

A Robust Optimization Approach to Inventory Theory

Dimitris Bertsimas

Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, E53-363,
Cambridge, Massachusetts 02139, dbertsim@mit.edu

Aurélie Thiele

Department of Industrial and Systems Engineering, Lehigh University, Mohler Building,
Bethlehem, Pennsylvania 18015, aurelie.thiele@lehigh.edu

We propose a general methodology based on robust optimization to address the problem of optimally controlling a supply chain subject to stochastic demand in discrete time. This problem has been studied in the past using dynamic programming, which suffers from dimensionality problems and assumes full knowledge of the demand distribution. The proposed approach takes into account the uncertainty of the demand in the supply chain without assuming a specific distribution, while remaining highly tractable and providing insight into the corresponding optimal policy. It also allows adjustment of the level of robustness of the solution to trade off performance and protection against uncertainty. An attractive feature of the proposed approach is its numerical tractability, especially when compared to multidimensional dynamic programming problems in complex supply chains, as the robust problem is of the same difficulty as the nominal problem, that is, a linear programming problem when there are no fixed costs, and a mixed-integer programming problem when fixed costs are present. Furthermore, we show that the optimal policy obtained in the robust approach is identical to the optimal policy obtained in the nominal case for a modified and explicitly computable demand sequence. In this way, we show that the structure of the optimal robust policy is of the same base-stock character as the optimal stochastic policy for a wide range of inventory problems in single installations, series systems, and general supply chains. Preliminary computational results are very promising.

Subject classifications: programming: linear, integer, stochastic; inventory, production: uncertainty.

Area of review: Financial Engineering.

History: Received April 2003; revisions received December 2003, May 2004, July 2004; accepted November 2004.

1. Introduction

Optimal supply chain management has been extensively studied in the past with much theoretical success. Dynamic programming has long emerged as the standard tool for this purpose, and has led to significant breakthroughs as early as 1960, when Clark and Scarf (1960) proved the optimality of base-stock policies for series systems in their landmark paper. Although dynamic programming is a powerful technique as to the theoretical characterization of the optimal policy for simple systems, the complexity of the underlying recursive equations over a growing number of state variables makes it ill suited for the computation of the actual policy parameters, which is crucial for real-life applications. Approximation algorithms have been developed to address those issues. These include stochastic approximation (see Kushner and Clark 1978) and infinitesimal perturbation analysis (IPA) (see Glasserman 1991, Ho and Cao 1991), where a class of policies, e.g., base-stock, characterized by a set of parameters, is optimized using simulation-based methods (see Fu 1994, Glasserman and Tayur 1995, Kapuscinski and Tayur 1999). IPA-based methods assume knowledge of the underlying probability distributions and restrict their attention to cer-

tain classes of policies that might be suboptimal for a general network problem. Another technique that has gained popularity in recent years is approximate dynamic programming, described by Bertsekas and Tsitsiklis (1996). Despite their promising potential, these methods remain hard to implement in practice. As a result, for implementation purposes, preference is given to more intuitive policies that are much easier to compute, but also suboptimal.

Zipkin (2000) describes policies widely used in practice, such as the economic-order-quantity model, where the demand is constant over time, and the dynamic-economic-lotsize model, which incorporates time-varying demands, for single installations. In both cases, the demand is considered to be without any uncertainty. Myopic policies, which minimize the cost solely at the current time period, are also often used as a substitute for the optimal policy obtained by dynamic programming. For supply chains more complex than series systems, what is commonly referred to as “the curse of dimensionality” plagues even the theoretical use of dynamic programming to find the structure of the optimal policy, thus making it necessary to resort to approximations.

Dynamic programming also assumes full knowledge of the underlying distributions, which further limits its prac-

tical usefulness. The first attempt to address the issue of imperfect information in inventory control is due to Scarf (1958), who studied the optimal policy for the most adverse distribution in a one-period, one-stage inventory model where only the mean and the variance of the demand are known. Moon and Gallego later extended this approach to single-period newsboy problems (1993) and to one-stage inventory models with a fixed reorder quantity and under periodic review (1994). Similar ideas were applied to finite-horizon inventory models by Gallego et al. (2001), under the assumption that demand is a discrete random variable taking values in a known countable set. However, their approach relies on dynamic programming and, as a result, suffers from similar practical limitations.

Hence, the need arises to develop a new optimization approach that incorporates the stochastic character of the demand in the supply chain without making any assumptions on its distribution, is applicable to a wide range of network topologies, is easy to understand intuitively, and combines computational tractability with the structural properties of the optimal policy. The goal of this paper is to present such an approach, based on robust linear and mixed-integer optimization.

Robust optimization addresses the problem of data uncertainty by guaranteeing the feasibility and optimality of the solution for the worst instances of the parameters. However, because it is intrinsically a worst-case approach, feasibility often comes at the cost of performance and generally leads to overconservative solutions. For instance, Soyster (1973) proposes a model to handle columnwise uncertainty in linear programming problems, where every uncertain parameter has to be taken equal to its worst-case value in the set. Subsequent research efforts, led by Ben-Tal and Nemirovski (1998, 1999, 2000), El Ghaoui et al. (1998), and El Ghaoui and Lebret (1997) to address overconservativeness, have applied robust optimization to linear programming problems with ellipsoidal uncertainty sets, thus obtaining conic quadratic programs. Bertsimas and Sim (2004) present a technique specifically tailored for polyhedral uncertainty that leads to linear robust counterparts while controlling the level of conservativeness of the solution. Moreover, their method readily extends to discrete optimization problems (Bertsimas and Sim 2003). This is the technique we will use here to develop a tractable approach to optimal supply chain control in the presence of demand uncertainty.

Specifically, the contributions of this paper are as follows:

1. We develop an approach that incorporates demand randomness in a deterministic manner, remains numerically tractable as the dimension of the problem increases, and leads to high-quality solutions without assuming a specific demand distribution. In particular, preliminary computational results are quite promising.

2. The robust problem is of the same class as the nominal problem, that is, a linear programming problem if there

are no fixed costs or a mixed-integer programming problem if fixed costs are present—independently of the topology of the network. Moreover, the optimal robust policy is identical to the optimal nominal policy for a modified demand sequence.

3. The optimal robust policy is qualitatively similar to the optimal policy obtained by dynamic programming when known. In particular, it remains base stock when the optimal stochastic policy is base stock, as well as in some other cases where the optimal stochastic policy is not known.

4. We derive closed-form expressions of key parameters defining the optimal policy. These expressions provide a deeper insight into the way uncertainty affects the optimal policy in supply chain problems.

Structure of the Paper

In §2, we present the general framework of robust linear optimization. In §3, we apply this technique to the inventory problem at a single station, and extend it to the network case in §4. We implement the method in §5. Finally, §6 contains some concluding remarks.

2. The Robust Optimization Approach

2.1. Uncertainty Structure

We rely extensively on the robust optimization tools developed by Bertsimas and Sim (2004) for linear programming problems. We consider the following problem subject to data uncertainty:

$$\begin{aligned} & \text{minimize} && \mathbf{c}'\mathbf{x} \\ & \text{subject to} && \mathbf{A}\mathbf{x} \leq \mathbf{b}, \\ & && \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}. \end{aligned}$$

We assume without any loss of generality that the data uncertainty only affects the elements in matrix \mathbf{A} . Indeed, if there is uncertainty on \mathbf{b} or \mathbf{c} , we can rewrite the linear programming problem as

$$\begin{aligned} & \text{minimize} && \tilde{\mathbf{c}}'\tilde{\mathbf{x}} \\ & \text{subject to} && \tilde{\mathbf{A}}\tilde{\mathbf{x}} \geq \mathbf{0}, \\ & && \tilde{\mathbf{l}} \leq \tilde{\mathbf{x}} \leq \tilde{\mathbf{u}}, \end{aligned}$$

with $\tilde{\mathbf{x}} = (z, \mathbf{x}, y)'$, $\tilde{\mathbf{c}} = (1, \mathbf{0}, 0)'$, $\tilde{\mathbf{l}} = (-M, \mathbf{l}, 1)'$, and $\tilde{\mathbf{u}} = (M, \mathbf{u}, 1)'$, where M is a large constant, and finally,

$$\tilde{\mathbf{A}} = \begin{pmatrix} 1 & -\mathbf{c}' & 0 \\ \mathbf{0} & -\mathbf{A} & \mathbf{b} \end{pmatrix}.$$

We model data uncertainty in \mathbf{A} as follows. Each uncertain coefficient a_{ij} is known to belong to an interval centered at its nominal value \bar{a}_{ij} and of half-length \hat{a}_{ij} , but its exact value is unknown. As much as it is unlikely

that all coefficients are equal to their nominal value, it is also unlikely that they are all equal to their worst-case value. For this reason, the “safest” approach, where all parameters are taken equal to their worst bound, leads to a severe deterioration of the cost without necessarily being justified in practice. Hence, we wish to adjust the level of conservativeness of the solution, so that a reasonable trade-off between robustness and performance is achieved. We now quantify the concept above in mathematical terms.

We define the scaled deviation of parameter a_{ij} from its nominal value as $z_{ij} = (a_{ij} - \bar{a}_{ij})/\hat{a}_{ij}$. The scaled deviation takes values in $[-1, 1]$. Moreover, we impose a budget of uncertainty in the following sense: The total (scaled) variation of the parameters cannot exceed some threshold Γ , not necessarily integer:

$$\sum_{(i,j) \in J} |z_{ij}| \leq \Gamma,$$

where J is the set of indices of the uncertain parameters. By taking $\Gamma = 0$ (respectively, $\Gamma = |J|$), we obtain the nominal (respectively, worst) case. Bertsimas and Sim (2004) show that having the threshold Γ vary in $(0, |J|)$ allows greater flexibility to build a robust model without excessively affecting the optimal cost. Intuitively, the budget of uncertainty rules out large deviations in $\sum_j a_{ij}x_j$, which play a predominant role in worst-case analysis, but actually occur with low probability, because large deviations in a_{ij} tend to cancel each other out as the number of parameters increases.

2.2. The Robust Approach

Let

$$\mathcal{A} = \left\{ \mathbf{A} \in \mathcal{R}^{m \times n} \mid a_{ij} \in [\bar{a}_{ij} - \hat{a}_{ij}, \bar{a}_{ij} + \hat{a}_{ij}] \ \forall i, j, \right. \\ \left. \sum_{(i,j) \in J} \frac{|a_{ij} - \bar{a}_{ij}|}{\hat{a}_{ij}} \leq \Gamma \right\}.$$

The robust problem is then formulated as

$$\begin{aligned} &\text{minimize} \quad \mathbf{c}'\mathbf{x} \\ &\text{subject to} \quad \mathbf{Ax} \leq \mathbf{b} \quad \forall \mathbf{A} \in \mathcal{A}, \\ &\quad \quad \quad \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}. \end{aligned} \tag{1}$$

THEOREM 2.1 (BERTSIMAS AND SIM 2003A). *The uncertain linear programming problem has the following robust, linear counterpart:*

$$\begin{aligned} &\text{minimize} \quad \mathbf{c}'\mathbf{x} \\ &\text{subject to} \quad \sum_j \bar{a}_{ij}x_j + q_i\Gamma + \sum_{j:(i,j) \in J} r_{ij} \leq b_i \quad \forall i, \\ &\quad \quad \quad q_i + r_{ij} \geq \hat{a}_{ij}y_j \quad \forall (i, j) \in J, \\ &\quad \quad \quad -\mathbf{y} \leq \mathbf{x} \leq \mathbf{y}, \quad \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \\ &\quad \quad \quad \mathbf{q} \geq 0, \quad \mathbf{r} \geq 0, \quad \mathbf{y} \geq 0. \end{aligned} \tag{2}$$

PROOF. The proof is based on strong duality. We outline it here because we utilize similar ideas in other parts of the paper. The i th constraint is $\max_{\mathbf{A} \in \mathcal{A}} \sum_j a_{ij}x_j \leq b_i$. As a result, we have to solve the following auxiliary problem associated with row i :

$$\begin{aligned} &\text{maximize} \quad \sum_j (\bar{a}_{ij} + z_{ij}\hat{a}_{ij})x_j \\ &\text{subject to} \quad \sum_{(l,j) \in J} |z_{lj}| \leq \Gamma, \\ &\quad \quad \quad 0 \leq |z_{lj}| \leq 1 \quad \forall (l, j) \in J. \end{aligned} \tag{3}$$

At optimality, only z_{lj} with $l = i$ will be nonzero. This maximization problem can then be reformulated as a minimization problem using strong duality because the feasible set is nonempty ($\mathbf{z} = \mathbf{0}$ is solution) and bounded. Its dual is

$$\begin{aligned} &\text{minimize} \quad q_i\Gamma + \sum_{j:(i,j) \in J} r_{ij} \\ &\text{subject to} \quad q_i + r_{ij} \geq \hat{a}_{ij}|x_j| \quad \forall j: (i, j) \in J, \\ &\quad \quad \quad q_i \geq 0, \quad r_{ij} \geq 0 \quad \forall j: (i, j) \in J. \end{aligned} \tag{4}$$

Reinjecting into the original problem, Bertsimas and Sim obtain the robust counterpart. \square

The robust counterpart is therefore of the same class as the nominal problem, that is, a linear programming problem. This is a highly attractive feature of this approach because linear programming problems are readily solved by standard optimization packages. Moreover, if in the original problem (1), some of the variables were constrained to be integers, then the robust counterpart (2) would retain the same properties, i.e., the robust counterpart of a mixed-integer programming problem is itself another mixed-integer programming problem.

