of Figure 5(a)).

We are confident that the TRO policy can handle practical problems with realistic sizes. We have tested the TRO policy in two case studies using data from very different industries. The case studies suggest that the TRO policy can be computed in a reasonable amount of time for monthly and daily ordering decisions, and it generates substantial savings over a myopic policy.

It is worth noting that the performance of the TRO policy can be improved if we know the demand distributions. Suppose $\tilde{d}_t^i$ has a cumulative distribution function $P_t^i(\cdot)$, we can define an adjustable uncertainty set for $\tilde{d}_t^i$ as $D_t^i(\gamma) := \left\{\underline{z}_t^i(\gamma) \le d_t^i \le \bar{z}_t^i(\gamma)\right\}$, where $\underline{z}_t^i(\gamma) := \inf\left\{d_t^i :\ P_t^i\left(d_t^i\right) \ge 0.5 - \gamma/2,\ d_t^i \ge \underline{d}_t^i\right\}$ and $\bar{z}_t^i(\gamma) := \sup\left\{d_t^i :\ P_t^i\left(d_t^i\right) \le 0.5 + \gamma/2,\ d_t^i \le \bar{d}_t^i\right\}$. In this case, Problem (7) can be solved by a binary search over $\gamma \in [0, 1]$.

Acknowledgments

The authors thank the associate editor and the two anonymous referees for their valuable comments, which have substantially improved the paper. This research was supported by the Neptune Orient Lines (NOL) Fellowship Program in Singapore.

References

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

Ang, M., Lim, Y. F. and Sim, M. (2012). Robust storage assignment in unit-load warehouses, *Management Science* **58**(11): 2114–2130.

Azoury, K. S. (1985). Bayes solution to dynamic inventory models under unknown demand distribution, *Mathematical Programming* **31**(9): 1150–1160.

Ben-Tal, A., Golany, B., Nemirovski, A. and Vial, J.-P. (2005). Retailer-supplier flexible commitments contracts: A robust optimization approach, *Manufacturing & Service Operations Management* **7**(3): 248–271.

Ben-Tal, A., Goryashko, A., Guslitzer, E. and Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs, *Mathematical Programming* **99**(2): 351–376.

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

Ben-Tal, A. and Nemirovski, A. (1999). Robust solutions of uncertain linear programs, *Operations Research Letters* **25**(1): 1–13.

Ben-Tal, A. and Nemirovski, A. (2000). Robust solutions of linear programming problems contaminated with uncertain data, *Mathematical Programming* **88**(3): 411–424.

Bertsimas, D., Goyal, V. and Lu, B. Y. (2013). A tight characterization of the performance of static solutions in two-stage adjustable robust linear optimization, *Under review, available at http://www.columbia.edu/~vg2277/* .

Bertsimas, D., Pachamanova, D. and Sim, M. (2004). Robust linear optimization under general norms, *Operations Research Letters* **32**(6): 510–516.

Bertsimas, D. and Sim, M. (2004). The price of robustness, *Operations Research* **52**(1): 35–53.

Bertsimas, D. and Sim, M. (2006). Tractable approximations to robust conic optimization problems, *Mathematical Programming* **107**(1): 5–36.

Bertsimas, D. and Thiele, A. (2004). A robust optimization approach to supply chain management, *Lecture Notes in Computer Science* **3064**: 86–100.

Bertsimas, D. and Thiele, A. (2006). A robust optimization approach to supply chain management, *Operations Research* **54**(1): 150–168.

Bienstock, D. and Özbay, N. (2008). Computing robust basestock levels, *Discrete Optimization* **5**(2): 389–414.

Birge, J. R. and Louveaux, F. (2011). *Introduction to Stochastic Programming*, 2nd edn, New York: Springer.

Chen, W. and Sim, M. (2009). Goal-driven optimization, *Operations Research* **57**(2): 342–357.

Chen, W., Sim, M., Sun, J. and Teo, C.-P. (2009). From cvar to uncertainty set: Implications in joint chance-constrained optimization, *Operations Research* **58**(2): 470–485.

Chen, X., Sim, M., Sun, P. and Zhang, J. (2008). A linear-decision based approximation approach to stochastic programming, *Operations Research* **56**(2): 344357.

El-Ghaoui, L. and Lebret, H. (1997). Robust solutions to least-squares problems with uncertain data, *SIAM Journal on Matrix Analysis and Applications* **18**(4): 1035–1064.

El-Ghaoui, L., Oustry, F. and Lebret, H. (1998). Robust solutions to uncertain semidefinite programs, *SIAM Journal on Optimization* **9**(1): 33–52.

