Supplier-Retailer Flexible Commitments Contracts : A Robust Optimization Approach

Tal, Aharon Ben; Golany, Boaz; Nemirovski, Arcadi; Vial, Jean-Philippe

Supplier-Retailer Flexible Commitments Contracts: A Robust Optimization Approach

Aharon Ben Tal†
Boaz Golany†
Arcadi Nemirovski†
Jean-Philippe Vial‡

March 4, 2003

†Faculty of Industrial Engineering and Management
Technion—Israel Institute of Technology
Haifa 32000, Israel

‡Department of Management Studies
University of Geneva, 102 Bd Carl Vogt
CH-1211 Geneva 4, Switzerland

Abstract

We propose the use of robust optimization, a powerful tool that has already proven itself in several application areas, to the problem of finding optimal values for quantity flexible commitment contracts between suppliers and retailers. We demonstrate the capabilities of the robust optimization methodology by solving model variations that could not be solved by other methods previously suggested in this area. In particular, we allow non-stationary demand and cost parameters and a more flexible way to describe the possible deviations in commitments and orders. We apply our models first to finite-horizon problems and then extend it to a new type of rolling-horizon settings which we denote as "folding-horizons".

Keywords: Flexible commitments, supply chain management, robust optimization, folding horizons

1 Introduction

We address retailer-supplier problems in which the supplier produces and ships items to a retailer who orders them to fulfill the demand of his end-customers. The demand is uncertain and this uncertainty is a source of contention for both the retailer and the supplier. Hence, in conventional settings, each side is trying to pass on as much of the uncertainty burden as possible to the the other side. From the retailer's perspective, it is desired to hold as little inventory as possible, order in each period and adjust the orders up or down according to the customers' demand. However, shortage costs force him to keep some inventory and order costs often cause him not to order in each period. When this happens, the situation worsens since the variability increases as information flows backwards in the chain leading to the well-documented "bullwhip" effect (see, e.g., Lee et al., 1996). The supplier naturally wishes the opposite – he would like to see the retailer buys the material in advance and keep a large inventory to meet the external demand. One of the traditional tools that suppliers employed to entice retailers in this direction is to offer quantity discounts (Monahan 1984).

In recent years several studies have pointed out the potential benefit to both suppliers and retailers from implementing a new approach, known as "flexible commitments". The purpose of a flexible commitment agreement between a retailer and a supplier is to assist both of them to face the uncertainty that is associated with the external demand. It is assumed that the retailer's position in the chain enables him to better forecast the future demand by the customers. Therefore, the retailer is expected to help the supplier by providing him with advanced information in the form of future commitments. These are estimates of his future orders for a given number of future periods. Now, given a vector of cost and penalty parameters that reflect the flexible arrangement that is proposed by a supplier and facing uncertain demand, the retailer's challenge is to find an optimal ordering policy. We shall refer to these problems as Retailer-Supplier Flexible Contracts (RSFC) models.

The literature on flexible commitment contracts, which we review in the next section, led to a number of interesting model variations. However, solving these models turn out to be a rather challenging task. In most cases, the models were presented as dynamic programming constructs and as such they suffer from the "curse of dimensionality" which prevents one from solving them in reasonable time. Therefore, with the exception of some special cases (which are not very practical), the only way to approach these problems is through special heuristics.

The main purpose of this paper is to present the methodology of Robust Optimization (RO) to the Operations Management community and demonstrate its potential benefits by implementing it to solve the complex problems of flexible commitment contracts that were proven to be difficult (if not impossible) to solve by other techniques. In particular, problems with large number of decision variables, uncertain parameters and many decision periods.

RO is a modelling methodology, combined with computational techniques, aimed at solving large-scale convex optimization problems, for which the data (parameters) are uncertain, and are only known to vary within a given *uncertainty set* $\mathcal{U}$. RO was successfully applied to some large scale and highly complex engineering design problems (Ben-Tal and Nemirovsky, 1997, Ben-Tal et al., 1999) and is gradually taking a place in optimization similar to the role of Robust Control in Control Theory. For a survey of the main results of RO, the reader is referred to Ben-Tal and Nemirovsky (2002) and references therein.

For Linear Programming (LP) problems, the RO methodology is particularly powerful; it allows uncertainty in any part of the LP data (the activity matrix, the cost coefficient and the right-hand side vector) and it can accommodate a rich variety of uncertainty sets, while keeping the ensuing mathematical program computationally tractable. Recently, RO was extended to deal with multi-period decision problems (Ben-Tal et al., 2002), such as the RSFC problem we study in this paper.

The rest of the paper is organized as follows. In §2 we provide a detailed literature review on flexible commitment contracts. In §3 we formulate the basic RSFC model (which is a fixed horizon problem) and transform it into a linear programming model. In §4 we introduce the RO principles and then implement them by constructing a robust counterpart to our problem. We demonstrate the outcomes of various simulation experiments with this model in §4. Then, in §5 we define an extension to the traditional rolling-horizon models, to be denoted as "folding horizons", and formulate its RO representation. Later, in §6 we report on a set of simulation experiments that were run with the folding horizon formulation. In §7 we discuss, without going into full detail, how our models can be extended to incorporate additional features (such as random yield in the fulfillment of the retailer's orders by the supplier). Finally, §8 offers conclusions and directions for future research.

2 Flexible Contracts: Literature Review

The area of flexible commitments was introduced by Anupindi and Bassok in a couple of working papers in 1995 and later in a series of publications from 1997 onwards. In Bassok and Anupindi (1997) the model specifies that the buyer must purchase at least a given total quantity of a single product over a finite time horizon. Additional volume is also available at the same price, beyond which a higher price is charged. The authors propose a dynamic programming methodology to derive the purchase policy that minimizes the buyer's total costs under stationary and random demand. One of the model variations in Anupindi and Bassok (1999) allows the buyer to submit to the supplier initial forecasts for his period-by-period purchases over a T-period horizon. Then, he may revise each period's purchase one time within specified percentage bounds. For this scenario, the authors propose a heuristic policy and discuss its effectiveness.