3. The Single-Station Case

3.1. The Uncapacitated Model

In this section, we apply the robust optimization framework to the problem of ordering, at a single installation, a single type of item subject to stochastic demand over a finite discrete horizon of T periods, so as to minimize a given cost function. We closely follow Bertsekas (1995) in our setting. We define, for $k = 0, \dots, T$,

- x_k : stock available at the beginning of the k th period,
- u_k : stock ordered at the beginning of the k th period,
- w_k : demand during the k th period.

The stock ordered at the beginning of the k th period is delivered before the beginning of the $(k + 1)$ st period, that is, all orders have a constant lead time equal to 0. Excess demand is backlogged. Therefore, the evolution of the stock over time is described by the following linear

equation:

$$x_{k+1} = x_k + u_k - w_k, \quad k = 0, \dots, T-1, \quad (5)$$

leading to the closed-form expression

$$x_{k+1} = x_0 + \sum_{i=0}^k (u_i - w_i), \quad k = 0, \dots, T-1. \quad (6)$$

Neither the stock available nor the quantity ordered at each period is subject to upper bounds. Section 3.2 deals with the capacitated case.

The demands w_k are random variables. To apply the approach outlined in §2, we model w_k for each k as an uncertain parameter that takes values in $[\bar{w}_k - \hat{w}_k, \bar{w}_k + \hat{w}_k]$. We define the scaled deviation of w_k from its nominal value to be $z_k = (w_k - \bar{w}_k)/\hat{w}_k$, which takes values in $[-1, 1]$. We impose budgets of uncertainty at each time period k for the scaled deviations up to time k . Hence, we now have the constraint $\sum_{i=0}^k |z_i| \leq \Gamma_k$ for all time periods $k = 0, \dots, T-1$. These budgets of uncertainty rule out large deviations in the cumulative demand, and as a result the robust methodology can be understood as a “reasonable worst-case” approach. The main assumption we make on Γ_k is that they are increasing in k , i.e., we feel that uncertainty increases with the number of time periods considered. We also constrain Γ_k to be increasing by at most 1 at each time period, i.e., the increase of the budgets of uncertainty should not exceed the number of new parameters added at each time period.

Finally, we specify the cost function. The cost incurred at period k consists of two parts: a purchasing cost, $C(u_k)$, and a holding/shortage cost resulting from this order, $R(x_k + u_k - w_k)$, which is computed at the end of the period, after the shipment u_k has been received and the demand w_k has been realized. Here, we consider a purchasing cost of the form

$$C(u) = \begin{cases} K + c \cdot u & \text{if } u > 0, \\ 0 & \text{if } u = 0, \end{cases} \quad (7)$$

with $c > 0$ the unit variable cost and $K \geq 0$ the fixed cost. If $K > 0$, a fixed positive cost is incurred whenever an order is made. The holding/shortage cost represents the cost associated with having either excess inventory (positive stock) or unfilled demand (negative stock). We consider a convex, piecewise linear holding/shortage cost

$$R(x) = \max(hx, -px), \quad (8)$$

where h and p are nonnegative. We assume that $p > c$, so that ordering stock remains a possibility up to the last period.

Using the piecewise linearity and convexity of the holding/shortage cost function, and modelling the fixed ordering cost with binary variables, the inventory problem we

consider can be written as a mixed-integer programming problem:

$$\text{minimize } \sum_{k=0}^{T-1} (cu_k + Kv_k + y_k) \quad (9)$$

subject to

$$y_k \geq h \left(x_0 + \sum_{i=0}^k (u_i - w_i) \right), \quad k = 0, \dots, T-1, \quad (10)$$

$$y_k \geq -p \left(x_0 + \sum_{i=0}^k (u_i - w_i) \right), \quad k = 0, \dots, T-1, \quad (11)$$

$$0 \leq u_k \leq Mv_k, \quad v_k \in \{0, 1\}, \quad k = 0, \dots, T-1, \quad (12)$$

where $w_i = \bar{w}_i + \hat{w}_i \cdot z_i$ such that $\mathbf{z} \in \mathcal{P} = \{|z_i| \leq 1 \forall i \geq 0, \sum_{i=0}^k |z_i| \leq \Gamma_k \forall k \geq 0\}$.

Data uncertainty now only affects the first two constraints of the mixed-integer programming problem. We isolate the effect of the uncertainty on the stock level by writing x_{k+1} under the form $x_{k+1} = \bar{x}_{k+1} + \sum_{i=0}^k \hat{w}_i z_i$, where \bar{x}_{k+1} is the inventory we would have by ordering the same quantities if there was no uncertainty on the demand, that is, $\bar{x}_{k+1} = x_0 + \sum_{i=0}^k (u_i - \bar{w}_i)$ for all k . Following the technique developed in §2, the robust approach consists here of maximizing the right-hand side of the constraints over the set of admissible scaled deviations. For the k th pair of holding/shortage constraints, this amounts to solving the auxiliary linear programming problem:

$$\begin{aligned} & \text{maximize } \sum_{i=0}^k \hat{w}_i z_i \\ & \text{subject to } \sum_{i=0}^k z_i \leq \Gamma_k, \\ & \quad 0 \leq z_i \leq 1 \quad \forall i. \end{aligned} \quad (13)$$

REMARKS. 1. This auxiliary problem arises from minimizing $\sum_{i=0}^k \hat{w}_i z_i$ in the holding constraint (10) and maximizing $\sum_{i=0}^k \hat{w}_i z_i$ in the shortage constraint (11) over \mathcal{P} , which is symmetric in z_i . As a result, the optimal z_i in $[-1, 1]$, $i = 0, \dots, k$, obtained in both cases are the *opposite* of each other and there is no feasible demand in the uncertain set that realizes both constraints of each pair.

2. This also illustrates why we allot thresholds to $\sum_{i=0}^k z_i$ for all time periods $k = 0, \dots, T-1$. If we only had a global threshold on $\sum_{i=0}^{T-1} z_i$, we would maximize $\sum_{i=0}^k \hat{w}_i z_i$ over $0 \leq z_i \leq 1$ for all i and $\sum_{i=0}^{T-1} z_i \leq \Gamma$. As a result, for all k such that $k \leq \lfloor \Gamma \rfloor$, we would have at optimality $z_i = 1$ for $i \leq k$, which is equivalent to taking all the w_i in the early time periods equal to their worst-case value. The early time periods would then be overprotected.

3. The robust methodology does not depend on \bar{x}_k , $k \geq 0$. Therefore, it can readily be extended to arbitrary constant lead times L , with $\bar{x}_{k+1} = x_0 + \sum_{i=0}^{k-L} u_i - \sum_{i=0}^k \bar{w}_i$ for all k .

Because the linear programming problem (13) is feasible and bounded, by strong duality the optimal cost of this problem is equal to the optimal cost of its dual. Reinjecting the dual of the auxiliary problem in Equations (10) and (11), we obtain the following robust formulation for the single-station inventory problem:

$$\begin{aligned}
 & \text{minimize} && \sum_{k=0}^{T-1} (cu_k + Kv_k + y_k) \\
 & \text{subject to} && \\
 & y_k \geq h \left(x_0 + \sum_{i=0}^k (u_i - \bar{w}_i) + q_k \Gamma_k + \sum_{i=0}^k r_{ik} \right) \quad \forall k, \\
 & y_k \geq p \left(-x_0 - \sum_{i=0}^k (u_i - \bar{w}_i) + q_k \Gamma_k + \sum_{i=0}^k r_{ik} \right) \quad \forall k, \\
 & q_k + r_{ik} \geq \hat{w}_i \quad \forall k, \forall i \leq k, \\
 & q_k \geq 0, \quad r_{ik} \geq 0 \quad \forall k, \forall i \leq k, \\
 & 0 \leq u_k \leq Mv_k, \quad v_k \in \{0, 1\} \quad \forall k,
 \end{aligned} \tag{14}$$

where M is a large positive number.

The variables q_k and r_{ik} quantify the sensitivity of the cost to infinitesimal changes in the key parameters of the robust approach, namely, the level of conservativeness and the bounds of the uncertain variables. At each time period k , $q_k \Gamma_k + \sum_{i=0}^k r_{ik}$ represents the worst-case deviation of the cumulative demand from its nominal value, subject to the budgets of uncertainty.

The robust problem is a linear programming problem if there is no fixed cost ($K = 0$) and a mixed-integer programming problem if fixed costs are present ($K > 0$). In both cases, this robust model can be readily solved numerically through standard optimization tools, which is of course very appealing. It is also desirable to have some theoretical understanding of the optimal policy, in particular with respect to the optimal nominal policy and, if known, the optimal stochastic policy. We address these questions next.

DEFINITION 3.1 ((S, S) AND (s, S) POLICIES). The optimal policy of a discrete-horizon inventory problem is said to be (s, S), or base stock, if there exists a threshold sequence (s_k, S_k) such that, at each time period k , it is optimal to order $S_k - x_k$ if $x_k < s_k$ and 0 otherwise, with $s_k \leq S_k$. If there is no fixed ordering cost ($K = 0$), $s_k = S_k$.

To analyze the optimal robust policy, we need the following lemma.

LEMMA 3.1 (OPTIMAL NOMINAL AND STOCHASTIC POLICY). (a) *The optimal policy in the stochastic case, where the cost to minimize is the expected value of the cost function over the random variables w_k is (s, S). As a result, the optimal policy for the nominal problem is also (s, S).*

(b) *For the nominal problem without fixed cost, the optimal policy is (S, S) with the threshold at time k being $S_k = \bar{w}_k$.*

(c) *For the nominal problem with fixed cost, if we denote by t_j ($j = 1, \dots, J$) the times where stock is ordered (which are obtained by solving (9) subject to (10)–(12)) and s_j, S_j the corresponding thresholds at time t_j , we have*

$$S_j = \sum_{i=0}^{I_j} \bar{w}_{t_j+i} \tag{15}$$

and

$$s_1 = x_0 - \sum_{i=0}^{t_1-1} \bar{w}_i, \quad s_j = - \sum_{i=L_{j-1}+1}^{L_{j-1}-1} \bar{w}_{t_{j-1}+i}, \quad j \geq 2, \tag{16}$$

where

$$L_j = t_{j+1} - t_j \quad \text{and} \quad I_j = \left\lfloor \frac{pL_j - c1_{\{j=J\}}}{h+p} \right\rfloor.$$

PROOF. (a) See Bertsekas (1995) for the optimality of base-stock policies in the stochastic case. The nominal problem is a special case where the random variables are equal to their nominal value with probability 1.

