

Retailer-Supplier Flexible Commitments Contracts: A Robust Optimization Approach

Aharon Ben-Tal, Boaz Golany, Arkadi Nemirovski

Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, Haifa 32000, Israel
{abental@ie.technion.ac.il, golany@ie.technion.ac.il, nemirovs@ie.technion.ac.il}

Jean-Philippe Vial

Department of Management Studies, University of Geneva, 40 Bd du Pont d'Arve, CH-1211 Geneva 4, Switzerland,
jean-philippe.vial@hec.unige.ch

We propose the use of robust optimization (RO) as a powerful methodology for multiperiod stochastic operations management problems. In particular, we study a two-echelon multiperiod supply chain problem, known as the retailer-supplier flexible commitment (RSFC) problem with uncertain demand that is only known to reside in some uncertainty set. We adopt a min-max criterion, whereby the cost function is minimized against the worst case demand occurrence. To solve the min-max RSFC problem we employ a recent extension of the RO method adapted to dynamic decision problems and known as the affinely adjustable robust counterpart (AARC) methodology. The AARC solution is tested by a large simulation study and found to provide excellent results.

Key words: robust optimization; affinely adjustable robust optimization; flexible commitment contracts; supply chain management; min-max criterion

History: Received: February 24, 2003; accepted: April 27, 2005. This paper was with the authors 9 months for 4 revisions.

1. Introduction

The field of operations management (OM) contains a vast collection of multiperiod problems concerning production, inventory, scheduling, and distribution decisions. Many of these problems involve uncertain quantities such as demand rates, cost coefficients, lead times, etc. Dynamic programming (DP) has long emerged as the leading methodology to address such problems and has led to significant breakthroughs in the early 1960s such as the optimality of base stock or (s, S) policies (Scarf 1960, Veinott 1966). Stochastic programming with recourse (SPR), which is the prevailing approach to uncertain mathematical programs (particularly linear programs) in the OR literature, has also been applied to OM problems (see, e.g., van Delft and Vial 2004).

DP is an attractive and powerful technique. It models the overall dynamic decision process as a sequence of simpler optimization problems. This feature makes it possible to reveal the theoretical structure of the optimal policy for *simple* systems. DP is also very

flexible: it enables the modeling of problems with nonlinear transitions or problems in which the actions influence the transition probabilities. DP also has its limitations, and some of them are quite severe. The main drawback is the complexity of the underlying recursive optimization problems that explode with the number of state variables, thus making DP impractical for computing the actual policy parameters of large problems. This phenomenon, which has become known as the “curse of dimensionality,” has motivated the development of numerous ad-hoc heuristic approaches that could only offer suboptimal solutions. DP also requires that the state variables correctly summarize the past history. This requirement is violated by time correlated uncertainties. This drawback can be circumvented with the help of auxiliary state variables, but this fix is of limited help because it intensifies the curse of dimensionality. Finally, the time separability assumption makes it difficult to model some global constraints, such as a bound on the risk measured by the expected shortfall,

because these constraints link several periods but are not separable.

The SPR approach can be considered as a viable alternative to *some* OM problems. Contrary to DP, SPR is not plagued with the curse of dimensionality in the state variables. It also handles global nonseparable constraints very naturally (van Delft and Vial 2004). However, its usefulness is limited to problems with very few periods, because it is adversely affected by the exponential explosion of the event tree when the number of periods increases.

Approximate dynamic programming (ADP) (e.g., Bertsekas and Tsitsiklis 1996, Schweitzer and Seidmann 1985) offers an alternative to the standard DP approach. It consists of approximating Bellman's value function and typically using either simulation-based learning or linear programming to obtain the parameters that define the approximation. In this way, it becomes possible to work with larger dimensional spaces. The simulation-based methods are involved and do not offer a complexity bound. The linear programming approach is relatively easier to implement and solves an optimization problem with few decision variables, yet still with a large number of constraints. This is in contrast with the small size of the polynomially solvable LP or conic-quadratic problems that arise in the methodology used in this paper.

Another difficulty associated with both stochastic DP and SPR, where the typical objective is to minimize expected cost, is the need to provide the probability distribution functions of the underlying stochastic parameters. This requirement creates a heavy burden on the user because in many real-world situations, such information is unavailable or hard (costly) to obtain. Thus, the need arises for a new optimization methodology that can address the uncertain nature of the problem without making specific assumptions on probability distributions, which is applicable to a wide range of OM problems and is computationally tractable for problems with a large number of state and decision variables, time periods, and stochastic parameters.

The main purpose of this paper is to introduce robust optimization (RO) as a general purpose computational approach that can solve complex OM problems. RO was originally designed to handle static

problems that can be formulated as linear programming (or conic-quadratic) problems with uncertain parameters, where these parameters can reside anywhere in the LP (in the cost vector, the right-hand side, or the activity matrix)—see Ben-Tal and Nemirovski (2002). The description that RO requires for the uncertainty is rather crude—the uncertain parameters are assumed to reside within a deterministic “uncertainty set.” RO adopts a min-max approach that addresses data uncertainty by guaranteeing the feasibility and optimality of the solution against all instances of the parameters within the uncertainty set.

