

A Robust Optimization Approach to Supply Chain Management

Dimitris Bertsimas and Aurélie Thiele

Massachusetts Institute of Technology, Cambridge MA 02139,
dbertsim@mit.edu, aurelie@mit.edu

Abstract. 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. The attractive features of the proposed approach are: (a) It incorporates a wide variety of phenomena, including demands that are not identically distributed over time and capacity on the echelons and links; (b) it uses very little information on the demand distributions; (c) it leads to qualitatively similar optimal policies (basestock policies) as in dynamic programming; (d) it is numerically tractable for large scale supply chain problems even in networks, where dynamic programming methods face serious dimensionality problems; (e) in preliminary computational experiments, it often outperforms dynamic programming based solutions for a wide range of parameters.

1 Introduction

Optimal supply chain management has been extensively studied in the past using dynamic programming, which leads to insightful policies for simpler systems (basestock policies for series systems; Clark and Scarf [7]). Unfortunately, dynamic programming assumes complete knowledge of the probability distributions and suffers from the curse of dimensionality. As a result, preference for implementation purposes is given to more intuitive policies that are much easier to compute, but also suboptimal (see Zipkin [12]).

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 that has witnessed increased research activity (Soyster [11], Ben-Tal and Nemirovski ([1,2,3]) and El-Ghaoui et. al. ([8,9], Bertsimas and Sim [5,6]). We utilize the approach in [5,6], which leads to linear robust counterparts while controlling the level of conservativeness of the solution.

The contributions of this paper are as follows: (a) We develop an approach that incorporates demand uncertainty 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. (b) 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. (c) The optimal robust policy is qualitatively similar to the optimal policy obtained by dynamic programming when known. In particular, it remains basestock when the optimal stochastic policy is basestock, as well as in some other cases where the optimal stochastic policy is not known. (d) 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.

2 The Robust Optimization Approach

We rely extensively on the robust optimization tools developed by Bertsimas and Sim in [5] for linear programming problems. We consider the following problem subject to data uncertainty:

$$\min \mathbf{c}'\mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u},$$

where we assume WLOG that only the matrix \mathbf{A} is subject to data uncertainty.

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

Γ is a parameter that controls the degree of conservatism. The robust problem is then formulated as:

$$\begin{aligned} \min \mathbf{c}'\mathbf{x} \\ \text{s.t. } \mathbf{A}\mathbf{x} \leq \mathbf{b}, \quad \forall \mathbf{A} \in \mathcal{A} \\ \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}. \end{aligned} \quad (1)$$

Theorem 1 (Bertsimas and Sim [5]). *The uncertain linear programming problem has the following robust, linear counterpart:*

$$\begin{aligned} \min \mathbf{c}'\mathbf{x} \\ \text{s.t. } \sum_j \bar{a}_{ij}x_j + q_i\Gamma + \sum_{j:(i,j) \in J} r_{ij} \leq b_i, \forall i \\ q_i + r_{ij} \geq \hat{a}_{ij}y_j, \quad \forall (i,j) \in J \\ -\mathbf{y} \leq \mathbf{x} \leq \mathbf{y} \\ \mathbf{l} \leq \mathbf{x} \leq \mathbf{u} \\ \mathbf{q} \geq 0, \mathbf{r} \geq 0, \mathbf{y} \geq 0. \end{aligned} \quad (2)$$

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, since 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 define, for $k = 0, \dots, T$:

x_k : the stock available at the beginning of the k th period,

u_k : the stock ordered at the beginning of the k th period,

w_k : the 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 leadtime equal to 0. Excess demand is backlogged. Therefore, the evolution of the stock over time is described by the following linear equation:

$$x_{k+1} = x_k + u_k - w_k, \quad k = 0, \dots, T - 1, \quad (3)$$

leading to the closed-form expression:

$$x_{k+1} = x_0 + \sum_{i=0}^k (u_i - w_i), \quad k = 0, \dots, T - 1. \quad (4)$$