(b) For the nominal case without fixed cost, the policy \mathbf{u} defined by

$$u_k = \begin{cases} \bar{w}_k - x_k & \text{if } x_k < \bar{w}_k, \\ 0 & \text{otherwise,} \end{cases}$$

is feasible and incurs the cost $COST = c(\sum_{i=0}^{T-1} \bar{w}_i - x_0) + h \sum_{k=0}^I (x_0 - \sum_{i=0}^k \bar{w}_i)$, where I is the largest integer k such that $x_0 - \sum_{i=0}^k \bar{w}_i \geq 0$. We assume that $I < T - 1$, otherwise the problem is trivial. We consider the dual of this linear programming problem:

$$\begin{aligned}
 & \text{maximize} && \sum_{k=0}^{T-1} \left(x_0 - \sum_{i=0}^k \bar{w}_i \right) (h\alpha_{k+1} - p\beta_{k+1}) \\
 & \text{subject to} && -h \sum_{j=k}^T \alpha_j + p \sum_{j=k}^T \beta_j \leq c, \quad k = 1, \dots, T, \\
 & && \alpha_k + \beta_k = 1, \quad \alpha_k \geq 0, \quad \beta_k \geq 0, \quad k = 1, \dots, T.
 \end{aligned}$$

The following solution is dual feasible with cost equal to $COST$, proving (b) by strong duality:

$$\alpha_k = \begin{cases} 1 & \text{if } x_0 - \sum_{i=0}^{k-1} \bar{w}_i \geq 0, \\ \frac{p-c1_{\{k=T\}}}{h+p} & \text{otherwise,} \end{cases} \quad \beta_k = 1 - \alpha_k.$$

(c) In the case with fixed cost, we consider the optimal ordering times as given (that is, \mathbf{v}^* is given). The problem becomes a linear programming problem. Let $t_j, j = 1, \dots, J$, be the times when an amount of stock u_j is ordered. The cost function can be decomposed in $J + 1$ pieces, the j th piece ($j = 0, \dots, J$) representing the cost incurred from time t_j up to t_{j+1} (not included),

with the conventions $t_0 = -1$ and $t_{J+1} = T$. The minimization problem is solved recursively backwards for $j = 1, \dots, J$, for the cumulative cost from step j onward. Let I_j be the greatest integer i in $[1, L_j]$ such that $x_{t_j+i} > 0$, where $L_j = t_{j+1} - t_j$. (If $x_{t_j+i} \leq 0$ for all $i \in [1, L_j]$, we take $I_j = 0$.) The cost function for the J th piece can be rewritten as

$$cu_J + h \sum_{k=0}^{I_j-1} \left(x_{t_j} + u_j - \sum_{i=0}^k \bar{w}_{t_j+i} \right) + p \sum_{k=I_j}^{L_j-1} \left(-x_{t_j} - u_j + \sum_{i=0}^k \bar{w}_{t_j+i} \right),$$

and is therefore linear in u_j with slope $(h + p)I_j + c - pL_j$, with u_j subject to the constraint $\sum_{i=0}^{I_j-1} \bar{w}_{t_j+i} < x_{t_j} + u_j \leq \sum_{i=0}^{I_j} \bar{w}_{t_j+i}$ from the definition of I_j . This function is minimized for

$$I_j^* = \left\lfloor \frac{pL_j - c}{h + p} \right\rfloor, \quad x_{t_j} + u_j^* = \sum_{i=0}^{I_j^*} \bar{w}_{t_j+i}. \quad (17)$$

Moreover, we have $x_{t_j+L_j} = x_{t_{j+1}} = x_{t_j} + u_j^* - \sum_{i=0}^{L_j-1} \bar{w}_{t_j+i} = -\sum_{i=I_j^*+1}^{L_j-1} \bar{w}_{t_j+i}$. At optimality, the cost function at the last time period is equal to $c(\sum_{i=0}^{I_j^*} \bar{w}_{t_j+i} - x_{t_j}) + h \sum_{k=0}^{I_j^*-1} (\sum_{i=k+1}^{I_j^*} \bar{w}_{t_j+i}) + p \sum_{k=I_j^*}^{L_j-1} (\sum_{i=I_j^*+1}^k \bar{w}_{t_j+i})$.

After step $j + 1$, we see that at optimality the cumulative cost at step $j + 1$ affects the cumulative cost at step j only through $-cx_{t_{j+1}}$, which depends on I_j and the data of the problem. The cumulative cost function at step j is minimized for

$$I_j^* = \left\lfloor \frac{pL_j}{h + p} \right\rfloor, \quad x_{t_j} + u_j^* = \sum_{i=0}^{I_j^*} \bar{w}_{t_j+i}. \quad (18)$$

Moreover, we have $x_{t_j+L_j} = x_{t_{j+1}} = x_{t_j} + u_j^* - \sum_{i=0}^{L_j-1} \bar{w}_{t_j+i} = -\sum_{i=I_j^*+1}^{L_j-1} \bar{w}_{t_j+i}$. Using the definition of s and S in an (s, S) policy, it follows immediately that

$$s_j = -\sum_{i=I_{j-1}^*+1}^{L_{j-1}-1} \bar{w}_{t_{j-1}+i} \quad \forall j \geq 2, \quad S_j = \sum_{i=0}^{I_j^*} \bar{w}_{t_j+i} \quad \forall j \geq 1, \quad (19)$$

with

$$I_j^* = \left\lfloor \frac{pL_j - c1_{\{j=J\}}}{h + p} \right\rfloor.$$

s_1 is equal to x_{t_1} and is obtained by the dynamics equation. \square

We next present the main result regarding the structure of the optimal robust policy.

THEOREM 3.2 (OPTIMAL ROBUST POLICY). (a) *The optimal policy in the robust formulation (14), evaluated at time 0 for the rest of the horizon, is the optimal policy for the nominal problem with the modified demand*

$$w'_k = \bar{w}_k + \frac{p-h}{p+h}(A_k - A_{k-1}), \quad (20)$$

where $A_k = q_k^* \Gamma_k + \sum_{i=0}^k r_{ik}^*$ is the deviation of the cumulative demand from its mean at time k , q^* and r^* being the optimal q and r variables in (14). (By convention, $q_{-1} = r_{-1} = 0$.) In particular, it is (S, S) if there is no fixed cost and (s, S) if there is a fixed cost.

(b) *If there is no fixed cost, the optimal robust policy is (S, S) with $S_k = w'_k$ for all k .*

(c) *If there is a fixed cost, the corresponding thresholds S_j, s_j , where $j = 1, \dots, J$, indexes the ordering times, are given by Equations (15) and (16) applied to the modified demand w'_k .*

(d) *The optimal cost of the robust problem (14) is equal to the optimal cost for the nominal problem with the modified demand plus a term representing the extra cost incurred by the robust policy, $(2ph)/(p+h) \sum_{k=0}^{T-1} A_k$.*

PROOF. Let $(\mathbf{u}^*, \mathbf{v}^*, \mathbf{q}^*, \mathbf{r}^*)$ be the optimal solution of (14). Obviously, setting the q and r variables to their optimal values q^* and r^* in (14) and resolving the linear programming problem will give \mathbf{u}^* and \mathbf{v}^* again. This enables us to focus on the optimal ordering policy only, taking the auxiliary variables q^*, r^* as given in the robust formulation (14). We then have to solve

$$\min_{\mathbf{u} \geq 0} \sum_{k=0}^{T-1} [cu_k + K1_{\{u_k > 0\}} + \max(h(\bar{x}_{k+1} + A_k), p(-\bar{x}_{k+1} + A_k))], \quad (21)$$

where $\bar{x}_{k+1} = x_0 + \sum_{i=0}^k (u_i - \bar{w}_i)$ and $A_k = q_k^* \Gamma_k + \sum_{i=0}^k r_{ik}^*$ for all k .

We define a modified stock variable x'_k , which evolves according to the linear equation

$$x'_{k+1} = x'_k + u_k - \underbrace{\left(\bar{w}_k + \frac{p-h}{p+h}(A_k - A_{k-1}) \right)}_{=w'_k}, \quad (22)$$

with $x'_0 = x_0$. Note that the modified demand w'_k is not subject to uncertainty. We have

$$\begin{aligned} & \max(h(\bar{x}_{k+1} + A_k), p(-\bar{x}_{k+1} + A_k)) \\ &= \max(hx'_{k+1}, -px'_{k+1}) + \frac{2ph}{p+h} A_k. \end{aligned} \quad (23)$$

The reformulation of the robust model, given the optimal q^* and r^* variables, as a nominal inventory problem in the modified stock variable x'_k (plus the fixed cost $(2ph/(p+h)) \sum_{k=0}^{T-1} A_k$) follows from injecting Equation (23) into formulation (21). This proves (a) and (d). We conclude that (b) and (c) hold by invoking Lemma 3.1. \square

REMARKS. 1. Using Problem (13) to define A_k by strong duality, and because $\Gamma_{k-1} \leq \Gamma_k$ for all k , we have $A_{k-1} \leq A_k$ for all k . Therefore, w'_k will be greater than \bar{w}_k if $p > h$ (that is, if shortage costs are more expensive than holding costs, we increase the “safety stock”), smaller than \bar{w}_k if $p < h$ (if holding costs are more expensive, we want to make sure that we will not be left with extra items), and equal to \bar{w}_k if $p = h$.

2. Using an argument similar to the one above, and because $\Gamma_k \leq \Gamma_{k-1} + 1$ for all k , we have $A_k \leq A_{k-1} + \hat{w}_k$ for all k . Therefore, at time k , w'_k belongs to

$$\begin{cases} [\bar{w}_k, \bar{w}_k + \frac{p-h}{p+h}\hat{w}_k] & \text{if } p \geq h \text{ and} \\ [\bar{w}_k + \frac{p-h}{p+h}\hat{w}_k, \bar{w}_k] & \text{if } p < h. \end{cases}$$

The extreme case is $p = h$, where $w'_k = \bar{w}_k$ for all k and Γ_k .

3. For the case without fixed cost, and for the case with fixed cost when the optimal ordering times are given, the robust approach leads to the thresholds in closed form. For instance, if the demand is i.i.d. ($\bar{w}_k = \bar{w}$, $\hat{w}_k = \hat{w}$ for all k), we have $A_k = \hat{w}\Gamma_k$ and, if there is no fixed cost,

$$S_k = w'_k = \bar{w} + \frac{p-h}{p+h}\hat{w}(\Gamma_k - \Gamma_{k-1}) \quad \text{for all } k.$$

Hence, the robust approach protects against the uncertainty of the demand while maintaining striking similarities with the nominal problem, remains computationally tractable, and is easy to understand intuitively.

3.2. The Capacitated Model

So far, we have assumed that there was no upper bound either on the amount of stock that can be ordered or on the amount of stock that can be held in the facility. In this section, we consider the more realistic case where such bounds exist. The other assumptions remain the same as in §3.1.

3.2.1. The Model with Capacitated Orders. The extension of the robust model to capacitated orders of maximal size d is immediate, by adding the constraint

$$u_k \leq d \quad \forall k \tag{24}$$

to formulation (14).

THEOREM 3.3 (OPTIMAL ROBUST POLICY). *The optimal robust policy is the optimal policy for the nominal problem with capacity d on the links and with the modified demand defined in Equation (20).*

PROOF. The reformulation of the robust model as a problem without uncertainty does not affect constraints on the orders u_k . Adding the constraints $u_k \leq d$ to this new problem, we obtain a deterministic model with capacity d on the links. \square

3.2.2. The Model with Capacitated Inventory. We now consider the case where stock can only be stored up to an amount C . This adds the following constraint to formulation (14):

$$x_0 + \sum_{i=0}^k (u_i - w_i) \leq C, \tag{25}$$

where $w_i = \bar{w}_i + \hat{w}_i \cdot z_i$ such that $\mathbf{z} \in \{|z_i| \leq 1 \forall i, \sum_{i=0}^k |z_i| \leq \Gamma_k \forall k\}$. This constraint depends on the uncertain parameters w_i . Applying the technique developed in §2 and using the same auxiliary problem (13) as before, we rewrite this constraint in the robust framework as

$$\bar{x}_{k+1} + q_k \Gamma_k + \sum_{i=0}^k r_{ik} \leq C \quad \forall k, \tag{26}$$

where q_k and r_{ik} are defined in (14).

We now analyze the optimal robust policy. We define the modified stock variables x'_k by $x'_{k+1} = x'_k + u_k - w'_k$ and $x'_0 = x_0$, with w'_k given by Equation (20) for all k . The inventory capacity constraint (26) becomes

$$x'_{k+1} \leq C - \frac{2p}{p+h}A_k \quad \forall k. \tag{27}$$

This deterministic problem in x'_k is not equivalent to a nominal problem with inventory capacity because the right-hand side in the new capacity constraint depends on the time period k , and worse, decreases with k . However, it never threatens the feasibility of the problem, in the following sense:

LEMMA 3.4. *For all k , if $x'_k \leq C_k$, then $x'_k - w'_k \leq C_{k+1}$, where $C_k = C - (2p/(p+h))A_{k-1}$.*

Therefore, if x'_k is feasible, it is always possible to satisfy the inventory capacity at time $k+1$ by not ordering.