The pioneering work of Anupindi and Bassok was joined by Tsay (1995) who assumed that the buyer first estimates a purchase quantity for a given selling season and then the supplier commits to production. Finally, the buyer makes his actual purchase in light of

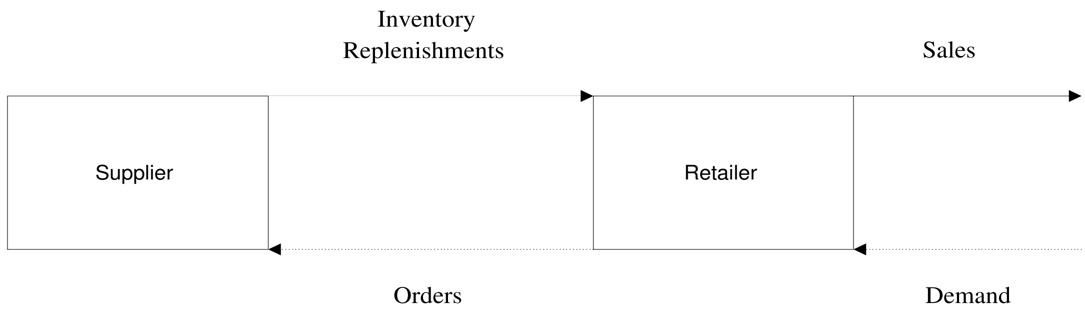

Figure 1: The Two-Stage System

updated information. He develops a contract which couples the buyer's commitment to purchase no less than a certain percentage below the initial estimate with the supplier's guarantee to deliver up to a certain percentage above the commitment. Tsay and Lovejoy (1999) studied a rolling-horizon contract in a multiple echelon setting, allowing for non-stationary demand with information updating and developed heuristics based on Open-Loop-Feedback-Control logic to solve it.

More recently, Urban (2000) offered another solution for the problem and suggested extensions to multi-product, multi-constraint models and Chen and Krass (2001) developed a total order quantity commitment model and studied its performance under various conditions.

All of the papers reviewed above suffer from the "curse of dimensionality" that is common to stochastic dynamic programming models. Therefore, the solutions they offer are either near optimal solution for small problems or heuristic solutions for larger problems. This is reflected in the words of Tsay and Lovejoy (1999, p. 99):

"Even with IID market demand and a G() of simple structure, (S-OLFC) is difficult to solve analytically due to the dimensionality and the constraint structure. Instead we have considered a number of computationally attractive, heuristic approaches based on relaxations of (S-OLFC)..."

3 Formulation of the RSFC Problem

3.1 Preliminaries

We assume a two-stage chain where a retailer who faces external demand for a single product orders from a unique supplier as depicted in the Figure 1:

This basic chain can be extended to multiple stages but this will not have a fundamental impact on the approach we propose here.

All the flows in the system (both material and information flows) are managed periodically. That is, at the end of each period (say, a month), the customers present their demand to the retailer who then determines (given his end-of-period inventory and other relevant information) his sales and order quantities for that period. The supplier receives the retailer's orders and treats them in the same manner as the retailer treated the customers' demand. This scenario is repeated every period.

The demand in every period is uncertain. Our approach enables us to model this uncertainty in various ways (e.g., box uncertainty, ellipsoid uncertainty) which we elaborate later on. All the other data parameters are assumed to be known in advance.

We first model these supplier-retailer relations for a fixed horizon scenario and then extend it to the more realistic case of rolling horizon orders and commitments.

In the fixed horizon scenario, the retailer-supplier relations follow a cyclic pattern. Once every T periods the retailer solves a model to generate a vector of T commitments and decision policy regarding his orders during these periods. The commitments, one for each period, are estimates for the future orders whose purpose is to enable the supplier to make the necessary preparations for production (e.g., order raw materials, secure the necessary capacity of labor and machines, etc.). The model also yields a firm order for the first period in the T periods horizon along with a decision rule (a policy) which the retailer will employ to compute the orders for future periods during that horizon (he is required to submit a firm order at the beginning of each period). Finally, the model also provides estimates of the inventory positions at the beginning of each period in the horizon. The actual inventories will be determined through an inventory balance equation after the demand is realized and a decision on the order quantity is placed.

From the supplier's perspective, two phenomena are desired:

- Small variability in the commitments' vector (which means better resource levelling for the supplier).
- No deviations between actual orders and earlier commitments.

To encourage the retailer to provide him with accurate and reliable advanced information, the supplier offers a favorable purchase price per-unit to the retailer. However, this is accompanied by two types of penalties (corresponding to the two objectives just mentioned):

- Penalties on upper and lower deviations of successive future commitments.
- Penalties on upper and lower deviations between actual orders and earlier commitments.

3.2 Initial model formulation

The model we start with is given below:

$$
(1) \min_{E,x,q,w} \left\{ E : \begin{array}{rrcl} (a): & E & \geq & sx_{T+1} + \sum\limits_{t=1}^{T} \Big[c_t q_t + h_t \max[x_{t+1}, 0] + p_t \max[-x_{t+1}, 0] \\ & & & +r_t^+ \max[q_t - w_t, 0] + r_t^- \max[w_t - q_t, 0] \\ & & & +v_t^+ \max[w_t - w_{t-1}, 0] + v_t^- \max[w_{t-1} - w_t, 0] \Big] \\ (b): & x_{t+1} & = & x_t + q_t - d_t \\ (c): & L_t & \leq & q_t \leq U_t \end{array} \right\}
$$

The notation we use here is:

- The quantities to be found are:
  - the total inventory management cost E;
  - the commitments w_t, $1 \leq t \leq T$ made at the beginning of the horizon ($t = 0$).
  - the orders q_t which the retailer orders from the supplier;
  - the states x_{t+1}, $t = 1, \ldots, T$ describing the state of the inventory at time $t + 1$.
- The certain data are:
  - the nonnegative coefficients c_t (purchasing costs per unit), h_t (holding costs per unit), p_t (shortage costs per unit of unsatisfied demand), s (salvage value per unit)[1];
  - the nonnegative coefficients r_t^+, r_t^- reflecting penalties on deviations between orders and commitments and v_t^+, v_t^- reflecting penalties on deviations between successive commitments;
  - the nonnegative lower (L_t) and upper ($U_t \geq L_t$) bounds on the retailers' orders;
  - the initial state of the inventory x_1.
- The uncertain data are the demands d_t.

We now turn to convert (1) into a linear programming problem. The conversion is done by adding auxiliary variables that eliminate piece-wise linear dependencies in the original formulation. This results in the LP model:

[1]The role of this parameter is to push the model towards holding some inventory at the end of the fixed horizon. To make sense in our settings, s must satisfy $(c_T - s) + h_T < p_T$. For a thorough discussion of ending inventory valuation the reader is referred to Fisher et al, (2001).

RSFC Model

$$(2)\quad \min_{x,q,w,y,z,E}\left\{E: \begin{array}{llllllll} (a): & E & \geq & sx_{T+1}+\sum_{t=1}^{T}[c_tq_t+y_t+u_t+z_t] & & & & \\ (b): & x_{t+1} & = & x_t+q_t-d_t & (c): & L_t & \leq & q_t\leq U_t \\ (d_1): & y_t & \geq & h_tx_{t+1} & (d_2): & y_t & \geq & -p_tx_{t+1} \\ (e_1): & u_t & \geq & r_t^+(q_t-w_t) & (e_2): & u_t & \geq & -r_t^-(q_t-w_t) \\ (f_1): & z_t & \geq & v_t^+(w_t-w_{t-1}) & (f_2): & z_t & \geq & -v_t^-(w_t-w_{t-1}) \end{array}\right\}$$

where the new variables y_t, u_t, z_t are upper bounds on the piece-wise linear components in $(1.a)$: y_t represents the holding and shortage costs; u_t represents the forecast reliability penalties; z_t represents the commitments consistency penalties. The relevant components are expressed equivalently by the respective constraints $(2.d-f)$.