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

The demands w_k are random variables. In order to apply the approach outlined in Sect. 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 the 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 the 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 $R(x_{k+1})$. Here, we consider a purchasing cost of the form:

$$C(u) = \begin{cases} K + c \cdot u, & \text{if } u > 0, \\ 0, & \text{if } u = 0, \end{cases} \quad (5)$$

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:

$$R(x) = \max(hx, -px), \quad (6)$$

where h and p are nonnegative. The holding/shortage cost for period k , y_k , is computed at the end of the period, after the shipment u_k has been received and the demand w_k has been realized. We assume $p > c$, 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:

$$\begin{aligned} \min \quad & \sum_{k=0}^{T-1} (cu_k + Kv_k + y_k) \\ \text{s.t.} \quad & y_k \geq h \left(x_0 + \sum_{i=0}^k (u_i - w_i) \right) \quad k = 0, \dots, T-1 \\ & y_k \geq -p \left(x_0 + \sum_{i=0}^k (u_i - w_i) \right) \quad k = 0, \dots, T-1 \\ & 0 \leq u_k \leq Mv_k, \quad v_k \in \{0, 1\} \quad k = 0, \dots, T-1, \end{aligned} \quad (7)$$

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

Theorem 2. *The robust formulation for the single-station inventory problem (7) is:*

$$\begin{aligned} \min \quad & \sum_{k=0}^{T-1} (cu_k + Kv_k + y_k) \\ \text{s.t.} \quad & y_k \geq h \left(x_0 + \sum_{i=0}^k (u_i - \bar{w}_i) + q_k \Gamma_k + \sum_{i=0}^k r_{ik} \right) \\ & y_k \geq p \left(-x_0 - \sum_{i=0}^k (u_i - \bar{w}_i) + q_k \Gamma_k + \sum_{i=0}^k r_{ik} \right) \\ & q_k + r_{ik} \geq \hat{w}_i \\ & q_k \geq 0, \quad r_{ik} \geq 0 \\ & 0 \leq u_k \leq Mv_k, \quad v_k \in \{0, 1\}, \end{aligned} \quad (8)$$

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, specifically the level of conservativeness and the bounds of the uncertain variables. $q_k \Gamma_k + \sum_{i=0}^k r_{ik}$ represents the extra inventory (or lack thereof) that we want to take into account in controlling the system from a worst-case perspective.

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 readily be 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 1 ((S,S) and (s,S) policies). *The optimal policy of a discrete-horizon inventory problem is said to be (s, S) , or basestock, 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 ($K = 0$), $s_k = S_k$.*

In order to analyze the optimal robust policy, we need the following lemma:

Lemma 1. (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 for the nominal case 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 and s_j, S_j the corresponding thresholds at time t_j , we have:*

$$S_j = \sum_{i=0}^{I_j} \bar{w}_{t_j+i}, \quad (9)$$

and

$$s_1 = x_0 - \sum_{i=0}^{t_1-1} \bar{w}_i, \quad s_j = - \sum_{i=I_{j-1}+1}^{L_{j-1}-1} \bar{w}_{t_{j-1}+i}, \quad j \geq 2, \quad (10)$$

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

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

Theorem 3 (Optimal robust policy).

(a) *The optimal policy in the robust formulation (8), evaluated at time 0 for the rest of the horizon, is the optimal policy for the nominal problem with the modified demand:*

$$w'_k = \bar{w}_k + \frac{p-h}{p+h} (A_k - A_{k-1}), \quad (11)$$

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 , \mathbf{q}^* and \mathbf{r}^* being the optimal \mathbf{q} and \mathbf{r} variables in (8). (By convention $q_{-1} = r_{\cdot,-1} = 0$.) 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 (9) and (10) applied to the modified demand w'_k .

(d) The optimal cost of the robust problem (8) 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, $\frac{2ph}{p+h} \sum_{k=0}^{T-1} A_k$.