PROOF. We need to show that

$$x'_k \leq C - \frac{2p}{p+h}A_{k-1} \quad \text{implies}$$

$$x'_k \leq C - \frac{2p}{p+h}A_k + w'_k \quad \text{for all } k.$$

Because

$$w'_k = \bar{w}_k + \frac{p-h}{p+h}(A_k - A_{k-1}),$$

it suffices to prove that $A_k - A_{k-1} \leq \bar{w}_k$. However, $A_k - A_{k-1} \leq \hat{w}_k$ from problem (13) and $\Gamma_k - \Gamma_{k-1} \leq 1$, and $\hat{w}_k \leq \bar{w}_k$ because the demand is always nonnegative. \square

We then have the following theorem.

THEOREM 3.5 (OPTIMAL ROBUST POLICY). *The optimal robust policy is the optimal policy for the nominal problem subject to the modified demand defined in Equation (20), with inventory capacity at time 0 equal to C , and with inventory capacity at time $k+1$, $k \geq 0$, equal to $C - (2p/(p+h))A_k$.*

PROOF. Follows from incorporating Equation (27) to formulation (14). \square

4. The Network Case

4.1. The Uncapacitated Model

We now extend the results of §3 to the network case. We first study the case of tree networks, which are well suited to describe supply chains because of their hierarchical structure: The main storage hubs (the sources of the network) receive their supplies from outside manufacturing plants and send items throughout the network, each time bringing them closer to their final destination, until they reach the stores (the sinks of the network). Let S be the number of sink nodes. When there is only one source and one sink, the tree network is called a series system.

We define echelon k for $k = 1, \dots, N$, with N the total number of nodes in the network, to be the union of all the installations, including k itself, that can receive stock from installation k , and the links between them. This is the definition used by Clark and Scarf (1960), as well as Zipkin (2000), when they consider tree networks. In the special case of series systems, we number the installations such that for $k = 1, \dots, N$, the items transit from installation $k + 1$ to k , with installation N receiving its supply from the plant and installation 1 being the only sink node, as in Clark and Scarf (1960). In that case, the demand at installation $k + 1$ at time t is the amount of stock ordered at installation k at the same time t .

We also define, for $k = 1, \dots, N$,

$I_k(t)$: stock available at the beginning of period t at installation k ,

$X_k(t)$: stock available at the beginning of period t at echelon k ,

$D_{i,k}(t)$: stock ordered at the beginning of period t at echelon k to its supplier i_k ,

$W_s(t)$: demand at sink node s during period t , $s = 1, \dots, S$.

Let $N(k)$ be the set of installations supplied by installation k and $O(k)$ the set of sink nodes in echelon k . We assume constant lead times equal to 0, backlog of excess demand, and linear dynamics for the stock at installation k over time ($k = 1, \dots, N$):

$$I_k(t+1) = I_k(t) + D_{i,k}(t) - \sum_{j \in N(k)} D_{kj}(t), \quad t = 0, \dots, T-1. \quad (28)$$

By convention, if k is a sink node s , $\sum_{j \in N(k)} D_{kj}(t) = W_s(t)$. This leads to the following dynamics for the stock at echelon k :

$$X_k(t+1) = X_k(t) + D_{i,k}(t) - \sum_{s \in O(k)} W_s(t), \quad t = 0, \dots, T-1. \quad (29)$$

Furthermore, the stock ordered by echelon k at time t is subject to the coupling constraint

$$\sum_{i \in N(k)} D_{ki}(t) \leq \max(I_k(t), 0) \quad \forall k, \forall t, \quad (30)$$

that is, the total order made to a supplier cannot exceed what the supplier currently has in stock, or, equivalently, the supplier can only send through the network items that it really has. Because the network was empty when it started operating at time $t_0 = -\infty$, it follows by induction on t that $I_k(t) \geq 0$ for all $k \geq 2$. Therefore, the coupling constraint between echelons is linear and can be rewritten as

$$\sum_{i \in N(k)} D_{ki}(t) \leq \bar{X}_k(t) - \sum_{i \in N(k)} \bar{X}_i(t) \quad \forall k, \forall t, \quad (31)$$

with $\bar{X}_i(t+1) = X_i(0) + \sum_{\tau=0}^t \{D_{ki}(\tau) - \sum_{s \in O(i)} \bar{W}_s(\tau)\}$ for all i and t , where k supplies i , i.e., $\bar{X}_i(t+1)$ is the level of inventory at time t if the demand is deterministic, equal to its nominal value. Neither the echelon inventories nor the orders are capacitated. Section 4.2 deals with the capacitated case.

Finally, we specify the cost function. We assume that each echelon k has the same cost structure as the single installation modelled in §3.1 with specific parameters (c_k, K_k, h_k, p_k) . We also keep the same notations and assumptions here as in §3.1 regarding the uncertainty structure at each sink node. In particular, each sink node s has its own threshold sequence $\Gamma_s(t)$ evolving over time that represents the total budget of uncertainty allowed up to time t for sink s . We have $W_s(t) = \bar{W}_s(t) + \widehat{W}_s(t) \cdot Z_s(t)$ such that the $Z_s(t)$ belong to the set $\mathcal{P}_s = \{|Z_s(t)| \leq 1 \forall t, \sum_{\tau=0}^t Z_s(\tau) \leq \Gamma_s(t) \forall t\}$. We assume that $0 \leq \Gamma_s(t) - \Gamma_s(t-1) \leq 1$ for all s and t , that is, the budgets of uncertainty are increasing in t at each sink node, but cannot increase by more than 1 at each time period.

Applying the robust approach developed in §2 to the holding/shortage constraints in the same manner as in §3, we obtain the mixed-integer programming problem

$$\begin{aligned} & \text{minimize} \quad \sum_{t=0}^{T-1} \sum_{k=1}^N \sum_{i \in N(k)} \{c_{ki} D_{ki}(t) + K_{ki} V_{ki}(t) + Y_i(t)\} \\ & \text{subject to} \\ & Y_i(t) \geq h_i \left\{ \bar{X}_i(t+1) + \sum_{s \in O(i)} \left(q_s(t) \Gamma_s(t) + \sum_{\tau=0}^t r_s(\tau, t) \right) \right\} \quad \forall i, \forall t, \\ & Y_i(t) \geq p_i \left\{ -\bar{X}_i(t+1) + \sum_{s \in O(i)} \left(q_s(t) \Gamma_s(t) + \sum_{\tau=0}^t r_s(\tau, t) \right) \right\} \quad \forall i, \forall t, \\ & \sum_{i \in N(k)} D_{ki}(t) \leq \bar{X}_k(t) - \sum_{i \in N(k)} \bar{X}_i(t) \quad \forall k, \forall t, \\ & q_s(t) + r_s(\tau, t) \geq \widehat{W}_s(\tau) \quad \forall s, \forall t, \forall \tau \leq t, \\ & q_s(t) \geq 0, \quad r_s(\tau, t) \geq 0 \quad \forall s, \forall t, \forall \tau \leq t, \\ & 0 \leq D_{ki}(t) \leq M V_{ki}(t), \quad V_{ki}(t) \in \{0, 1\} \quad \forall k, \forall i \in N(k), \forall t. \end{aligned} \quad (32)$$

As in the single-station case, an attractive feature of this approach is that the robust model of a supply chain remains of the same class as the nominal model, that is, a linear programming problem if there are no fixed costs, and a mixed-integer programming problem if fixed costs are present. Therefore, the proposed methodology is numerically tractable for very general topologies.

First, we study the optimal policy in the nominal case. We have the following result.

LEMMA 4.1 (OPTIMAL NOMINAL POLICY). (a) *For the problem without uncertainty, the optimal policy for each echelon k is the optimal policy obtained for a single installation with time-varying capacity on the orders, subject to the demand $\sum_{s \in O(k)} \bar{W}_s(t)$ at time t .*

(b) *In the case without fixed costs, it is also the optimal policy obtained for a single installation with new, time-varying cost coefficients, without capacity, subject to the demand $\sum_{s \in O(k)} \bar{W}_s(t)$ at time t .*

(c) *In the case without fixed costs, the optimal policy at each echelon is base stock for the new parameters of the system.*

PROOF. Let us first analyze the series system case. The coupling constraint for echelon k at time t is then $D_k(t) \leq I_{k+1}(t)$. We analyze the optimal orders by setting the right-hand sides of the coupling constraints to their optimal values. The coupling constraint for echelon k at time t becomes $D_k(t) \leq C_k(t)$ for some given $C_k(t)$. Hence, the problem is now decoupled in the echelons and is equivalent to solving a capacitated single-station inventory problem with or without fixed cost at each echelon, subject to the nominal demand at the sink node and with the original cost parameters. This proves (a) for series systems.

In the general network case, because the coupling constraints (31) bound the total order made at an installation by its customers, it cannot be directly interpreted as a capacity on the orders made by each customer. We analyze the nominal problem by duplicating the coupling constraint and writing it as

$$D_{ki}(t) \leq \bar{X}_k(t) - \sum_{i \in N(k)} \bar{X}_j(t) - \sum_{j \in N(k), j \neq i} D_{kj}(t) \quad \forall k, \forall t \tag{33}$$

for each echelon i supplied by installation k , and setting the right-hand side of this new constraint to its optimal value, to obtain a time-varying capacity on the orders made by each echelon. (a) follows immediately.

For any network, the inventory problem in the case without fixed costs is a linear programming problem. We dualize the coupling constraints (31) through a Lagrangean multiplier approach. The feasible set of the relaxation is now separable in the echelons, and the cost function of the relaxation can be rewritten as the sum of separable single-installation problems, with new cost parameters that incorporate the Lagrangean multipliers. It follows from the

theory of Lagrangean relaxation for linear programming problems that the cost of the relaxation of the problem is equal to the cost of the original problem. This proves (b).

(c) follows from applying Lemma 3.1 to (b). \square

We now give the main theorem regarding the optimal policy in the robust approach.

THEOREM 4.2 (OPTIMAL ROBUST POLICY). (a) *The optimal policy in the robust formulation (32) for echelon k is the optimal policy obtained for the supply chain subject to the modified, deterministic demand at sink node s (for $s \in O(k)$):*

$$W'_{s,k}(t) = \bar{W}_s(t) + \frac{p_k - h_k}{p_k + h_k} (A_s(t) - A_s(t-1)), \tag{34}$$

where $A_s(t) = q_s^*(t)\Gamma_s(t) + \sum_{\tau=0}^t r_s^*(\tau, t)$, q_s^* and r_s^* being the optimal q and r variables associated with sink node s in (32).

(b) *The optimal cost in the robust case for the tree network is equal to the optimal cost of the nominal problem for the modified demands, plus a term representing the extra cost incurred by the robust policy, $\sum_{k=1}^N (2p_k h_k / (p_k + h_k)) \sum_{t=0}^{T-1} \sum_{s \in O(k)} A_s(t)$.*


PROOF. We reformulate the problem as a nominal problem in the same way as in the proof of Theorem 3.2, and invoke Lemma 4.1. \square

The case of more general supply chains is complex because they cannot be reduced to a tree network: The need might arise to order from a more expensive supplier when the cheapest one does not have enough inventory. We can still define echelons for those networks in a similar manner as before, and the evolution of the stock at echelon k , which is supplied by the set of installations $I(k)$ and has the set $O(k)$ as its sink nodes, is described by the following linear equation:

$$X_k(t+1) = X_k(0) + \sum_{\tau=0}^t \left\{ \sum_{i \in I(k)} D_{ik}(\tau) - \sum_{j \in O(k)} W_j(\tau) \right\} \quad \forall t. \tag{35}$$

With the standard cost assumptions used before, the echelons cannot be studied independently and the optimal policy is not necessarily base stock, even in the simple case of demand without uncertainty. This is illustrated in Figure 1.

Figure 1. A network for which the optimal policy is not base stock.



The network in Figure 1 has two installations, and therefore two echelons. Echelon 1 can be supplied by installation 2 at a unit cost $c_1 = 1$, without any fixed ordering cost, and has the option to order directly from the plant for the same unit cost $c_2 = 1$, but with an additional fixed cost $K_2 = 4$ incurred whenever an order is made. This option is attractive only if installation 2 does not have enough stock in inventory. The holding and shortage unit costs at echelon 1 are $h_1 = p_1 = 2$. The horizon is one time period, and the demand at time 0 is deterministic, equal to 10 units. Echelon 1 has only five units in inventory at time 0.