4 Robust Optimization Formulation of the RSFC Problem

4.1 The general RO methodology

The RSFC problem (2) is an example of an *uncertain LP*, that is, a family of deterministic LPs:

$$(3)\qquad \left\{\min_x \left\{c^Tx : Ax\leq d\right\}\right\}_{\xi\equiv(A,c,d)\in\mathcal{U}}$$

where $\mathcal{U}$ is a given *uncertainty set*, in which the data (A,c,d) may vary. The RO methodology associates with the family (3) a single deterministic problem—its Robust Counterpart (RC):

RC Model

$$(4)\qquad \min_{x,E}\left\{E : c^Tx\leq E,\ Ax\leq d,\ \forall\,\xi\equiv(A,c,d)\in\mathcal{U}\right\}.$$

A solution (x^*,E^*) of (4) is called *robust*. Such a solution satisfies the constraints for all possible realization of the data in $\mathcal{U}$, and guarantees an optimal objective function value not worse than E^*.

Problem (4), being a semi-infinite LP seems computationally intractable; yet, for a wide variety of compact, convex uncertainty sets $\mathcal{U}$, it turns out that the RC model is a tractable (polynomially solvable) convex mathematical problem, typically an LP or a conic quadratic problem (see Ben-Tal and Nemirovski, 1998, 2002).

The formulation of the RC model implicitly implies that the entire decision vector x is to be determined before actual realization of the uncertain data (A,c,d) occurs. This is certainly *not* the case for the RSFC problem (as well as other multi-period decision problems). Indeed, only some of the decision variables are actual "here and now" decisions; In the RSFC problem, the latter are: q_1—the first period order and the commitments $w_1, w_2, \dots, w_T$. Other variables are "wait and see" decisions, i.e., they can be adjusted as part of the uncertain data become known. In the RSFC problem these are the orders $q_2, \dots, q_T$; indeed, for $t > 1$, q_t may depend on the past demands $d_1, d_2, \dots, d_{t-1}$. For uncertain LPs with the above features, a better representation is

$$\min_{x,w,E} \left\{ E \mid a^T x + b^T w \le E,\ \ Ax + Bw \le d \right\}_{\xi=(a,b,A,B,d)\in\mathcal{U}} \tag{5}$$

where x corresponds to the *adjustable* ("wait and see") variables, and w to the *nonadjustable* ("here and now") ones.

For uncertain LPs of the type (5), a recent extension of the RC approach was developed (see [ ]), leading to an Adjustable Robust Counterpart (ARC) model:

<u>ARC Model</u>

$$\min_{w,E} \left\{ E \mid \forall\, \xi \in \mathcal{U} : \exists\, x : a^T x + b^T w \le E,\ \ Ax + Bw \le d \right\}. \tag{6}$$

The ARC model is clearly more flexible (less conservative) than the RC model; it has a larger feasible set, thus allowing a better optimal value, while still satisfying the constraints for all possible realization of the data $\xi = (a, b, A, B, d)$ varying in the uncertainty set $\mathcal{U}$. There is, however, a price to pay for this added flexibility: Problem (6) is often computationally intractable (NP-hard) even for simple uncertainty sets (e.g., a general polyhedral set).

The remedy for the above situation is offered in [ ] by restricting the dependence of the vector x of adjustable variables on the uncertainty vector ξ to be affine; i.e.:

$$x = L\xi + \lambda\,. \tag{7}$$

for some matrix L and vector λ of appropriate dimensions.

This idea of focusing on *linear decision rules* is quite common in many branches of science and engineering (linear controllers or linear feedback rules in control theory, linear estimators/predictors in signal processing, etc.). With the linear decision rule (7), the ARC model becomes the Affinely Adjustable RC (AARC) model:

<u>AARC Model</u>

$$\begin{array}{l} \min\limits_{w,E,L,\lambda} E \\ \text{subject to} \\ \left.\begin{array}{ll} (i) & a^T(L\xi+\lambda) + b^T w \le E \\ (ii) & A(L\xi+\lambda) + Bw \le d \end{array}\right\} \forall\, \xi \in \mathcal{U} \end{array} \tag{8}$$

An important special case of the AARC model is when the parameters (a, A), corresponding to the adjustable variables x in the uncertain LP (5), are *not* uncertain. Such uncertain LPs (which include our RSFC problem) are said to be with *fixed recourse*; for them, AARC depends affinely on the remaining uncertain parameters (b, B) and its mathematical structure is similar to the usual RC model. As a result, AARC is computationally tractable for a wide class of uncertainty sets $\mathcal{U}$ (see Ben-Tal et al., 2002).

4.2 Applying the AARC model to the RSFC problem

The first step in converting the RSFC problem to its AARC formulation is to determine the specific form of the linear decision rules for the adjustable variables $q_t (t > 1)$. Here we set

$$(9) \qquad \begin{array}{ll} (i) & q_t = q_t^0 + \sum_{\tau=1}^{t-1} q_t^\tau d_\tau \\ (ii) & q_t^\tau = 0 \quad \text{if} \ \ (\tau, t) \in J \end{array}$$

where J denotes the pairs (τ, t), $1 \leq \tau < t$ such that q_t does not depend on the demand d_τ at period τ. (J may include, for example, periods τ which are too distant in the past.)

With q_t being affine functions of the demand d_τ, $\tau < t$, the constraints (2.b) enforce the state variables x_{t+1} to become affine function of the d_τ's:

$$(10) \qquad x_{t+1} = x_{t+1}^0 + \sum_{\tau=1}^{t} x_{t+1} d_\tau,$$

with the same effect on the auxiliary variables, y_t, u_t:

$$(11) \qquad \begin{array}{rcl} y_t & = & y_t^0 + \sum_{\tau=1}^{t-1} y_t^\tau d_\tau \\ u_t & = & u_t^0 + \sum_{\tau=1}^{t-1} u_t^\tau d_\tau. \end{array}$$

The variables z_t, on the other hand, are affected only by the nonadjustable variables w_t and so they are not substituted by an affine function of the demands.

The second step in building the AARC formulation is the selection of the uncertainty set for each of the constraints of problem (2) involving the uncertain demands; these are the constraints (2.a), (2.b), (2.c), (2.d) and (2.e).