Proof. Let $(\mathbf{u}^*, \mathbf{v}^*, \mathbf{q}^*, \mathbf{r}^*)$ be the optimal solution of (8). Obviously, setting the \mathbf{q} and \mathbf{r} variables to their optimal values \mathbf{q}^* and \mathbf{r}^* in (8) 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 $\mathbf{q}^*, \mathbf{r}^*$ as given in the robust formulation (8). We have then to solve:

$$\min_{\mathbf{u} \geq 0} \sum_{k=0}^{T-1} [cu_k + K1_{\{u_k > 0\}} + \max(h(\bar{x}_{k+1} + A_k), p(-\bar{x}_{k+1} + A_k))] \quad (12)$$

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:

$$x'_{k+1} = x'_k + u_k - \underbrace{\left(\bar{w}_k + \frac{p-h}{p+h} (A_k - A_{k-1}) \right)}_{=w'_k}, \quad (13)$$

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

$$\max(h(\bar{x}_{k+1} + A_k), p(-\bar{x}_{k+1} + A_k)) = \max(hx'_{k+1}, -px'_{k+1}) + \frac{2ph}{p+h} A_k. \quad (14)$$

The reformulation of the robust model, given the optimal \mathbf{q}^* and \mathbf{r}^* variables, as a nominal inventory problem in the modified stock variable x'_k (plus the fixed cost $\frac{2ph}{p+h} \sum_{k=0}^{T-1} A_k$) follows from injecting (14) into (12). This proves (a) and (d). We conclude that (b) and (c) hold by invoking Lemma 1. \square

Remark: 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, $S_k = w'_k = \bar{w} + \frac{p-h}{p+h} \hat{w} (\Gamma_k - \Gamma_{k-1})$ 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 Sect. 3.1.

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

$$u_k \leq d, \quad \forall k, \quad (15)$$

to (8). We next study the structure of the optimal policy.

Theorem 4 (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 (11). As a result, the optimal policy remains (S, S) (resp (s, S)) in the case without (resp with) fixed cost.*

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 (8):

$$x_0 + \sum_{i=0}^k (u_i - w_i) \leq C, \quad (16)$$

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 Sect. 2, we rewrite this constraint in the robust framework as:

$$\bar{x}_{k+1} + q_k \Gamma_k + \sum_{i=0}^k r_{ik} \leq C, \quad \forall k, \quad (17)$$

where q_k and r_{ik} are defined in (8). We next study the optimal policy.

Theorem 5 (Optimal robust policy). *The optimal robust policy is the optimal policy for the nominal problem subject to the modified demand defined in (11), and with inventory capacity at time 0 equal to C , and inventory capacity at time $k + 1$, $k \geq 0$, equal to $C - \frac{2p}{p+h} A_k$.*

As a result, the optimal policy remains (S, S) (resp (s, S)) in the case without (resp with) fixed purchasing cost.

4 The Network Case

4.1 The Uncapacitated Model

We now extend the results of Sect. 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 sink node, 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. 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 [7]. 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 $k = 1, \dots, N$:

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

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

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

$W_s(t)$: the 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 leadtimes equal to 0, backlog of excess demand, and linear dynamics for the stock at installation k at time $t = 0, \dots, T - 1$:

$$I_k(t + 1) = I_k(t) + D_{i_k k}(t) - \sum_{j \in N(k)} D_{kj}(t), \quad (18)$$

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 at time $t = 0, \dots, T - 1$:

$$X_k(t + 1) = X_k(t) + D_{i_k k}(t) - \sum_{s \in O(k)} W_s(t). \quad (19)$$

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

$$\sum_{i \in N(k)} D_{ki}(t) \leq \max(I_k(t), 0), \quad \forall k, \quad \forall t, \quad (20)$$

that is, the total order made to a supplier cannot exceed what the supplier has currently in stock, or, equivalently, the supplier can only send through the network items that it really has. Since 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:

$$\sum_{i \in N(k)} D_{ki}(t) \leq \bar{X}_k(t) - \sum_{i \in N(k)} \bar{X}_i(t), \quad \forall k, \forall t. \quad (21)$$