Goh, J. and Sim, M. (2010). Distributionally robust optimization and its tractable approximations, *Operations Research* **58**(4): 902–917.

Miller, B. L. (1986). Scarf's state reduction method, flexibility, and a dependent demand inventory model, *Operations Research* **34**(1): 83–90.

Nahmias, S. (2009). *Production and Operations Analysis*, 1st edn, McGraw-Hill, New York.

Scarf, H. (1958). A min-max solution of an inventory problem, *Studies in the mathematical theory of inventory and production*, K.J. Arrow and S. Karlin and H. Scarf, eds., Stanford University Press, Stanford, CA, pp. 201–209.

Scarf, H. (1959). Bayes solutions of the statistical inventory problem, *Annals of Mathematical Statistics* **30**(2): 490–508.

Scarf, H. (1960). Some remarks on bayes solution to the inventory problem, *Naval Research Logistics* **7**(4): 591596.

See, C.-T. and Sim, M. (2010). Robust approximation to multiperiod inventory management, *Operations Research* **58**(3): 583–594.

Shapiro, A. (2006). On complexity of multistage stochastic programs, *Operations Research Letters* **34**(1): 1–8.

Song, M., Klabjan, D. and Simchi-Levi, D. (2012). Robust stochastic lot-sizing by means of histograms, *Production and Operations Management* .

Veinott, A. F. (1966). The status of mathematical inventory theory, *Management Science* **12**(11): 745–777.

Wang, C., Ong, C.-J. and Sim, M. (2009). Convergence properties of constrained linear system under MPC control law using affine disturbance feedback, *Automatica* **45**(7): 1715–1720.

Wang, C., Ong, C.-J. and Sim, M. (2010). Model predictive control using segregated disturbance feedback, *IEEE Transactions on Automatic Control* **55**(4): 831–840.

Zipkin, P. H. (2000). *Foundations of Inventory Management*, McGraw-Hill, New York.

A Proof of Theorem 2

Proof. Since $P = 1$ and $T = 1$, Problem (8) is equivalent to Problem (7) and both reduce to

$$\gamma^* = \gamma'' = \max \quad \gamma \tag{21a}$$

$$s.t. \quad cx + \theta \leq \tau; \tag{21b}$$

$$\theta \geq h\left(y - \hat{d} + (x + \gamma \underline{z})\right); \tag{21c}$$

$$\theta \geq -b\left(y - \hat{d} + (x - \gamma \bar{z})\right); \tag{21d}$$

$$x \geq 0; \tag{21e}$$

$$0 \leq \gamma \leq 1; \tag{21f}$$

which is equivalent to Problem (14). Therefore, we have $\gamma^\dagger = \gamma^*$ in this special case. There are two cases for the solution to Problem (21): (a) $\gamma^* = 1$ and (b) $0 \leq \gamma^* < 1$.

(a) $\gamma^* = 1$. As described at the end of Section 4.1, the optimal solution x in this case can be determined by solving the following problem:

$$\min \quad cx + \theta \tag{22a}$$

$$s.t. \quad \theta \geq h\left(y - \hat{d} + (x + \underline{z})\right); \tag{22b}$$

$$\theta \geq -b\left(y - \hat{d} + (x - \bar{z})\right); \tag{22c}$$

$$x \geq 0. \tag{22d}$$

Recall that $c - b < 0$, otherwise no orders should be placed all the time. Define $g_1(x) := (h + c)x + h\left(y - \hat{d} + \underline{z}\right)$, $g_2(x) := (c - b)x - b\left(y - \hat{d} - \bar{z}\right)$, and $g(x) := \max\{g_1(x), g_2(x)\}$. Problem (22) is equivalent to $\min\{g(x) | x \geq 0\}$. According to the initial inventory level y, there are two possible cases for the solution of this problem:

(a.1) If the initial inventory level y is sufficiently low such that $g_1(0) \leq g_2(0)$, then the functions $g_1(x)$ and $g_2(x)$ intersect with each other at some $x \geq 0$ as shown in Figure 13(a). In this case, the minimum of $g(x)$ occurs at x_{opt} where $g_1(x_{opt}) = g_2(x_{opt})$, which implies

$$x_{opt} = \hat{d} + \frac{b\bar{z} - h\underline{z}}{h + b} - y = Y(1) - y. \tag{23}$$