The constraint (2.a) corresponds to the objective function of our problem ("minimum cost") and is thus different in nature from the "real constraints" (2.b)–(2.e). For this constraint, we choose an *ellipsoidal* uncertainty set, specifically

$$(12) \qquad \mathcal{U}_1 = \left\{ d \in \mathbb{R}^T : (d - d^n) S^{-1} (d - d^n) \leq \rho^2 \right\}$$

where d^n is some fixed "nominal demand", S is a $T \times T$ symmetric positive definite matrix, and $\rho > 0$ is the "safety parameter". Natural choices of d^n and S, when the demand vector is stochastic are

$$d^n = E(d), \qquad S = \mathrm{cov}(d)$$

(see [,]). For each of the remaining constraints (2.b)–(2.e), we choose the *box* uncertainty set:

$$\mathcal{U}_2 = \left\{ d \in \mathbb{R}^T : |d_t - d_t^u| \le \rho G_t, \quad t = 1, 2, \ldots, T \right\} \tag{13}$$

where the positive numbers G_t represent "uncertainty scale".

With the above choices of the linear decision rules (eqs. (9)–(11)) and the uncertainty sets (12)–(13), the AARC formulation corresponding to our RSFC problem (2) is finally the following:

$$\min_{\substack{q_t^\tau, x_{t+1}^\tau, y_t^\tau, \\ u_t^\tau, w_t, z_t, E}} E$$

s.t.

$$\begin{array}{ll}
(a): \; E \ge s\left[x_{T+1}^0 + \sum_{t=1}^{T} x_{T+1}^\tau d_\tau\right] + \sum_{t=1}^{T}\left[c_t q_t^0 + y_t^0 + u_t^0 + z_t + \sum_{\tau=1}^{t}[c_t q_t^\tau + y_t^\tau + y_t^\tau + u_t^\tau] d_\tau\right], & \\
 & \forall\, d \in \mathcal{U}_1 \\
(b): \; x_{t+1}^0 + \sum_{\tau=1}^{t} x_{t+1}^\tau d_\tau = x_t^0 + \sum_{\tau=1}^{t-1} x_t^\tau d_\tau + q_t^0 + \sum_{\tau=1}^{t} q_t^\tau d_\tau - d_t, & \forall\, d \in \mathcal{U}_2 \\
(c): \; L_t \le q_t^0 + \sum_{\tau=1}^{t} q_t^\tau d_\tau \le U_t, & \forall\, d \in \mathcal{U}_2 \\
(d_1): \; y_t \ge h_t\left[x_{t+1}^0 + \sum_{\tau=1}^{t} x_{t+1}^\tau d_\tau\right] \qquad (d_2): \; y_t \ge -p_t\left[x_{t+1}^0 + \sum_{\tau=1}^{t} x_{t+1}^\tau d_\tau\right], & \forall\, d \in \mathcal{U}_2 \\
(e_1): \; u_t \ge r_t^+\left[q_t^0 + \sum_{\tau=1}^{t} q_t^\tau d_\tau - w_t\right] \qquad (e_2): \; u_t \ge -r_t^-\left[q_t^0 + \sum_{\tau=1}^{t} q_t^\tau d\tau - w_t\right], & \forall\, d \in \mathcal{U}_2 \\
(f_1): \; z_t \ge v_t^+(w_t - w_{t-1}) \qquad (f_2): z_t \ge -v_t^-(w_t - w_{t-1}) & \\
(g): \; q_t^\tau = 0, \; (t, \tau) \in J. &
\end{array} \tag{14}$$

4.3 Deriving a Computationally-Tractable Optimization Problem Equivalent to the AARC Problem (14)

Let v denote the vector consisting of all the decision variables q_t^T, x_{t+1}^τ, y_t^τ, u_t^T, w_t, z_t, E, $(1 \le t \le T, \; 1 \le \tau < t)$. To denote the fact that a function f depends affinely on v, we write $f[v]$.

Now, it is easy to verify that each of the constraints (2.a)–(2.e) (including constraint (2.b), after being converted to two inequalities) has the general form

$$z_0[v] + \sum_{t=1}^{T} z_t[v] d_t \leq 0, \quad \forall, d \in \mathcal{U}_1 \quad (\text{or } d \in \mathcal{U}_2). \tag{15}$$

For the box uncertainty set (13), inequality (15) is equivalent to (we temporarily write $z_t = z_t[v]$)

$$\max_{d \in \mathbb{R}^T} \left\{ z_0 + \sum_{t=1}^{T} z_t d_t \mid |d_t - d_t^n| \leq \rho G_t,\ t = 1, \ldots, T \right\} \leq 0\,. \tag{16}$$

The optimal solution of the problem in the left-hand side of (16) is $d_t = d_t^n + \text{sign}(z_t)\rho G_t$ and so (16) becomes

$$z_0 + \sum_{t=1}^{T} (z_t d_t^n + \rho G_t |z_t|) \leq 0$$

which can be further expressed by the following linear inequalities

$$\left\{ \begin{array}{l} z_0 + \sum_{t=1}^{T} (z_t d_t^n + \rho G_t \eta_t) \leq 0 \\ -\eta_t \leq z_t \leq \eta_t, \quad t = 1, \ldots, T. \end{array} \right. \tag{17}$$

Recalling that $z = z[v]$ is an affine function of v, it follows that (17) remains a linear system of inequalities in the original vector of variables v.

For the ellipsoidal uncertainty set $\mathcal{U}_1$ in (12), inequality (15) is equivalent to

$$\max_{d \in \mathbb{R}^T} \left\{ z_0 + \sum_{t=1}^{T} z_t d_t \mid (d - d^n)^T S^{-1} (d - d^n) \leq S^2 \right\} \leq 0. \tag{18}$$

The problem on the left-hand side of (18) can be solved easily using the Karush-Kuhn-Tucker conditions; the optimal d is given by

$$d = d^n + \frac{\rho}{\sqrt{z^T S z}} S z$$

and so (18) becomes

$$z_0 + \sum z_t d_t^n + \rho \sqrt{z^T S z} \leq 0\,; \tag{19}$$

this is a conic-quadratic constraint in the z_t's. When making an affine substitution of variables $z = z[v]$, (19) becomes a conic quadratic inequality in the original decision vector v.

The analysis we have carried out shows that each of the semi-infinite constraints in the AARC problem (14) can be written equivalently either as a finite set of linear constraints, or a single conic quadratic constraint. Problems with these types of constraints are called Conic-Quadratic (CQ). They are known to be polynomially solvable (see e.g., Ben-Tal and Nemirovski, 2001) and very efficient algorithms (and software, e.g., MOSEK) are available to find their solution.