Finally, we specify the cost function. We assume that each echelon k has the same cost structure as the single installation modelled in Sect. 3.1 with specific parameters (c_k, K_k, h_k, p_k) . We also keep here the same notations and assumptions as in Sect. 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 $0 \leq \Gamma_s(t) - \Gamma_s(t-1) \leq 1$ 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 Sect. 2 to the holding/shortage constraints in the same manner as in Sect. 3, we obtain the mixed integer programming problem:

$$\begin{aligned} \min & \sum_{t=0}^{T-1} \sum_{k=1}^N \sum_{i \in N(k)} \{c_{ki}D_{ki}(t) + K_{ki}V_{ki}(t) + Y_i(t)\} \\ \text{s.t.} & Y_i(t) \geq h_i \left\{ \bar{X}_i(t+1) + \sum_{s \in O(i)} \left(q_s(t)\Gamma_s(t) + \sum_{\tau=0}^t r_s(\tau, t) \right) \right\}, \\ & Y_i(t) \geq p_i \left\{ -\bar{X}_i(t+1) + \sum_{s \in O(i)} \left(q_s(t)\Gamma_s(t) + \sum_{\tau=0}^t r_s(\tau, t) \right) \right\}, \\ & \sum_{i \in N(k)} D_{ki}(t) \leq \bar{X}_k(t) - \sum_{i \in N(k)} \bar{X}_i(t), \\ & q_s(t) + r_s(\tau, t) \geq \widehat{W}_s(\tau), \\ & q_s(t) \geq 0, r_s(\tau, t) \geq 0, \\ & 0 \leq D_{ki}(t) \leq MV_{ki}(t), V_{ki}(t) \in \{0, 1\}, \end{aligned} \quad (22)$$

with $\bar{X}_i(t+1) = X_i(0) + \sum_{\tau=0}^t \left\{ D_{ki}(\tau) - \sum_{s \in O(i)} \bar{W}_s(\tau) \right\}$ for all i and t , where k supplies i .

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. The main result is as follows.

Theorem 6 (Optimal robust policy).

(a) The optimal policy in (22) for echelon k is the optimal policy obtained for the supply chain subject to the modified, deterministic demand at sink node s (for $s \in O(k)$):

$$W'_{s,k}(t) = \overline{W}_s(t) + \frac{p_k - h_k}{p_k + h_k} (A_s(t) - A_s(t - 1)), \quad (23)$$

where $A_s(t) = q_s^*(t)\Gamma_s(t) + \sum_{\tau=0}^t r_s^*(\tau, t)$, \mathbf{q}_s^* and \mathbf{r}_s^* being the optimal \mathbf{q} and \mathbf{r} variables associated with sink node s in (22).

(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,

$$\sum_{k=1}^N \frac{2p_k h_k}{p_k + h_k} \sum_{t=0}^{T-1} \sum_{s \in O(k)} A_s(t).$$

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:

$$X_k(t + 1) = X_k(0) + \sum_{\tau=0}^t \left\{ \sum_{i \in I(k)} D_{ik}(\tau) - \sum_{j \in O(k)} W_j(\tau) \right\}. \quad (24)$$

With the standard cost assumptions used before, the echelons cannot be studied independently and the optimal policy is not necessarily basestock, even in the simple case of demand without uncertainty. This is illustrated by the following example.

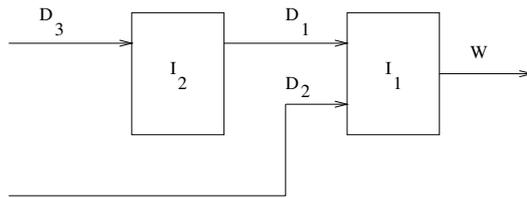


Fig. 1. A network for which the optimal policy is not basestock.