Comparing the two options, it is easy to see that it is optimal for echelon 1 to order five units from 2 if 2 has five units in inventory at time 0, two units from 2 and none from the plant if 2 has two units, and five units from the plant if 2 has no item in inventory. Therefore, the optimal amount of stock on hand and on order at echelon 1 at time 0 is 10, respectively, 7, 10 units if installation 2 has 5, respectively, 2, 0 units in inventory at time 0. Thus, the optimal policy is not base stock.

Also, while we can reformulate the robust problem as a new problem with modified demand in the same fashion as before, it loses some of its meaning because distinct echelons can now “see” the same sink node but different demands at this node (because of the cost parameters specific to each echelon, which appear in the expression of the modified demand). Hence, it is not clear how they can work together to meet this demand optimally.

However, the proposed robust methodology remains numerically tractable in a wide range of settings, in particular with a holding/shortage cost at the installation level instead of the echelon. This illustrates the applicability of the proposed approach to different cost structures.

4.2. The Capacitated Model

We now refine our description of the inventory problem in a supply chain by introducing upper bounds on the amount of stock that can be ordered and/or held at any time and at any echelon. As explained in §3.2, an upper bound on the maximal order can be directly introduced in the proposed approach, by adding the constraint

$$D_{ki}(t) \leq d_{ki} \quad \forall k, \forall i \in \mathcal{N}(k), \forall t, \quad (36)$$

to formulation (32). Inventory capacity, however, requires further manipulation, because the level of inventory held at an echelon at any time depends on the demand, which is subject to uncertainty. Manipulations similar to those in §3.2 lead to the constraint

$$\bar{X}_k(t+1) + \sum_{s \in \mathcal{O}(k)} \left(q_s(t) \Gamma_s(t) + \sum_{\tau=0}^t r_s(\tau, t) \right) \leq C_k \quad \forall k, \forall t, \quad (37)$$

to be added to the formulation, $q(t)$ and $r(\tau, t)$ being defined as in (34).

We next study the structure of the optimal policy.

THEOREM 4.3 (OPTIMAL POLICY WITH LINK AND ECHELON CAPACITIES). *The optimal policy is the optimal policy of the nominal problem with the modified demand defined in Equation (34), time-varying echelon capacities*

$$C'_k(t+1) = C_k - \frac{2p_k}{p_k + h_k} \sum_{s \in \mathcal{O}(k)} A_s(t),$$

where C_k is the original capacity at echelon k , and link capacities $D_{ki}(t) \leq d_{ki}$ for all $k, i \in \mathcal{N}(k)$, and t .

PROOF. Follows from incorporating Equations (36) and (37) into formulation (32). \square

5. Numerical Implementation

The purpose of this computational study is to investigate the performance of the robust approach when the underlying distributions are not perfectly known. In particular, its aim is to provide some insight into the relative performance of the robust approach and traditional methods: dynamic programming in the single-station case and myopic policies for more complex networks. The topics we seek to address are

- the robustness of both methods with respect to changes in probability distributions,
- the effect of various parameters (holding, shortage, ordering-cost parameters, and time horizon) on the relative performance of the two methods.

To meet these goals, we consider various values for the parameters of the system, and study several (assumed and realized) distributions. Section 5.1 describes how to select the budgets of uncertainty. Section 5.2 applies the methodology to the single-station case, and §5.3 presents results for series systems and tree supply chains. The conclusions of these numerical experiments are summarized in §5.4.

5.1. Selection of the Budgets of Uncertainty

We next investigate how to select the budgets of uncertainty $\Gamma_s = (\Gamma_s(0), \dots, \Gamma_s(T-1))$ to guarantee performance while protecting against the uncertainty of the demands $W_s(t)$, where $s = 1, \dots, S$ are the sink nodes, with a limited amount of information on the distributions. We feel that this partial knowledge of the distributions describes well the situation in practice. We assume that the random variables are uncorrelated and that we know the mean $\bar{W}_s(t)$ and the variance $\sigma_s^2(t)$ of the demand at each sink and each time period. Let \mathcal{W} be the set of random vectors with admissible distributions, and let $C_W(\Gamma)$ be the cost incurred in real life by the optimal policy in the robust problem (32) at Γ given, when the random demand sequence takes the value $\mathbf{W} = (W_s(0), \dots, W_s(T-1))_s$. Bertsimas and Popescu (2002) describe how to incorporate additional information on other moments of the distribution. We have,

for any admissible distribution,

$$E[C_W(\Gamma)] = \sum_{t=0}^{T-1} \sum_{k=1}^N \sum_{i \in N(k)} \left\{ c_{ki} D_{ki}(t) + K_{ki} V_{ki}(t) + h_i \bar{X}_i(t+1) + (h_i + p_i) \cdot E \max \left(0, \sum_{\tau=0}^t \sum_{s \in O(i)} W_s(\tau) - \left(X_i(0) + \sum_{\tau=0}^t D_{ki}(\tau) \right) \right) \right\}. \tag{38}$$

We will need the following lemma, which was also derived by Moon and Gallego (1994). An equivalent tight lower bound for $E \min(X, a)$ is due to Scarf (1958).

LEMMA 5.1 (OPTIMAL UPPER BOUND; SEE LO 1987 AND BERTSIMAS AND POPESCU 2002). *Let X be a nonnegative random variable with mean μ and variance σ^2 , which we denote as $X \sim (\mu, \sigma^2)^+$, and let $a > 0$. We have*

$$\max_{X \sim (\mu, \sigma^2)^+} E[\max(0, X - a)] = f(a - \mu, \mu, \sigma^2), \tag{39}$$

where f is the convex function defined by

$$f(x, \mu, \sigma^2) = \begin{cases} \frac{1}{2}[-x + \sqrt{\sigma^2 + x^2}] & \text{if } x \geq \frac{\sigma^2 - \mu^2}{2\mu}, \\ -x \frac{\mu^2}{\mu^2 + \sigma^2} + \mu \frac{\sigma^2}{\mu^2 + \sigma^2} & \text{if } x < \frac{\sigma^2 - \mu^2}{2\mu}. \end{cases} \tag{40}$$

PROOF. See Bertsimas and Popescu (2002), which also describes how to incorporate moments of higher order. \square

Our goal is therefore to minimize in Γ :

$$\sum_{t=0}^{T-1} \sum_{k=1}^N \sum_{i \in N(k)} \{ c_{ki} D_{ki}(t) + K_{ki} V_{ki}(t) + h_i \bar{X}_i(t+1) + (h_i + p_i) f(\bar{X}_i(t+1), M_i(t+1), S_i(t+1)) \} \tag{41}$$

with $M_i(t+1) = \sum_{\tau=0}^t \sum_{s \in O(i)} \bar{W}_s(\tau)$ and $S_i(t+1) = \sqrt{\sum_{\tau=0}^t \sum_{s \in O(i)} \sigma_s^2(\tau)}$, where the $D_{ki}(t)$ depend on Γ as the optimal solution of the robust problem (32). As before, $V_{ki}(t) = 1$ if $D_{ki}(t) > 0$ and 0 otherwise, and Γ is subject to the constraints $0 \leq \Gamma_s(t) - \Gamma_s(t-1) \leq 1$ for all s and t .

Because of the complex dependence of the $D_{ki}(t)$ and $\bar{X}_i(t)$ in Γ , and of the presence of the binary variables $V_{ki}(t)$, we will instead consider the following problem:

$$\text{minimize } \sum_{t=0}^{T-1} \sum_{k=1}^N \sum_{i \in N(k)} \left\{ c_{ki} W'_{s,i}(t) + h_i \bar{X}_i(t+1) + (h_i + p_i) \cdot f(\bar{X}_i(t+1), M_i(t+1), S_i(t+1)) \right\} \tag{42}$$

$$\text{subject to } \bar{X}_i(t+1) = \sum_{\tau=0}^t \sum_{s \in O(i)} (W'_{s,i}(\tau) - \bar{W}_s(\tau)) \quad \forall i, t,$$

$$0 \leq \Gamma_s(t) - \Gamma_s(t-1) \leq 1 \quad \forall s, t.$$

This assumes that the fixed costs K_{ki} do not play a significant role in the function to minimize, and that $\bar{X}_i(t+1)$ can be approximated by $\sum_{\tau=0}^t \sum_{s \in O(i)} (W'_{s,i}(\tau) - \bar{W}_s(\tau))$. If the fixed costs are indeed significant, we could also solve the problem for prespecified ordering times, for instance, those of the nominal problem, which would take into account the fixed costs.

REMARK. From Theorem 3.2, Problem (42) yields the exact optimal Γ for single-station problems without fixed cost and with nonpositive initial inventory.

ALGORITHM 5.2 (SELECTION OF THE BUDGETS OF UNCERTAINTY). At time 0, we select Γ according to the following procedure:

If the demands are i.i.d., we solve

$$\begin{aligned} \text{minimize } & \sum_{k=1}^N \sum_{i \in N(k)} c_{ki} \frac{p_i - h_i}{p_i + h_i} \sum_{s \in O(i)} \widehat{W}_s \Gamma_s(T-1) \\ & + \sum_{t=0}^{T-1} \sum_{k=1}^N \sum_{i \in N(k)} \{ h_i \bar{X}_i(t+1) + (h_i + p_i) \\ & \cdot f(\bar{X}_i(t+1), M_i(t+1), S_i(t+1)) \} \end{aligned} \tag{43}$$

$$\text{subject to } \bar{X}_i(t+1) = \frac{p_i - h_i}{p_i + h_i} \sum_{s \in O(i)} \widehat{W}_s \Gamma_s(t) \quad \forall i, t,$$

$$0 \leq \Gamma_s(t) - \Gamma_s(t-1) \leq 1 \quad \forall s, t,$$

where the function f has been defined in Lemma 5.1.

If the demands are not i.i.d., we replace \widehat{W}_s by $\sum_{\tau=0}^{T-1} \widehat{W}_s(\tau)/T$ in the cost function of problem (43) and by $\sum_{\tau=0}^t \widehat{W}_s(\tau)/(t+1)$ in the definition of $X_i(t+1)$ for all t and $s \in O(i)$.

In the case with i.i.d. demand, we have used that if there is no fixed cost, $S_i(t) = \sum_{s \in O(i)} \bar{W}_s(t) + (p_i - h_i)/(p_i + h_i) \cdot \sum_{s \in O(i)} \widehat{W}_s(t) \Gamma_s(t)$, which is a direct extension of Theorem 3.2.

REMARK. The approximations are only used to find the optimal budgets of uncertainty Γ . The robust problem (32) itself does not rely on approximations; therefore, the optimal orders $D_{ki}(t)$ obtained by solving (32) take into account fixed costs and non-i.i.d. demand, when applicable.

5.2. Example of a Single Station

We now apply the proposed methodology to the example of minimizing the cost at a single station. The horizon is $T = 20$ time periods; the station has zero initial inventory, with an ordering cost per unit $c = 1$, a holding cost $h = 4$, and a shortage cost $p = 6$, in appropriate measurement units. There is no fixed ordering cost. The stochastic demand is i.i.d., with mean $\bar{w} = 100$ and standard deviation $\sigma = 20$ (unless specified otherwise). In the robust framework, we take $\widehat{w} = 2 \cdot \sigma$, that is, the demand belongs to the interval $[\bar{w} - 2 \cdot \sigma, \bar{w} + 2 \cdot \sigma]$.

In the first set of experiments, the stochastic policy is computed using a binomial distribution. In the second set of experiments, the stochastic policy is computed using an approximation of the Gaussian distribution on five points ($\bar{w} - 2\sigma, \bar{w} - \sigma, \bar{w}, \bar{w} + \sigma, \bar{w} + 2\sigma$). In both cases, the realized distribution is gamma, lognormal or Gaussian, with the same mean \bar{w} and standard deviation σ . These choices of assumed and realized distributions are motivated by the fact that we want to study the sensitivity of the performance of the robust approach relative to dynamic programming for several combinations of assumed and realized distributions. Note that we implement dynamic programming with the assumed distribution, but estimate its performance (as well as the performance of the robust approach) using the realized distribution.

The key metric we consider is the relative performance of the robust policy compared to the stochastic policy obtained by dynamic programming, as measured by the ratio $R = 100 \cdot (E(DP) - E(ROB))/E(DP)$, in percent. The expectations are computed with respect to the realized probability distribution, on a sample of size 1,000. In particular, when $R > 0$ the robust policy leads to lower costs on average than with the stochastic policy. We are also interested in the sample probability distribution of the costs DP and ROB .