5 Analysis of the RSFC Model

5.1 Preliminaries

In this section we provide a glimpse into a large set of experiments that we did to test the usefulness of the RO approach to the RSFC problem. The basic set of data we used is given in the Appendix. During the experiments we implemented numerous versions of this basic data where the changes in the data are explained in each relevant sub-section.

Solving a single RSFC model with ellipsoidal demand scenarios for 12 periods with linear decision rules whose length is also 12 periods leads to a mathematical program with 1656 variables, 2231 linear constrains and one conic constraint. To run such programs efficiently thousands of times and compare them to various other models we developed the Robust Commitment Optimization (RobCop) package – a dedicated software that employs Matlab and Mosek at its core. RobCop is composed of two components. The first component uses a friendly user interface to assist users in constructing data sets with various cost parameters, demand uncertainty sets, etc. The second component again employs a friendly user interface to help the user select a desired mode for running the program (e.g., with or without comparisons to dynamic programming or to order-up-to-level policies) and then runs the optimization. A typical run of $20 - 30$ simulations (including various comparisons to alternative policies) requires solving over 1000 problems of the size described above. RobCop executes that task on a Pentium 4 laptop computer in 10-20 seconds. Even when the problem sizes are doubled (say, to 24 periods) solutions are obtained in about 30 seconds. Hence, computation time is not an issue here.

A typical outcome of the RSFC model is depicted in Figure 2 (such figures are automatically generated by our software). The color code we use in the figure is,

- yellow: bounds on demand,
- green: realized demand,

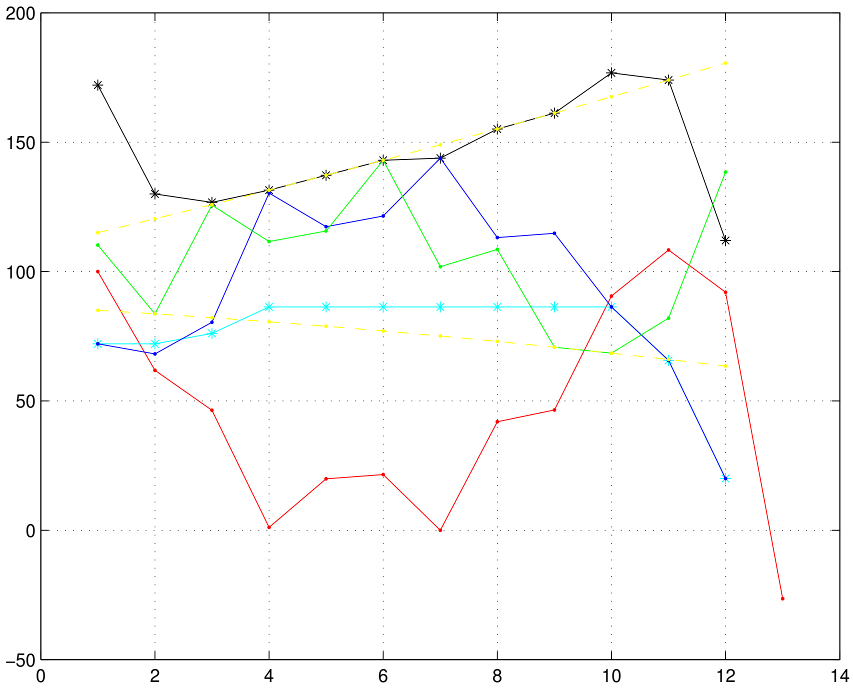

Figure 2: Outcome of a typical RC model run

- blue: orders,
- cyan: commitments,
- red: inventory,
- black: total costs.

The particular outcome depicted in Figure 2 corresponds to a variation of the basic data in which the initial demand variance was ±20% and this variance was expanded gradually in time through the 12-period horizon. The purchase costs in this data fluctuated in a cyclic manner around 40. We see that the retailer's inventory, whose initial value was $x_1 = 100$ fluctuates quite significantly (even becoming zero in periods 4 & 7 and negative in period 13). The commitments follow the fluctuations in the purchase costs and the orders fluctuate at about the same magnitude as the commitments with an exception occurring towards the end of the horizon where the program tries to avoid unwanted inventory.

5.2 Calibration with respect to dynamic programming

Before going into more elaborate experiments we wish to compare the performance of the RO approach with an ordinary dynamic programming (DP) model. However, due to the limitations of the DP approach, we had to make some alterations in the initial model (1).

Attempting to implement DP to that model in its original form would entail a state vector of $T + 1$ dimensions (T commitments and the inventory position at the beginning of the horizon). Since the problems we are typically interested in are associated with horizons whose length can easily exceed 12 periods, this would result in an intractable DP.

Therefore, we formulated a version of model (1) in which there are no commitments and penalty components are only applied on deviations between the sizes of successive orders (i.e., on $|q_t - q_{t-1}|$). We solved this version of the model with stationary demand and obtained almost *identical* results to both the adapted RSFC model and its corresponding DP.

Demand scenario	DP		Adapted RSFC	
	Worst-case	Average	Worst-case	Average
box	53660	44728	53610	44756
ellipsoidal	53660	44706	44827	44662

Table 1: Comparison of DP and RC Solutions

Table 5.2 summarizes the results that were obtained for two demand uncertainty sets (box and ellipsoidal). The "worst case" columns in the table correspond to the basic solution of both DP and adapted RSFC where in both cases the respective programs guarantee that the constraints are always met. Notice that there is no difference in the "box" case (the minor deviation is due to the discretization that we used on the values of x_t to facilitate the use of DP). However, in the "ellipsoidal" case the DP is unaware of the truncation of extreme demand situations and therefore it computes the same worst-case while the RC program takes advantage of this truncation and returns a much better solution. Each of the "average" columns in the table corresponds to 100 simulation runs of 12-period cycles. Notice that in these columns there is no significant difference, in both demand scenarios, between the two technique (again, the minor deviations are due to rounding).

Hence, we conclude that in the relaxed version of our problem, RC is able to obtain the optimal values as verified by DP.

5.3 Comparison with the order-up-to-level policy

The order-up-to-level (OL) policy is a well known and extensively used method for inventory control. In fact, it was shown (Veinott 1966) that under regular conditions this is an optimal policy for stationary demand single item problems. The OL policy parameter ($\mathcal{S}$) is computed on the basis of an optimal fractal (F^*) from the demand distribution:

$$F^* = \frac{\frac{p}{t} - \frac{h}{2}}{\frac{p}{t} + \frac{h}{2}} \tag{20}$$
$$\mathcal{S} = D_{F^*}(t)$$

where t is the time between orders (in our case, $t = 1$) and $D_f(t)$ is the fractal f value of the distribution of demand over a period t.