The network in Fig. 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 1 time period, and the demand at time 0 is deterministic, equal to 10 units. Echelon 1 has only 5 units in inventory at time 0.

Comparing the two options, it is easy to see that it is optimal for echelon 1 to order 5 units from 2 if 2 has 5 units in inventory at time 0, 2 units from 2 and none from the plant if 2 has 2 units, and 5 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, resp. 7, 10, units if installation 2 has 5, resp 2, 0 units in inventory at time 0. Thus, the optimal policy is not basestock.

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 since 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 Sect. 3.2, an upper bound on the maximal order can be directly introduced in the proposed approach, by adding the constraint:

$$D_{ki}(t) \leq d_{ki} \quad \forall k, \quad \forall i \in \mathcal{N}(k), \quad \forall t, \quad (25)$$

to (22). Inventory capacity, however, requires further manipulation, since the level of inventory held at an echelon at any time depends on the demand, which is subject to uncertainty. Similar manipulations as in Sect. 3.2 lead to the constraint $\forall k, \quad \forall t$:

$$\bar{X}_k(t+1) + \sum_{s \in \mathcal{O}(k)} \left(q_s(t) \Gamma_s(t) + \sum_{\tau=0}^t r_s(\tau, t) \right) \leq C_k \quad (26)$$

to be added to the formulation, $q(t)$ and $r(\tau, t)$ being defined as in (23).

We next study the structure of the optimal policy.

Theorem 7. *The optimal policy at each echelon remains basestock in presence of link and echelon capacities. It is identical to the optimal policy of a nominal problem at a single station subject to the modified demand defined in (23), 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 that incorporate $D_{ki}(t) \leq d_{ki}$ for all $k, i \in \mathcal{N}(k)$ and t , as well as the capacity induced by the coupling constraint (21).

5 Numerical Implementation

We now apply the proposed methodology to the example of minimizing the cost at a single station. The horizon is $T = 10$ time periods, the initial inventory is $x_0 = 150$, with an ordering cost per unit $c = 1$, a holding cost $h = 2$ and a shortage cost $p = 3$, in appropriate measurement units. There is no fixed ordering cost. The stochastic demand is i.i.d. with mean $\bar{w} = 100$. In the robust framework, we consider that the demand belongs to the interval $[0, 200]$, that is $\hat{w} = 100$. We compare the expected costs of the robust policy and of the stochastic policy obtained by dynamic programming as a function of the standard deviation σ of the distribution.

To select the budgets of uncertainty we minimize a tight upper bound on the expected cost of the system over the set of nonnegative distributions of given mean and variance, using the results in [4]. This yields for all k :

$$\Gamma_k = \min \left(\frac{\sigma}{\hat{w}} \sqrt{\frac{k+1}{1-\alpha^2}}, k+1 \right), \quad (27)$$

and the modified demand at time k is in this example:

$$w'_k = \bar{w} + \frac{\alpha \sigma}{\sqrt{1-\alpha^2}} (\sqrt{k+1} - \sqrt{k}), \quad (28)$$

with $\alpha = \frac{p-h}{p+h}$. Expected costs are computed using the mean of a sample of size 1,000.

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 seven points $(\bar{w} - 3\sigma, \bar{w} - 2\sigma, \dots, \bar{w} + 2\sigma, \bar{w} + 3\sigma)$. In both cases, the actual distribution is Gamma, Lognormal or Gaussian, with the same mean \bar{w} and standard deviation σ . The impact of the mistake on the demand distribution is measured by the ratio $(DP - ROB)/DP$, with DP , resp. ROB , the expected cost obtained using dynamic programming, resp. the robust approach. The results are shown in Fig. 2. In the first case, where the distributions are very different beyond their first moments, the impact of the ratio increases as the standard deviation increases and the robust policy outperforms dynamic programming by up to 8%. In the second case, the two methods are equivalent in terms of performance, since the robust policy outperforms dynamic programming by at most 0.3%, which is not statistically significant.