The RO modeling technique has been successfully applied to some large-scale and highly complex engineering design problems (Ben-Tal and Nemirovski 1997, Ben-Tal et al. 1999) and is gradually taking a place in optimization similar to the role of robust control in control theory. Recently, Bertsimas and Thiele (2004) have developed RO models for various supply chain settings and showed the advantages of the technique over DP in situations where the underlying probability distribution of the uncertain parameters is not known exactly. For a summary of the state of the art in RO, the reader is referred to Ben-Tal and Nemirovski (2002) and references therein.

The original RO model deals with static problems where all the decision variables have to be determined before any of the uncertain parameters are realized. This is not the typical situation in most OM problems that are multiperiod in nature, and where a decision at any period can and should account for data realizations in previous periods. Recognizing the need to address such dynamic environments, RO was recently extended into a new paradigm termed as affinely adjustable robust counterpart (AARC); see Ben-Tal et al. (2004). A main feature in AARC is that part of the decision variables can be determined after a portion of the uncertain data is realized. In the AARC method, the dependence of these “adjustable” variables on the realized data is restricted to be in the form of affine functions. This restriction is imposed to achieve tractability. Indeed, in the AARC method, the family of uncertain linear programs is replaced by a *single tractable deterministic problem* (either a linear or a conic-quadratic one). The price to pay for using

only suboptimal policies is, of course, problem specific but, at least for the problem we discuss here, a large simulation study reveals that it is very low.

To demonstrate the use of AARC to OM settings, we consider in this paper a two-echelon T -period supply chain problem known as the retailer-supplier flexible commitment (RSFC) model. This problem involves a retailer who faces uncertain demand for a product by end customers. The retailer orders from a supplier and operates under a contract whereby he commits (at time zero) to a vector of order quantities over a fixed time horizon, and then dynamically replaces these committed quantities with actual orders. Besides the usual costs (ordering, holding, and shortage) and ending salvage values, the retailer incurs additional cost stemming from penalties on deviations between actual and committed orders and deviations between successive committed/actual orders.

Like any other method, AARC also has some limitations. At present, it is applicable only to multistage *finite* horizon problems whose deterministic versions are *linear* programs. Also, the method works only on problems where the uncertainty is exogenously given, rather than influenced by endogenous decision variables. For example, it will not be applicable to multistage pricing problems where the market price chosen today influences the uncertainty associated with the demand tomorrow. The AARC method uses a min-max objective function, which is reasonable if no reliable knowledge of the probability distributions of the uncertain parameters is available. Otherwise, an objective function based on expectation, median, etc., may be more appropriate and will result in a less conservative policy. However, we argue that even in the latter case, a decision maker who wants a strong guarantee on the performance of his chosen policy in the face of uncertainty may adopt a min-max approach. Finally, the AARC uses a linear decision rule to obtain an approximated solution and there is no guarantee that the optimal solution (which is often impossible to find) is close to a linear decision rule.

The rest of this paper is organized as follows. In §2 we describe in detail the RSFC model and review the relevant literature on this model. We first consider the RSFC problem in the case of known (certain) demand and model it as a linear programming problem. For

uncertain data we develop in §2.2 the min-max-based adjustable robust optimization model (called min-max RSFC) that allows decisions at period t to be functions of past demands (i.e., *policies*). We then outline in §2.3 the DP approach to solve the min-max RSFC problem and discuss its limitations. For the case where the uncertainty is a T -dimensional box, we show that an optimal solution of the min-max RSFC problem can be obtained by solving a large-scale linear program (with design dimension of order 2^T). In §3 we present our general methodology for solving min-max multiperiod uncertain linear programming problems. We start with a nonadjustable robust counterpart (RC) formulation, develop the corresponding *adjustable* RC (ARC) model, and then approximate it through the AARC method. We then apply the general AARC methodology to the specific RSFC problem and derive a single deterministic convex optimization problem that is either a linear program or a conic-quadratic one—thus, a tractable problem even for large-scale instances. In §4 we report on a large simulation experiment that we designed to test the performance of the AARC. First, we benchmark the AARC and the RC solutions against the optimal solution of a min-max RSFC problem with box uncertainties. Then, we analyze the actual cost that AARC yields through its mean value across many simulations, and compare it to the mean solution that might be generated under perfect information conditions. This is followed by several additional analyses in which we tested other aspects of the proposed approach. These aspects include an investigation of the linear decision rule policies that AARC yields, a comparison to base stock policies, an investigation of the effects of box versus ellipsoidal uncertainty sets, a study of the effects of information gaps, analysis of the trade-off between various parameters, and a study of a folding-horizon version of the problem. Then, in §5 we briefly discuss possible extensions of the basic RSFC model to incorporate additional features such as multiproduct settings, randomness in the cost coefficients, random yield in the fulfillment of the retailer's orders by the supplier, and additional types of penalties on deviations from cumulative commitments. Finally, §6 offers a summary and concluding remarks.