Some comparative results appear in Table 2. the left side of the table contains information on the demand characteristics assumed for each run: μ_{init} is the mean demand for the first period; μ_{trend} is a linear change in the mean periodic demand; σ_{init} is the initial standard deviation of demand; σ_{trend} is the periodic change in the standard deviation and η is the correlation between demands in successive periods. The middle part of the table describes the RSFC solutions: RC_sub gives a subtotal of the cost that contains only the purchase, holding, penalty and salvage components; RC_tot gives the total cost (including the penalty components) and the "Utopian" column corresponds to average of results achieved by storing in memory the realization of demand in each run and then employing a deterministic program to solve the problem for the known demand. The right part of the table reports the results for the OL policy: O_level is the order-up-to-level value, O_tot is the total cost of the OL policy. To compute this total cost in a way comparable to the RC_tot we had to artificially generate some commitments and then use them to compute the relevant penalties. We did this in a way that is most favorable to the OL policy – after the OL orders became known, we computed for them the *best* possible commitments and then calculated the penalties. Finally, O_sub are the subtotals of the costs for comparison purposes with the RC_sub column.

The first few runs were done with stationary cost parameters (in particular, the purchase cost (c_t) was fixed to 40). Next we let the purchase cost change in a cyclical manner around 40 and finally we also let the penalty costs (p_t) change cyclically around 10.

Not surprisingly, as long as we maintain stationary demand and cost parameters, the OL policy yields excellent results (that is, in these cases the OL results are better than the RSFC results and are not too far from the "Utopian" results). However, when the problem is no longer stationary, we observe rapid deterioration in the performance of the OL policy which quickly becomes inferior to the RSFC model.

5.4 Effects of demand variability

Since demand variability is the source of our problem, it is interesting to see the extent of deterioration in the performance of the RSFC model as we increase the variance in demand. To test this phenomena we ran simulations in which demand was realized from a series of distributions with ever-increasing variance. In each case we saved the vector of realized demand values and later ran the model in a deterministic mode using the saved

Demand Characteristics					RSFC Solutions			OUTL Solutions		
μ_{init}	μ_{trend}	σ_{init}	σ_{trend}	η	RC_sub	RC_tot	Utopian	O_level	O_tot	O_sub
Constant $c_t = 40$										
100	0	20	0	0	45950	46109	44665	102.68	45193	44984
100	3	20	0	0	53692	52632	51803	119.61	52054	52157
100	3	20	5	0	56023	53894	51809	124.37	52038	52124
100	3	20	5	0.2	62355	58271	55500	123.6	55386	55365
100	3	20	5	-0.2	61330	54755	51733	124.1	55050	53608
100	3	20	5	0.5	66653	63976	61816	122.1	60840	60919
Variable $c_t \approx 40$										
100	0	20	0	0	46705	45954	45821	102.87	46469	46321
100	0	20	4	0	49430	46785	46236	105.5	47386	46796
Variable $c_t \approx 40$ **and** $p_t = 5, 15, 15, 5, 15, 15, \ldots$										
100	0	20	0	0	46396	45510	45099	118.9	46653	45970

Table 2: Comparison between RC and Order-Up-To-Level results

demand values. We call the outcomes of such runs "utopian" solutions. We emphasize that comparing the outcome of the RSFC model with the corresponding utopian result is a highly conservative analysis approach since the latter solution is totally theoretic and since there is no technique to accurately forecast the demands in advance, there is no way anyone can achieve such results.

The evaluation was done with data set containing a single set of stationary cost parameters and demand whose mean was held constant at 100 and its variability was gradually increased. Each scenario was repeated over 30 simulations with 12 periods in each run. The results of this comparison appear in Table 3 below.

Variability	RSFC Solution	RSFC Simulation	Utopian	Effect (%)
± 10	49469	48011	47729	0.59
± 20	51100	49333	48839	1.01
± 30	52753	49894	49173	1.46
± 40	54405	49358	48256	2.28
± 50	56058	50636	49346	2.61
± 60	59668	48716	47136	3.35
± 60	63278	48274	46448	3.93

Table 3: Effects of demand variability

As expected, the "utopian" solution is almost indifferent to the variability of demand. The RSFC optimal solution is positively correlated with the demand variability. But, the increase in optimal RSFC cost is not dramatic. In fact, the changes in the cost seem to follow a step function – the first few times the increase in cost per %10 increase in demand variability is about 1630 and the last few times it is about 3600. It is also interesting to point out that in the actual simulations of the RSFC approach, the mean total cost stays stable (in the range $[48000, 50000]$) even though demand variability changes. The relative difference between the RSFC simulation results and the utopian results appear in the "effect" column on the right hand side of the table. While the variability was increased from $\pm10\%$ to $\pm60\%$ the gap between the actual results that one would get with our RSFC model and a utopian solution climbs from less than 1% to less than 4%. This means that while the conservative initial cost estimation must go up with the variability, the implementation of the policy rules that are determined through the RSFC model lead to very reasonable results that are consistently within $\%1 - \%4$ of the utopian results.

5.5 Tradeoff between purchase cost and & forecast reliability penalties

As explained in §1, the motivation for the flexible commitment approach was to assist both the retailer and the supplier to combat uncertainty. The supplier's interest is to encourage the retailer to "absorb" as much as possible the demand variability – that is to make him order consistent orders and tackle the demand uncertainty by holding inventory. To do so, one would expect the supplier to set relatively large penalties on consistent ordering (r_t^+, r_t^-) and on forecast reliability (v_t^+, v_t^-) and offset them by offering discounts on purchase costs (making it easier to the retailer to hold more inventory).

The tool we have developed might thus be used during *contract negotiations* by either the retailer or the supplier in an attempt to identify fair combinations of purchase costs and penalty values that will best serve their wishes vis-á-vis the handling of demand variability. To demonstrate how such tradeoffs might be analyzed we ran a number of "iso-cost" simulations in which decreases in purchase costs were offset by increases in the forecast reliability penalties v while the overall cost to the retailer remained stable. Table 4 presents the results of these runs. Each run consisted of 10 simulations with a cycle of 12 periods in each simulation. The mean total cost (per cycle) to the retailer in each one of the runs reported in this table was 47955 ± 20. The penalties for both upper and lower deviations were kept identical throughout these runs and therefore in the table below $v_t = v_t^+ = v_t^-$

c_t	v_t	mean $(q_t - w_t)^+$
40.2	0.8	294
40.0	2	194
39.4	6	167.1
39.0	11	46.2
38.9	12	24.8
38.8	13	0

Table 4: Tradeoffs between purchase costs and forecast reliability penalties

The table starts with a mean $(q_t - w_t)^+$ deviation per cycle of 294 (equivalent to a mean deviation of 24.5 per period). To cut this value by nearly a half (to 167.1 per cycle) we needed to increase the penalty on these deviation by a factor of 7.5 (from 0.8 to 6). To maintain the same total cost, the purchase cost was reduced from 40.2 to 39.4. The same pattern continued until the penalty costs were set at $v_t^+ = v_t^- = 13$ where the deviations became zero.