In Figs. 3-5, we study the impact of the cost parameters c , h and p in the settings described above, where we vary one parameter at a time. The impact

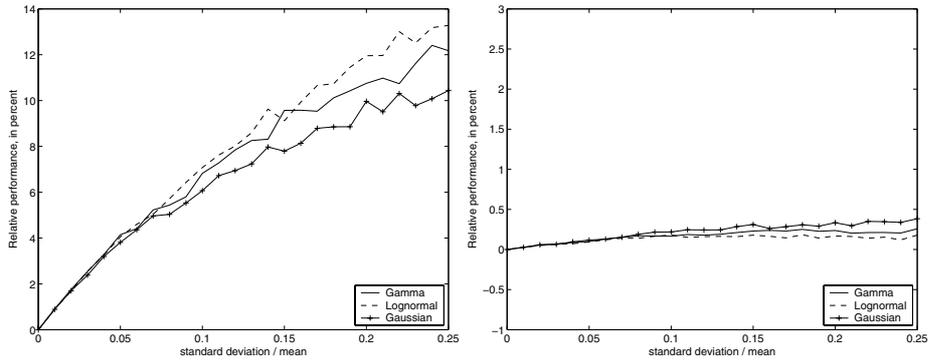


Fig. 2. Impact of the standard deviation on performance.

of a change in the parameters is qualitatively similar in both cases, with little dependence on the actual distribution of the demand. The robust approach outperforms the stochastic policy for a wide range of parameters, although the stochastic policy leads to better results for large values of the ratio p/h (greater than about 3). The exact numbers depend on the distribution used to compute the stochastic policy.

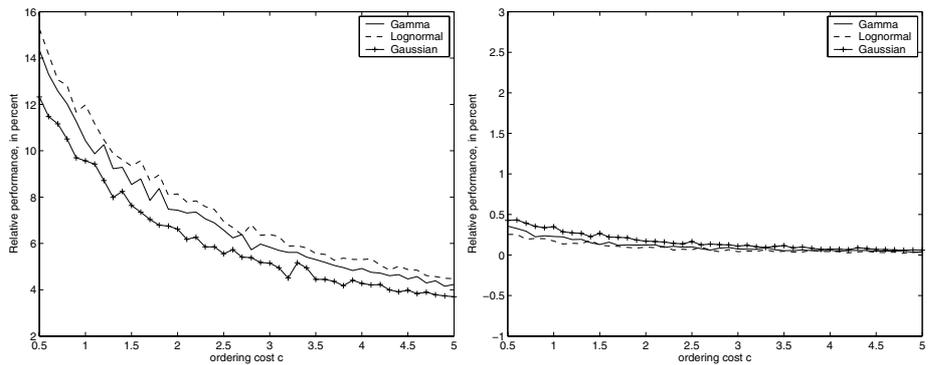


Fig. 3. Impact of the ordering cost.

Overall the numerical evidence suggests that the robust policy performs significantly better than dynamic programming when assumed and actual distributions differ widely despite having the same mean and standard deviation, and performs similarly to dynamic programming when assumed and actual distributions are close. The results are thus quite promising.

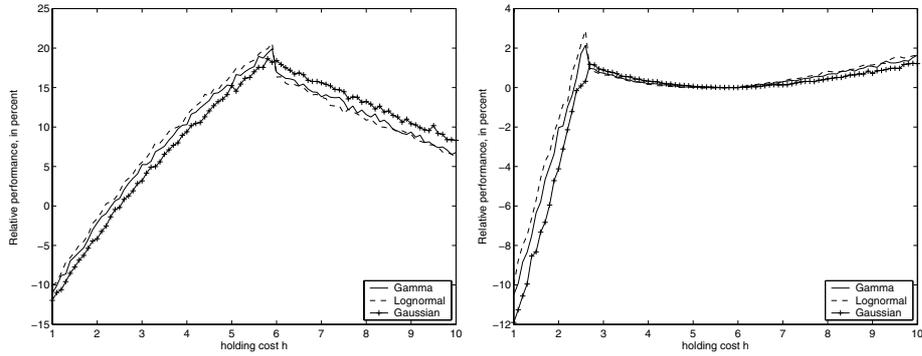


Fig. 4. Impact of the holding cost.