The numerical experiments aim to provide some insight into the relationship between the performance of the robust policy and

- the assumed and realized distributions,
- the standard deviation of the demand,
- the cost parameters c, h , and p .

Budgets of Uncertainty. The budgets of uncertainty Γ_k are computed from Algorithm 5.2. For $c = 0$, we have

$$\Gamma_k = \min\left(\frac{\sigma}{\bar{w}} \sqrt{\frac{k+1}{1-\alpha^2}}, k+1\right), \tag{44}$$

with $\alpha = (p - h)/(p + h)$. Therefore, the budgets Γ_k will be equal to their worst-case value $k + 1$ for $k \leq \bar{k}$ with

$$\bar{k} = \frac{1}{1-\alpha^2} \left(\frac{\sigma}{\bar{w}}\right)^2 - 1 \quad \text{and to} \quad \frac{\sigma}{\bar{w}} \sqrt{\frac{k+1}{1-\alpha^2}}$$

otherwise (assuming $T \geq \bar{k} + 1$). Note that if the time horizon is smaller than $\bar{k} + 1$, then the robust approach is the same as the worst-case approach. However, this happens for extreme values of the parameters. For example, when $\alpha^2 = 0.9$ ($p/h = 37.97$) and $\bar{w} = 2\sigma$, then $\bar{k} + 1 = 3$. In other words, even in the extreme case that the shortage cost is close to 40 times the holding cost, the robust approach will differ from the worst-case approach if the time horizon is more than three periods.

Figure 2 (respectively, Figure 3) shows Γ_k as a function of k for p varying and $h = 4$ (respectively, h varying and $p = 4$), with $c = 0$ on the left panel and $c = 1$ on the right panel. The general trend is that Γ_k evolves as $\sqrt{k+1}$, although there are differences for the last few time periods.

In Figures 4–10 we report the performance of the robust and dynamic programming policies when the assumed distribution is binomial (respectively, Gaussian) on the left (respectively, right) panel of each figure.

Impact of the Standard Deviation. Figure 4 shows how the ratio $R = 100 \cdot (E(DP) - E(ROB))/E(DP)$, in percent, evolves as the ratio σ/\bar{w} increases (i.e., as the standard deviation increases because we keep \bar{w} constant). When the assumed distribution is binomial, the ratio R increases as the standard deviation increases and the robust policy outperforms dynamic programming by up to 10% to 13%, depending on the realized distribution. When the assumed distribution is Gaussian, the two methods are equivalent because the robust policy outperforms dynamic programming by at most 0.4%. In both cases, R shows the same qualitative trend as σ increases for the three realized distributions implemented here, although spreads in numerical values increase as σ increases.

Figure 2. Budgets of uncertainty for $c = 0$ (left) and $c = 1$ (right), as p varies with $h = 4$.


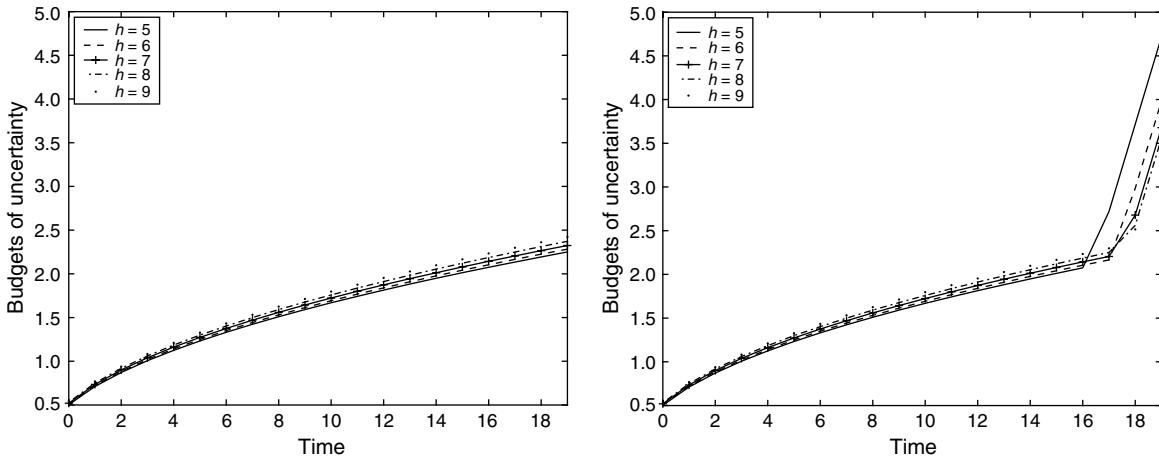


Figure 3. Budgets of uncertainty for $c = 0$ (left) and $c = 1$ (right), as h varies with $p = 4$.



A Probabilistic View of Performance. Figure 5 shows the sample probability distribution of the costs *ROB* and *DP*. Note that when a binomial distribution is assumed (the graph on the left), *ROB* yields significantly lower cost than *DP*, when the realized distribution is different. In contrast, when a Gaussian distribution is assumed (the graph on the right), the two approaches lead to very close costs.

Impact of the Cost Parameters. In Figures 6–10, we study the impact on performance of the cost parameters c , h , and p for $\sigma = 20$. We change one parameter (c , h , or p) and show the results in terms of the performance ratio $R = 100 \cdot (E(DP) - E(ROB))/E(DP)$ (Figures 6–8) and of the fill rates obtained by the robust approach and by dynamic programming (Figures 9 and 10). As the fill rates are almost identical for the three realized distributions, for clarity we only show the graphs obtained when the distribution is gamma. The numerical evidence suggests that:

- The key insight is that if $h < h_0$ ($h_0 \approx 2.5$, when $p = 6$), then *DP* outperforms *ROB*, while if $h > h_0$, *ROB*

outperforms *DP*. This remains valid independently of the assumed and realized distributions.

- Whether *DP* or *ROB* performs better is by and large not influenced by the ordering cost c . In Figure 6, *ROB* outperforms *DP* (because $h > h_0$), but the degree of out-performance decreases as c increases.

- If $h < h_1$, then the fill rate based on *DP* is larger than the fill rate based on *ROB*. The exact value of h_1 depends on the assumed distribution. If $h > h_1$, then the fill rate based on *DP* is by and large smaller than the fill rate based on *ROB*. For example, when a binomial distribution is assumed, $h_1 = p = 6$, and when $h > h_1$, the fill rate of *DP* drops below 0.8, compared to 0.92 for *ROB*.

5.3. Examples of Networks

5.3.1. A Series System. In this section, we apply the proposed approach to the series system in Figure 11 over $T = 10$ time periods. Each order made at echelon 1 or 2 incurs a fixed cost $K_1 = K_2 = 10$ and a cost per unit ordered

Figure 4. Impact of the standard deviation on performance.

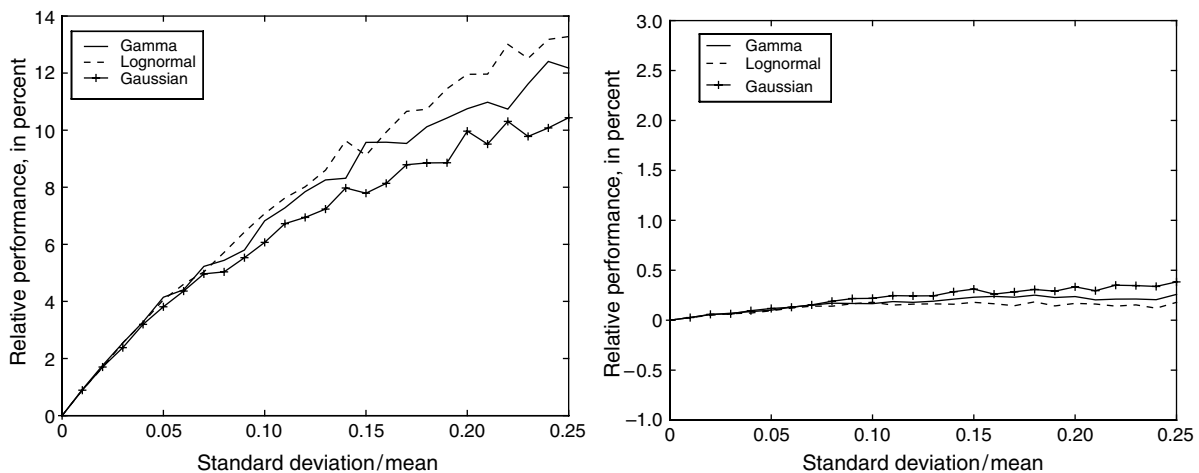


Figure 5. Sample probability distributions.

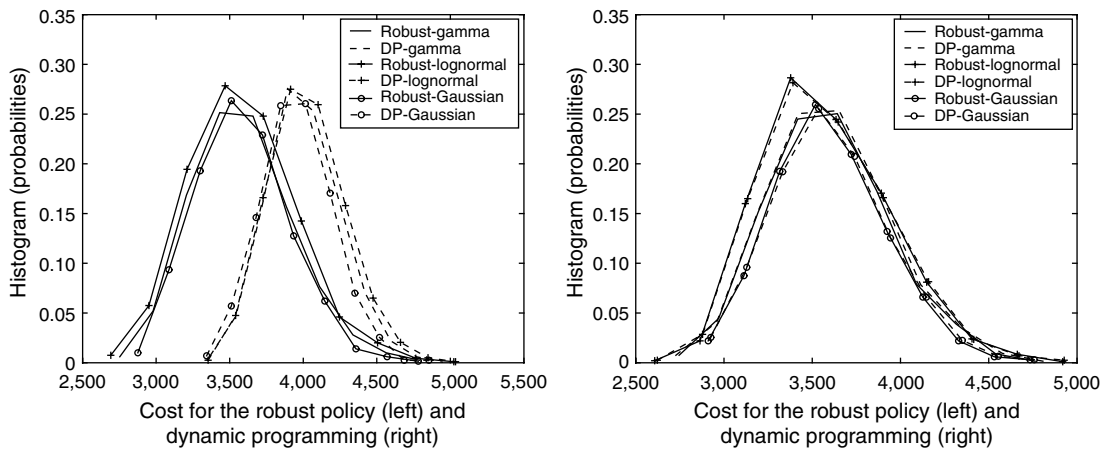


Figure 6. Impact of the ordering cost on performance.

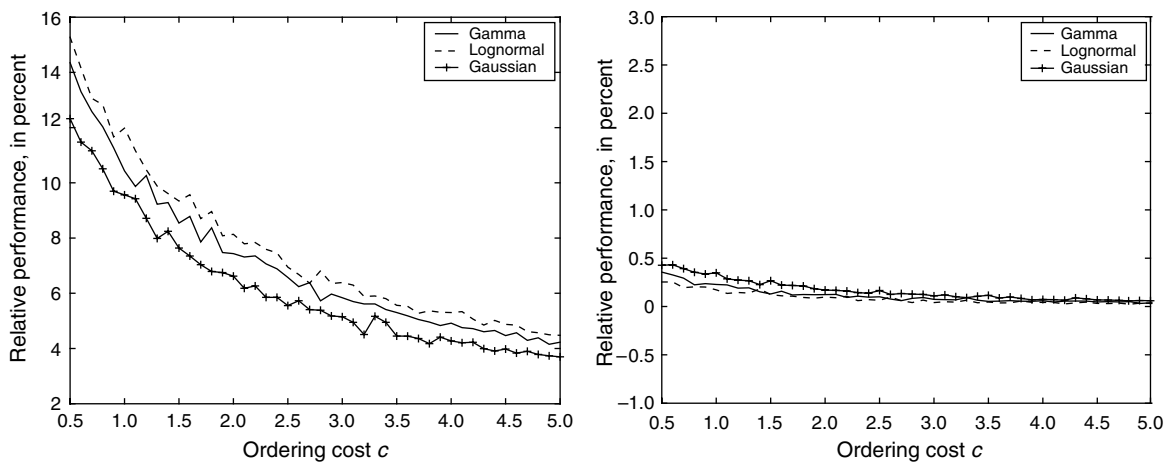


Figure 7. Impact of the holding cost on performance.




Figure 8. Impact of the shortage cost on performance.