Thus, using such analyses, a supplier might find suitable combinations of purchase cost and penalty values that will encourage the retailer to stick to his original commitments without changing the overall cost.

5.6 Tradeoff between purchase cost and & minimum order quantities

The lower bounds L_t in the RSFC model (2) play a crucial role in the supplier-retailer relations. From the supplier's perspective, it is desirable to see large values of these minimum order quantities. From the retailer's point of view, however, the larger these bounds are the less flexibility he has and the larger is the probability that he will be stuck with unnecessary inventories. Once again, our tool may be used during contract negotiation stages, by either of the negotiating partners, to determine appropriate combination of values for these bounds and related cost parameters.

We demonstrate the tradeoff and the possible gains it may bring to both the retailer and the supplier by repeatedly applying our formulation to find alternative combinations of purchase cost and minimum order quantities that simultaneously cause a gradual increase in the purchase cost (which represents the revenue of the supplier) and a gradual decrease in the total cost (which represents the retailer's expenses) . To make the comparison valid, we generated a single realization of demand for a path of 10 cycles with 12 periods in each cycle and used it throughout this experiment. The tradeoff is shown in Table 5 below.

c_t	L_t	$\sum q_t$	purchase cost (supplier's revenue)	total cost (retailer's expenses)
40	50	1130	45233	47837
39.5	60	1150	45457	47312
39	70	1180	46052	46826
38.8	80	1191	46204	46608

Table 5: Tradeoffs between purchase costs and minimum order quantities

Initially, with purchase cost of 40 and minimum order quantity of 50, the retailer's mean total order per cycle is 1130 and his mean total cost is 47837. Raising the lower bound to 60 and offsetting it with a 0.5 discount in the purchase cost, the supplier is able to make the retailer order 1150 items per cycle (thus increasing his revenue from 45233 to 45457) while causing a slight decrease in the retailer's total cost (to 47312). The same pattern continues until we reach an near-perfect equilibrium in which the total cost to the retailer is nearly equal to the supplier's revenues.

6 Folding Horizon Formulation

In this section we discuss an extension of the finite horizon model we analyzed above. Recall that the outcome of our formulation is a vector of T commitments, a firm order for the first period (q_1) and a policy (linear decision rule) for determining the future orders ($q_2, ..., q_T$). Let's assume that the commitments must be submitted at the beginning of the horizon to enable the supplier to make the necessary preparations. But, why shouldn't the retailer try to improve the quality of his policy regarding future orders in response to observed demand? This would lead to an approach which we denote as "folding horizon". According to this approach, at the first period the retailer will solve a finite horizon problem of T periods to determine his commitments, his first order and his policy for future orders. At the end of the first period he will solve a problem of $T-1$ periods to determine q_2 and a new policy for future orders (the commitments are no longer decision variables since they were already determined in the first period). This will be repeated until the end of period $T-1$ where he will solve a single period problem to determine q_T.

The operations management literature usually categorizes horizons into finite and infinite types with further refinements regarding forecast, solution and study horizons. These issues were recently analyzed in a survey by Chand et al., (2002). Although the possibility of rolling horizons with varying length is mentioned in the survey (ibid. p. 27), there is no discussion of a folding horizon similar to the one we explore here.

Folding horizons are common in the fashion industry where due to long lead times production commitments must be made a year or two prior to actual sales and the salvage value of the leftover inventory after the sales period is over is rather small. Fisher and Raman (1996) describe the difficulties associated with such scenarios and provide the motivation for a folding horizon approach. However, due to the complexity of the resultant formulation, they propose a model that assumes a 2-period horizon only.

The implementation of the folding horizon method requires only slight adjustments in the RSFC model. The first run of any folding horizon model is identical to the corresponding run of a fixed horizon model. Then, to execute the second run we implement the following adjustments:

- The horizon is shortened by one period.
- The commitments (w_t) are fixed according to their optimal values in the first run.
- The starting inventory is fixed according to the realization of demand and order replenishment from the previous period.

The above adjustments are repeated each time we run the model until the entire original horizon is exhausted.

7 Analysis of the Folding Horizon Model

The folding horizon approach allows the retailer to solve his problem again and again, each time with a shorter horizon and with more accurate information on his starting position. Hence, we can only expect it to improve the results obtained through the fixed horizon approach. To test this hypothesis we constructed three data files (denoted below as "stat", "non-stat1" and "stat2"). The "stat" file was simpler in the sense that it contained stationary demand ($D \sim [100, \pm 20\%]$) and stationary cost parameters. The "non-stat" file contained non-stationary demand (mean starting at 100 and increasing by 2 every period and uncertainty starting at ±15% and increasing by 3% every period) and non-stationary cost parameters. The cost parameters in the "stat2" file were identical to those in the "non-stat" file and its demand was stationary but with variability that was three times larger than the one in "stat1" (that is, $D \sim [100, \pm 60\%]$).

We ran the RobCop software with these files under two uncertainty sets (box and ellipsoid) and two types of horizons (12 and 24 periods). Each run included 20 simulations and so the number of RO solutions required for this experiment was 3800! Table 6 reports the results of these comparisons. To ensure fair comparisons, the realizations of demand generated for each simulation of the folding horizon model were stored in memory and later used for the corresponding simulation of the fixed horizon model (that is, the demand data is identical within each line of this table but it is different across different lines).

Data & uncertainty sets			Folding Horizons			Fixed Horizon		
Length	Uncertainty	File	Min	**Mean**	Max	Min	**Mean**	Max
12	box	stat	42748	**47874**	56064	43703	**48501**	56155
12	box	non-stat	46627	**52529**	58627	47830	**53423**	59374
12	box	stat2	36964	**47348**	56671	40006	**49228**	56419
12	ellipsoid	stat	44250	**50121**	55726	44207	**50172**	55825
12	ellipsoid	non-stat	45574	**53185**	61029	45579	**53183**	61310
24	box	stat	87033	**95920**	106395	87884	**96484**	106628
24	box	non-stat	108408	**118881**	131934	111150	**121160**	133640
24	box	stat2	88785	**96498**	106547	90927	**98438**	107017
24	ellipsoid	stat	90919	**97738**	107283	90836	**97702**	107147
24	ellipsoid	non-stat	114288	**122452**	130594	114307	**122540**	130691

Table 6: Comparison between fixed and folding horizon models

For each run, Table 6 reports the minimum, mean and maximum values that were recorded during the 20 simulations. Starting with the mean values (which appear in bold font) we observe that for box uncertainty sets the folding horizon model yields consistently better results than the fixed horizon model. For stationary demand (file "stat"), the gap between these two models is not significant and it stays rather stable (in relative terms) as the length of the horizon increases [2]. However, with larger demand variability coupled with non-stationary costs (file "stat2"), the advantage of the folding horizon over the fixed horizon model becomes more evident.