Recall that Equation (23) holds if $g(x_{opt}) \leq \tau$ and $g_1(0) \leq g_2(0)$, which imply

$$y \geq y_a \text{ and } y \leq Y(1), \tag{24}$$

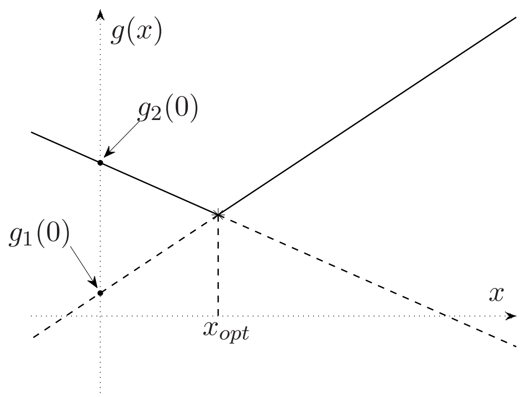

(a) Case a.1

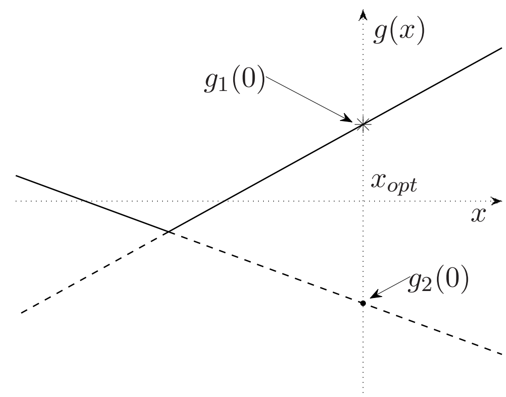

(b) Case a.2

Figure 13: **Two cases for $\gamma^* = 1$.**

respectively.

(a.2) If the initial inventory level y is sufficiently high such that $g_1(0) > g_2(0)$, then the functions $g_1(x)$ and $g_2(x)$ intersect with each other at some $x < 0$ as shown in Figure 13(b). In this case,

$$x_{opt} = 0, \tag{25}$$

and $\min\{g(x)|x \geq 0\} = g_1(0)$. The order quantity $x_{opt} = 0$ is optimal if $g_1(0) \leq \tau$ and $g_1(0) \geq g_2(0)$, which imply

$$y \geq Y(1) \text{ and } y \leq y_b. \tag{26}$$

(b) $0 \leq \gamma^* < 1$. In this case, Constraint (21b) is always binding because γ^* can no longer be increased. Thus, given $\gamma^* < 1$ the optimal solution x to Problem (21) can also be obtained by solving the following problem:

$$\min \quad cx + \theta \tag{27a}$$

$$s.t. \quad \theta \geq h\left(y - \hat{d} + (x + \gamma^* \underline{z})\right); \tag{27b}$$

$$\theta \geq -b\left(y - \hat{d} + (x - \gamma^* \bar{z})\right); \tag{27c}$$

$$x \geq 0. \tag{27d}$$

Note that Problem (27) has the same structure as Problem (22) and its solution can be derived in a similar procedure:

(b.1) If y is sufficiently small such that $h\left(y-\hat{d}+\gamma^*\underline{z}\right) \le -b\left(y-\hat{d}-\gamma^*\bar{z}\right)$, which is equivalent to $y \le Y(\gamma^*)$, the optimal order quantity is

$$x_{opt} = \hat{d} + \gamma^* \frac{b\bar{z} - h\underline{z}}{h+b} - y = Y(\gamma^*) - y. \tag{28}$$

Since Constraint (21b) is binding, we have $(h+c)x_{opt} + h(y - \hat{d} + \gamma^*\underline{z}) = \tau$. By substituting x_{opt} into this equation, we have

$$\gamma^* = \left(\tau + cy - c\hat{d}\right) / \left(\frac{c+h}{b+h} b\bar{z} + \frac{b-c}{b+h} h\underline{z}\right). \tag{29}$$

Together with $0 \le \gamma^* < 1$, Equation (29) yields

$$\underline{y} \le y < y_a. \tag{30}$$

(b.2) If y is sufficiently large such that $h\left(y-\hat{d}+\gamma^*\underline{z}\right) > -b\left(y-\hat{d}-\gamma^*\bar{z}\right)$, which is equivalent to $y > Y(\gamma^*)$, we have $x_{opt} = 0$. Furthermore, $\tau = h\left(y - \hat{d} + \gamma^*\underline{z}\right)$, which implies

$$\gamma^* = \frac{1}{h\underline{z}} \left(\tau - hy + h\hat{d}\right). \tag{31}$$

Together with $0 \le \gamma^* < 1$, Equation (31) yields

$$y_b < y \le \bar{y}. \tag{32}$$

Combining Cases (a) and (b), the optimal objective value and replenishment policy for the single-product, single-period problem can be summarized as Equations (16) and (17) respectively. □