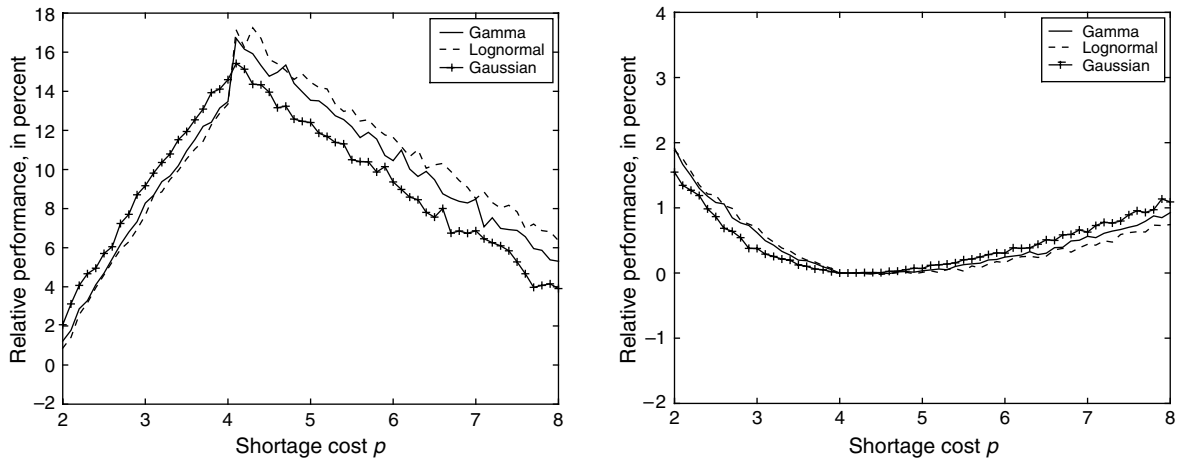


Figure 9. Impact of the holding cost on the fill rates.

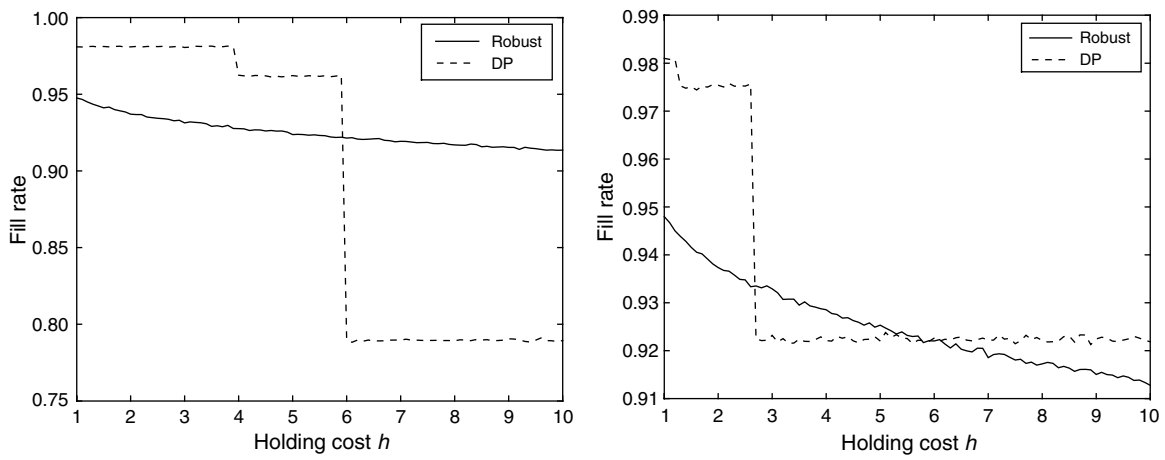


Figure 10. Impact of the shortage cost on the fill rates.

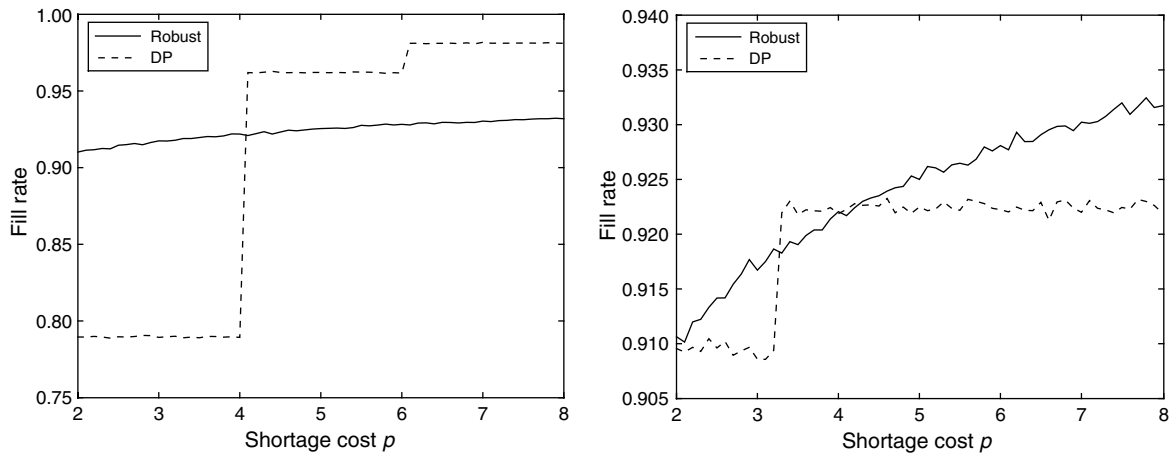



Figure 11. A series system.



$c_1 = c_2 = 1$. The holding and shortage costs at echelon 1 are $h_1 = 4$, $p_1 = 12$, while holding and shortage are penalized equally at echelon 2: $h_2 = p_2 = 4$. The stochastic demand is i.i.d. with mean $\bar{W} = 100$ and $\sigma = 20$. In the robust model, we take $\hat{W} = 2 \cdot \sigma$. Echelons 1 and 2 hold zero inventory at time 0.

Here we are interested in comparing the performance of the robust policy and the myopic policy obtained for the same two assumed distributions, as in §5.2 (binomial and approximate Gaussian on five points). The realized distributions remain gamma, lognormal, or Gaussian with the same first two moments.

The budgets of uncertainty at time 0 for all remaining time periods are computed using Algorithm 5.2. We reoptimize the problem at each time period for both the robust and the myopic approach. For the robust approach, we update the budgets of uncertainty at time t , $t = 1, \dots, T - 1$, as follows:

$$\Gamma^{(t)}(\tau) = \Gamma^{(t-1)}(\tau + 1) - \Gamma^{(t-1)}(1) \quad (45)$$

for $\tau = 0, \dots, T - t - 1$. (If we reoptimized the budgets of uncertainty, we would consistently overprotect the first time period in the horizon.) This updating rule does not change $\Gamma^{(t)}(\tau) - \Gamma^{(t)}(\tau - 1)$, which we use to define the modified demand, and therefore is consistent with the implementation of the base-stock levels in §5.2.

The performance is measured by the sample probability distributions of the costs *MYO* and *ROB*, and the ratio $r = 100 \cdot (MYO - ROB)/MYO$, in percent. As before,

the left (respectively, right) panels of the figures show the results obtained assuming a binomial distribution (respectively, approximate Gaussian).

Costs of the Robust and the Myopic Policies. The realized distribution of the demand and, more surprisingly, the distribution assumed to compute the myopic policy, do not appear to play a significant role in the sample probability distribution of the costs of the two policies shown in Figure 12. The robust policy clearly outperforms the myopic policy because it leads to costs with a lower mean and variance.

Sample Distribution of Performance Ratio. When the assumed distribution is binomial, all values of r are positive and the sample probability distribution of r is mostly concentrated on the range [20%, 45%] as shown in Figure 13. When the assumed distribution is Gaussian, $r < 0$ has a very low probability, and the sample probability distribution of r is mostly concentrated on the range [15%, 35%]. Therefore, it appears that having similar assumed and realized distributions does reduce the relative difference between the two policies, but the very use of a myopic policy leads to a significant cost disadvantage. This further demonstrates the high potential of the robust policy when applied to supply chains.

5.3.2. A Tree Network. In this section, we apply the proposed approach to the supply chain in Figure 14 over $T = 5$ time periods. Echelons 1 and 2 consist of installations 1 and 2, respectively. Echelon 3 consists of installations 1, 2, and 3, and the links in-between. The unit ordering costs are $c_1 = c_2 = c_3 = 1$. There are no fixed costs. Holding and shortage are penalized equally at echelons 1 and 2: $h_1 = p_1 = 8$, $h_2 = p_2 = 8$, while the holding and shortage costs at echelon 3 are $h_3 = 5$, $p_3 = 7$. Demands at installations 1 and 2 are i.i.d., with the same nominal demand $\bar{W} = 100$ and the same standard deviation σ . What distinguishes echelons 1 and 2 is their initial inventory: 150 items at echelon 1, 50 at echelon 2, and the realized

Figure 12. Costs of robust and myopic policies.

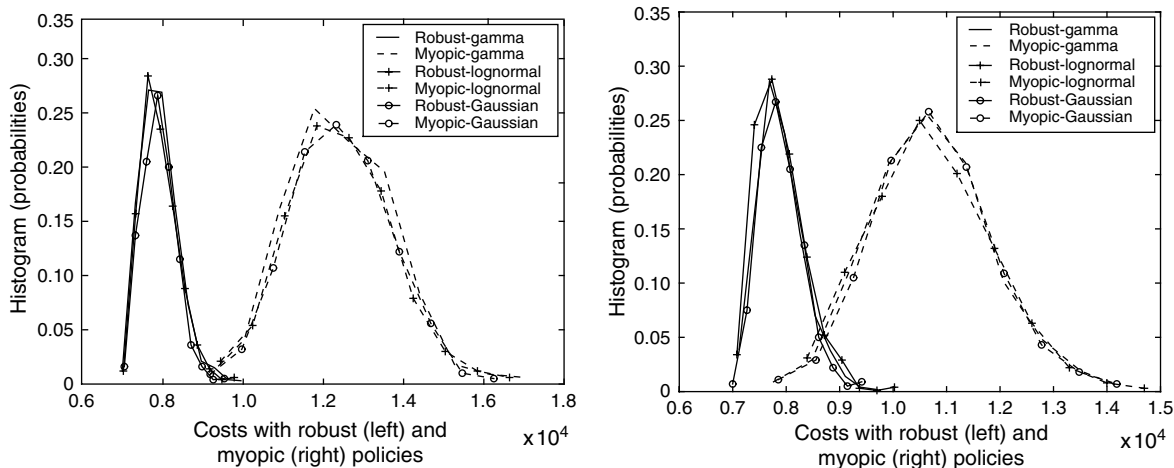
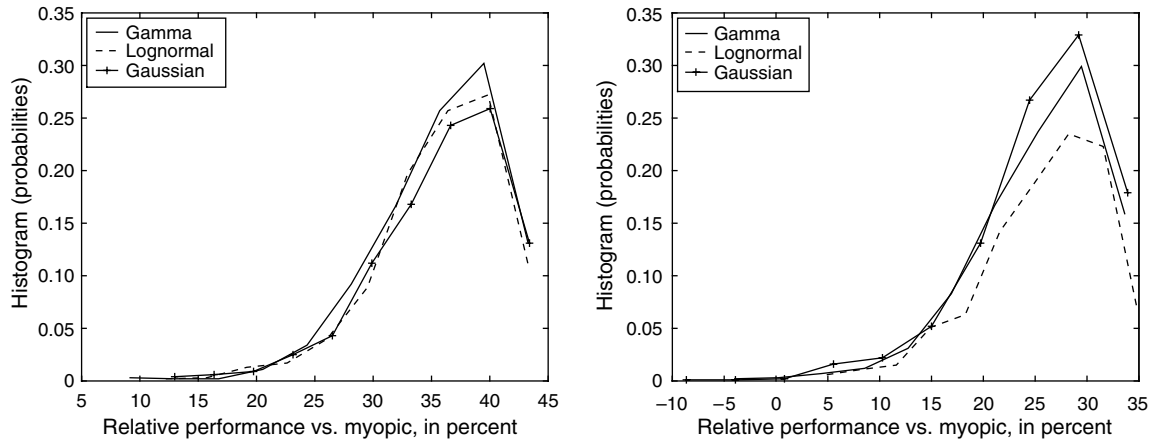


Figure 13. Sample distribution of relative performance.




demand distributions, which are either gamma, lognormal or Gaussian, with the same mean and standard deviation. Echelon 3 initially holds 300 items. Expected costs are computed using a sample size of 1,000.