For ellipsoid uncertainty sets, on the other hand, there is virtually no difference between the two models. The explanation to this phenomenon is that in the "box" scenarios the fixed horizon solution must protect itself against very rare events and so it produces order policies that are suboptimal for the more likely events. When the assumed uncertainty is an ellipsoid, the rare events are truncated anyway and then the fixed horizon produces order policies that are identical to those produced by the folding horizon models.

Turning now to the "Min" and "Max" values, we observe that the range between them doesn't change much within each line in the table as the range follows very closely the gradual shifts in the values of the means. In relative terms, these ranges are about 15% of the mean values. Therefore, one must replicate these simulations a sufficient number of times in order to gain confidence in the results obtained.

8 Extensions

The robust optimization approach we propose herein enables the extension of the fixed and rolling horizon models in directions that could not be pursued with the previous methods suggested for the flexible commitments class of problems. We discuss some of these possible extensions below.

8.1 Penalty on deviations from total commitment

An interesting variation of our RSFC model is a formulation in which at the beginning of the finite horizon the retailer generates a single value of total commitment (W) for the entire horizon rather than a vector of periodical commitments. Penalties on deviations from this total commitment are computed only at the last period of the horizon on the basis of the difference between the total accumulated orders and the total commitment. This formulation can be implemented in a finite or folding horizon mode just like the RSFC model we have presented above. All that is required to implement this model variation is a slight adaptation of the objective function in model (2). The terms that involve v_t^+, v_t^- disappear, and the terms that involve r_t^+, r_t^- are taken out of the summation

[2] The table reports on costs-per-cycle. The mean cost for a cycle of 12 periods in the first line is 47874 and the corresponding cost for a cycle of 24 periods is 95920. Similar effect happens with other pairs of related lines in the table. Hence, we conclude that the cost is linear with the length of the horizon.

and replaced by: $r^+max[\sum_{t=1}^{T} q_t - W, 0] + r^-max[W - \sum_{t=1}^{T} q_t, 0]$. The rest of the steps in transforming this adapted formulation to an AARC model is identical to the ones shown above.

8.2 Random yield

Bassok et al. (1997) list a number of key elements in a typical supply contract. In the list they include *quality* which they define as "the upper limit on the percentage of defective units in the supply that is acceptable to the buyer." But, they do not include this element (which was also referred to in the operations management literature as "random yield" – see, e.g., Wang and Gerchak 1996) in their model nor have others tried to incorporate it in any of the flexible commitment models. The reason for this omission is quite simple – including a random yield phenomena in these models would have rendered them intractable.

treatment of random yield in our case can be obtained by defining that the actual quantities which the retailer would receive in period t is given by $\omega \cdot q_t$ where ω is a random variable defined over the interval $[0, 1]$ and q_t is the order for the relevant period.

9 Conclusions

References

[1] Anupindi, R. and Y. Bassok. 1998. Supply Contracts with Quantity Commitments and Stochastic Demand. In S. Tayur, M. Magazine and R. Ganeshan (Eds.) *Quantitative Models for Supply Chian Management*, Kuwer Academic Publishers.

[2] Bassok, Y. and R. Anupindi. 1997. Analysis of Supply Contracts with Total Minimum Commitment. *IIE Transactions*, **29**, 373-381.

[3] Bassok, Y., A. Bixby, R. Srinivasan and H.Z. Wiesel. 1997. Design of Component-Supply Contracts with Commitment-Revision Flexibility. *IBM Journal of Research and Development*, **41(6)**, 1- 14.

[4] Ben-Tal, A., A. Goryashko, E. Guslitzer and A. Nemirovski. 2002. Adjusting Robust Solutions of Uncertain Linear Programs. Submitted to *Math Programming*. Also available at http://iew3.technion.ac.il/Labs/Opt/

[5] Ben-Tal, A., M. Kocvara, A. Nemirovski and J. Zowe. 1999. Free Material Design via Semidefinite Programming, The Multi Load Case with Contact Conditions. *SIAM Journal on Optimization*, **9**, 813-832.

[6] Ben-Tal, A. and A. Nemirovski. 1997. Robust Truss Toplogy Design via Semidefinite Programming. *SIAM Journal on Optimization*, **7**, 991-1016.

[7] Ben-Tal, A. and A. Nemirovski. 2001. *Lectures on Modern Convex Optimization: Analysis, Algorithms and Engineering Applications*, SIAM-MPS Series in Optimization, SIAM, Philladelphia.

[8] Ben-Tal, A. and A. Nemirovski. 2002. A Robust Optimization - Methodology and Applications. *Math Programming (Series B)*, **92**, 453-480.

[9] Chand, S., V.N. Hsu and S Sethi. 2002. Forecast, Solution and Rolling Horizons in Operations Management Problems: A Classified Bibliography. *Manufacturing & Service Operations Management*, **4**, 25-43.

[10] Chen, F.Y. and D. Krass. 2001. Analysis of Supply Contracts with Minimum Total Commitment and Non-Stationary Demands. *European Journal of Operational Research*, **131**, 309-323.

[11] Fisher, M. and A. Raman. 1996. Reducing the Cost of Demand Uncertainty Through Accurate Response to Early Sales. *Operations Research*, **44**, 87-99.

[12] Fisher, M., K. Ramdas and Y.S. Zheng. 2001. Ending Inventory Valuation in Multiperiod Production Scheduling. *Management Science*, **45**, 679-692.

[13] Lee, H., P. Padmanabhan and S. Whang. 1996. Information Distortion in a Supply Chain: The Bullwhip Effect. *Management Science*, **43**, 546-558.

[14] Monahan, J.P. 1984. Quantity Discount Pricing Models to Increase Vendor Profits. *Management Science*, **30**, 720-726.

[15] Tsay, A.A. 1995. *Supply Chain Control with Quantity Flexibility*. Ph.D. dissetation, Graduate School of Business, Stanford University, Stanford, CA.

[16] Tsay, A.A. and W.S. Lovejoy. 1999. Quantity Flexible Contracts and Supply Chain Performance. *Manufacturing & Service Operations Management*, **1**, 89-111.

[17] Urban, T.L. 2000. Supply Contracts with Periodic Stationary Commitments. *Production and Operations Management*, **9**, 400-413.

[18] Veinott, A.F. 1966. State of Mathematical Inventory Theory. *Management Science*, **12**, 747-777.

[19] Wang, Y.Z. and Y. Gerchak. 1996. Periodic review production models with variable capacity, random yield, and uncertain demand. *Management Science*, **42(1)**, 130-137.