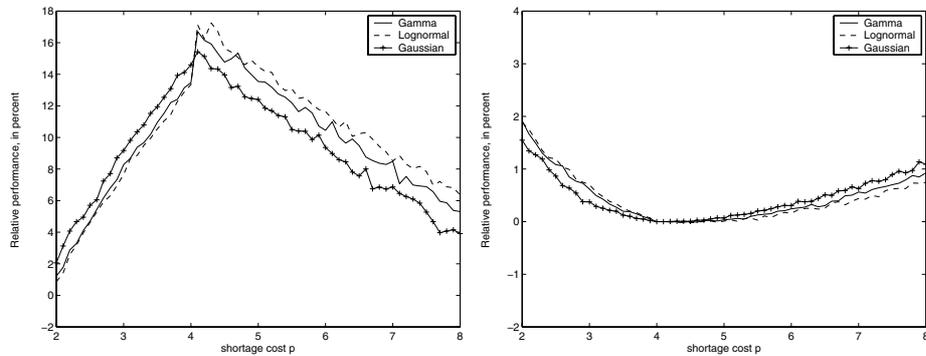


Fig. 5. Impact of the shortage cost.

6 Conclusions

We have proposed a deterministic, numerically tractable methodology to address the problem of optimally controlling supply chains subject to uncertain 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.

The key attractive features of the proposed approach are: (a) It incorporates a wide variety of phenomena, including demands that are not identically distributed over time and capacity on the echelons and links; (b) it uses very

little information on the demand distributions; (c) it leads to qualitatively similar optimal policies (basestock policies) as in dynamic programming; (d) it is numerically tractable for large scale supply chain problems even in networks, where dynamic programming methods face serious dimensionality problems; (e) in preliminary computational experiments, it often outperforms dynamic programming based solutions for a wide range of parameters.

References

1. Ben-Tal, A., Nemirovski, A. (2000): Robust solutions of linear programming problems contaminated with uncertain data, *Math. Program., Ser. A* 88:411-424.
2. Ben-Tal, A., Nemirovski, A. (1999): Robust solutions of uncertain linear programs, *Oper. Research Lett.* 25, 1-13.
3. Ben-Tal, A., Nemirovski, A. (1998): Robust convex optimization, *Math. Oper. Res.* 23, 769-805.
4. Bertsimas, D., Popescu, I. (2002): On the relation between option and stock prices: a convex optimization approach, *Oper. Res.* 50. 358-374.
5. Bertsimas, D., Sim, M. (2003): The price of robustness, to appear in *Operations Research*.
6. Bertsimas, D., Sim, M. (2003): Robust discrete optimization and network flows, *Mathematical Programming*, Vol. 98, No. 1-3, 49-71.
7. Clark, A., Scarf, H. (1960): Optimal policies for a multi-echelon inventory problem, *Management Science* Vol.6 No.4, 475-490.
8. El-Ghaoui, L., Lebret, H. (1997): Robust solutions to least-square problems to uncertain data matrices, *SIAM J. Matrix Anal. Appl.* 18, 1035-1064.
9. El-Ghaoui, L., Oustry, F., Lebret, H. (1998): Robust solutions to uncertain semidefinite programs, *SIAM J. Optim.* 9, 33-52.
10. Scarf, H. (1958): A min-max solution of an inventory problem, *Studies in the mathematical theory of inventory and production*, 201-209.
11. Soyster, A.L. (1973): Convex programming with set-inclusive constraints and applications to inexact linear programming, *Oper. Res.* 21. 1154-1157.
12. Zipkin, P. (2000): Foundations of inventory management, McGraw-Hill Higher Education.