As in §5.3.1, we are interested in comparing the performance of the robust policy and the myopic policy obtained for the same two assumed distributions as in §5.2 (binomial and approximate Gaussian on five points). Here, we assume the same distribution at stations 1 and 2 when computing the myopic policy. The realized distributions remain gamma, lognormal, or Gaussian with the same first two moments, and can differ between stations. In the robust framework, $\hat{W} = 2 \cdot \sigma$ and the budgets of uncertainty are computed at time 0 for all remaining time periods using Algorithm 5.2. We reoptimize the problem at each time period for both the robust and the myopic approach, and update the budgets of uncertainty as in §5.3.1.

The performance is measured by the sample probability distributions of the costs *MYO* and *ROB*, and the ratio $r = 100 \cdot (MYO - ROB)/MYO$, in percent. As before, the left (respectively, right) panels of the figures show the results obtained assuming a binomial distribution (respectively, approximate Gaussian).

Figure 14. A supply chain.



Sample Distribution of Performance Ratio. As shown in Figure 15, the robust policy performs significantly better than the myopic policy, independently of the realized probability distributions, although there is a small positive probability that the myopic policy will perform better. This last point is not surprising given the short time horizon we are considering.

Costs of the Robust and the Myopic Policies. For clarity, and because the realized distribution does not seem to play a major role, in Figure 16 we only show the costs for gamma distributions at both sink nodes. The robust policy has lower mean and variance than the myopic policy when the assumed distribution is binomial or Gaussian.

Impact of the Horizon. We study here the impact of the time horizon T on the relative performance of the robust and myopic policies as measured by the ratio r for gamma demand distributions at both stations 1 and 2. The results presented in Figure 17 confirm the intuition that the robust policy tends to perform better as the horizon T increases. Although the peak of the relative performance r stays around 20% for a sample probability of about 0.3, the spread of the sample distribution seems to be reduced as T increases, thus making it more and more likely that the robust policy will outperform the myopic policy as the horizon increases.

5.4. Summary of Results

The numerical evidence we have presented in §5 suggests that:

- The robust approach leads to high-quality solutions and often outperforms dynamic programming-based policies in single stations and myopic policies in more complex supply chains.
- For single stations, the robust approach outperforms dynamic programming, when the holding cost is above a certain threshold (relative to the shortage cost) and is outperformed by dynamic programming otherwise. This conclusion remains valid for a wide variety of assumed and

Figure 15. Relative performance.

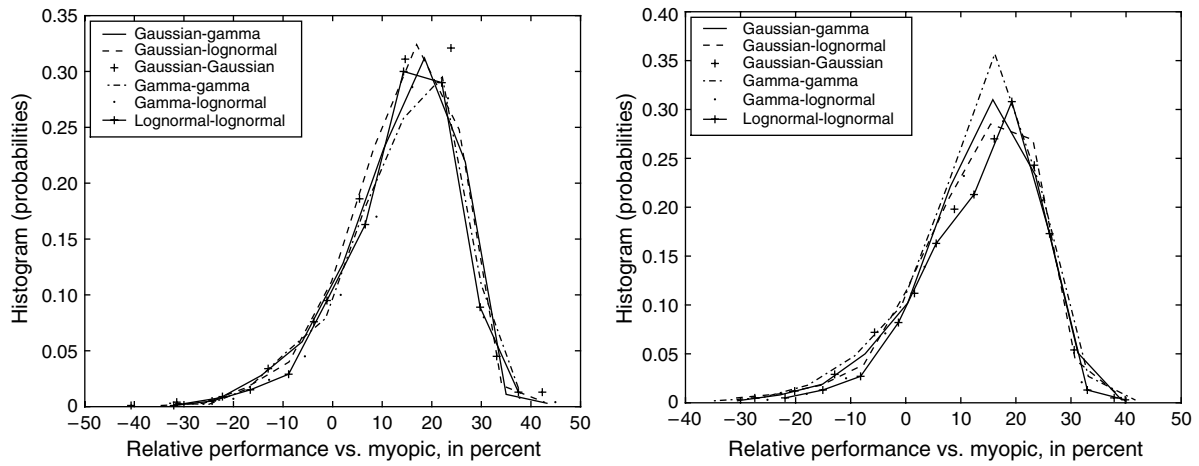


Figure 16. Sample probability distributions of costs.

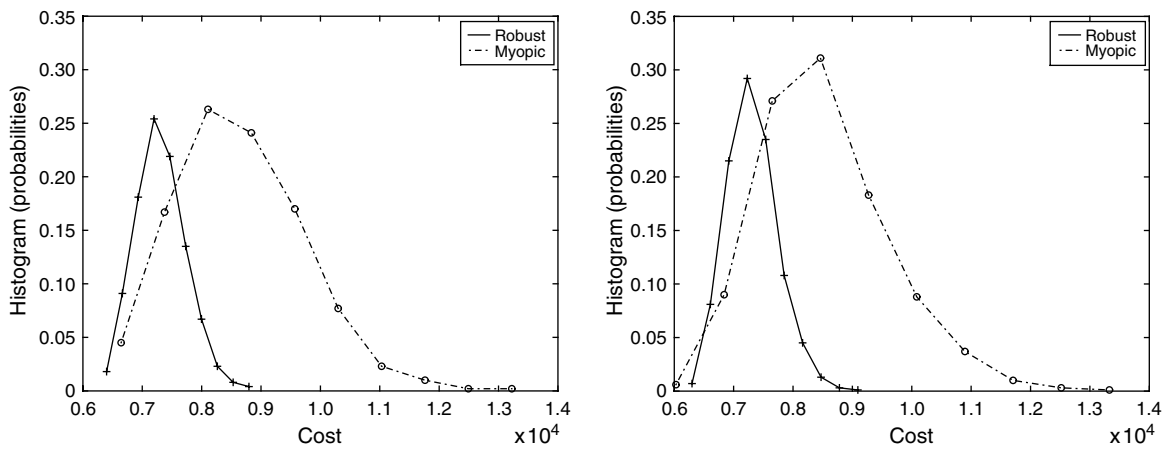
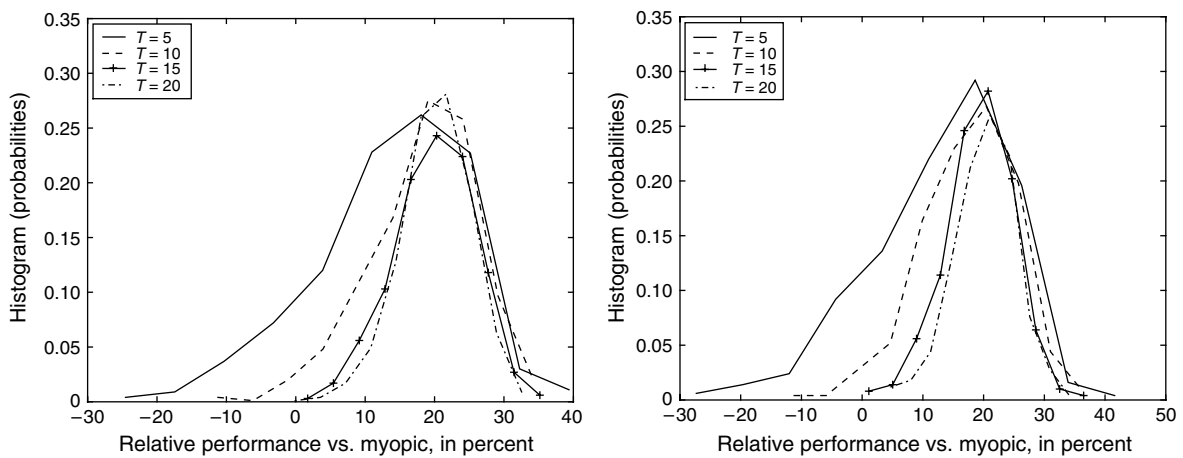


Figure 17. Impact of the horizon.



realized distributions as well as other parameters (like the ordering cost).

- For more complex supply chains, the robust policy performs significantly better than a myopic policy, in particular over many time periods, even when assumed and realized demand distributions are close.

6. Conclusions

In this paper, we have proposed a deterministic, numerically tractable methodology to address the problem of optimally controlling supply chains subject to random demand. Using robust optimization ideas, we have built an equivalent model without uncertainty of the same class as the nominal problem, with a modified demand sequence. Specifically, the proposed model is a linear programming problem if there are no fixed costs throughout the supply chain, and a mixed-integer programming problem if fixed costs are present.

This model incorporates a wide variety of phenomena, including demands that are not identically distributed over time and capacity on the echelons and links. When the parameters are chosen appropriately, the proposed approach preserves performance while protecting against uncertainty. One of its most appealing features is that it uses very little information on the demand distributions, and therefore is widely applicable. In particular, if we only know the mean and the variance of the distributions, the robust policy often outperforms the nominal policy, as well as policies computed assuming full but erroneous knowledge of the distributions for the correct mean and variance.

This approach also provides valuable theoretical insights. We have derived the expression of key parameters of the robust policy, and shown optimality of base-stock policies in the proposed framework when the optimal stochastic policy is base stock, but also in other instances where the optimal stochastic policy is not known. Hence, the methodology is not only validated on benchmark problems where the optimal stochastic policy is already known (yet in general hard to compute numerically), but also provides a framework to analyze complex supply chains for which the traditional tools of stochastic optimization face serious dimensionality problems.

Acknowledgments

The authors thank the associate editor and the reviewers of the paper for their comments, which improved the paper substantially. This research was partially supported by the Singapore-MIT Alliance.

References

- Ben-Tal, A., A. Nemirovski. 1998. Robust convex optimization. *Math. Oper. Res.* **23** 769–805.
- Ben-Tal, A., A. Nemirovski. 1999. Robust solutions of uncertain linear programs. *Oper. Res. Lett.* **25** 1–13.
- Ben-Tal, A., A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming Ser. A.* **88** 411–424.
- Bertsekas, D. 1995. *Dynamic Programming and Optimal Control*, Vol. 1. Athena Scientific, Belmont, MA.
- Bertsekas, D., J. N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA.
- Bertsimas, D., I. Popescu. 2002. On the relation between option and stock prices: A convex optimization approach. *Oper. Res.* **50** 358–374.
- Bertsimas, D., M. Sim. 2003. Robust discrete optimization and network flows. *Math. Programming Ser. B.* **98** 48–71.
- Bertsimas, D., M. Sim. 2004. The price of robustness. *Oper. Res.* **52** 35–53.
- Clark, A., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6**(4) 475–490.
- El Ghaoui, L., H. Lebret. 1997. Robust solutions to least-square problems to uncertain data matrices. *SIAM J. Matrix Anal. Appl.* **18** 1035–1064.
- El Ghaoui, L., F. Oustry, H. Lebret. 1998. Robust solutions to uncertain semidefinite programs. *SIAM J. Optim.* **9** 33–52.
- Fu, M. C. 1994. Sample path derivatives for (s, S) inventory systems. *Oper. Res.* **42** 351–364.
- Gallego, G., J. Ryan, D. Simchi-Levi. 2001. Minimax analysis for finite-horizon inventory models. *IIE Trans.* **33** 861–874.
- Glasserman, P. 1991. *Gradient Estimation via Perturbation Analysis*. Kluwer Academic Publishers, Boston, MA.
- Glasserman, P., S. Tayur. 1995. Sensitivity analysis for base stock levels in multi-echelon production-inventory systems. *Management Sci.* **41** 263–281.
- Ho, Y. C., X. R. Cao. 1991. *Discrete Event Dynamic Systems and Perturbation Analysis*. Kluwer Academic Publishers, Boston, MA.
- Kapuscinski, R., S. Tayur. 1999. Optimal policies and simulation based optimization for capacitated production inventory systems. S. Tayur, R. Ganeshan, M. J. Magazine, eds. Chapter 2 in *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, Boston, MA.
- Kushner, H., D. Clark. 1978. *Stochastic Approximation for Constrained and Unconstrained Systems*. Springer-Verlag, New York.
- Lo, A. 1987. Semiparametric upper bounds for option prices and expected payoffs. *J. Financial Econom.* **19** 373–388.
- Moon, I., G. Gallego. 1993. The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.* **44** 825–834.
- Moon, I., G. Gallego. 1994. Distribution free procedures for some inventory models. *J. Oper. Res. Soc.* **45** 651–658.
- Scarf, H. 1958. A min-max solution of an inventory problem. K. J. Arrow, S. Karlin, H. E. Scarf, eds. *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, Stanford, CA, 201–209.
- Soyster, A. L. 1973. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.* **21** 1154–1157.
- Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill Higher Education, Boston, MA.

Copyright 2006, by INFORMS, all rights reserved. Copyright of Operations Research is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.