2. The RSFC Problem

2.1. Literature Review

Demand uncertainty is a source of contention among supply chain parties. In conventional settings, each party tries to pass on as much of the uncertainty burden as possible to other parties. From a retailer's perspective, it is desirable to hold as little inventory as possible. A retailer orders in each period and adjusts the orders up or down according to the customers' actual demand. However, shortage costs motivate the retailer to keep some inventory, and order costs often cause him not to order in each period. When this happens, the situation worsens because the variability increases as information flows backward in the chain. This contributes to the celebrated "bullwhip" effect (see, e.g., Lee et al. 1996). A supplier naturally wishes the opposite—he would like to see the retailer buy the material in advance and keep a large inventory to meet the fluctuations in external demand. One of the traditional tools that suppliers have 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 coordination mechanism, known as "flexible commitments" (Bassok et al. 1997, Bassok and Anupindi 1998, Anupindi and Bassok 1999). The purpose of a flexible commitment agreement between a retailer and a supplier is to assist both of them in facing 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 of his 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. In return, the supplier offers him a discounted cost as long as his actual orders do not differ too drastically from the original commitments and charges certain penalties otherwise. Now, given these cost and penalty parameters and facing uncertain demand, the retailer's challenge is to find an optimal ordering policy minimizing the maximal (with respect to all possible demand) total cost. We refer to this class of problems as retailer-supplier flexible contracts (RSFC) models.

The concept of flexible commitments was proposed in several variants by its original developers. In Bassok and Anupindi (1997), the authors specify a model in which the buyer commits to purchasing at least a given total quantity of a single product over a finite time horizon. The model specifies an additional volume that is also available at the same price and a higher price is charged for any quantity larger than that volume. Another variant (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 RSFC model in its different variants is an NP-hard problem. Although it is possible to formulate it by means of standard inventory recursions, it is impractical to solve it through a standard dynamic programming procedure when the number of decision and state variables is large. Recognizing this difficulty, researchers have developed various heuristics to solve the problem. Bassok et al. (1997) developed lower and upper bounds and then showed that the gap between them is small enough to allow using the upper bound as an approximate solution. The lower bound is based on relaxing the constraints in the model that relate order quantities to commitments—yielding a standard newsvendor problem. The upper bound is generated by maximizing the probability of being able to raise the inventory in each period to the base stock level (as obtained by the newsvendor solution). For the T -period RSFC problem with minimum total commitments, Bassok and Anupindi (1997) showed that the optimal policy is given by T critical numbers, consisting of T order-up-to levels (S_1, S_2, \dots, S_T) and an additional order-up-to level S^M that corresponds to a single-period newsvendor problem with zero purchase costs. Until the minimum commitment is met (say, at period t), the retailer orders according to S^M , and from that period onward he orders according to S_{t-1}, \dots, S_1 . The T order-up-to levels are computed through a stochastic dynamic program, and S^M is found through the ordinary newsvendor solution.

The pioneering work of Bassok and Anupindi was joined by Tsay (1995) and later by Tsay and Lovejoy

(1999) who studied a rolling-horizon contract in a multiple echelon setting, allowing for nonstationary demand with information updating. They defined two types of supply nodes: a “flex” node, which deals with internal demand from a downstream production stage, and a “semiflex” node that deals with external demand. The authors developed an open loop feedback control (OLFC) heuristic to determine the actions of a flex node and showed that a deterministic OLFC model yields a “minimum commitment” policy, where at any period the present decisions minimize the exposure to future costs subject to meeting service requirements. For a semiflex node they again formulate an OLFC model. However, the latter does not induce a deterministic formulation. Hence, the authors had to develop heuristics that relaxed the model and led to a “sequential fractile” policy that is a generalization of a multiperiod newsvendor problem. More recently, Urban (2000) suggested extensions to multiproduct, multiconstraint models, and Chen and Krass (2001) developed a total order quantity commitment model and studied its performance under various conditions.

All the exact models formulated in the papers reviewed above suffer from the “curse of dimensionality” that is common to stochastic DP models. Therefore, optimal solutions for these models can be obtained, at best, for small size problems; while for more realistic problem sizes they offer only approximate solutions that are based on heuristic procedures. Furthermore, most of the heuristics mentioned above are based to a certain extent on solutions of the well-known newsvendor problem. When the structure of the problem deviates from the classical assumptions pertaining to the newsvendor model, the performance of these heuristics may be questionable. In contrast, the RO methodology we propose hereafter does not depend on the classical assumptions and is quite capable of addressing a wide variety of flexible commitment arrangements at the expense of more sophisticated computation.

2.2. Model Formulation

We consider a single-product, two-echelon, multiperiod supply chain in which inventories are managed periodically over a finite horizon of T periods. At the beginning of the planning horizon the retailer

specifies a vector of commitments $w = w_1, \dots, w_T$ for the product. These commitments serve as forecasts for the supplier who uses them to determine his production capacity. At the beginning of each period t , the retailer has an inventory of size x_t and he orders a quantity q_t from the supplier at a unit cost c_t . The customers then place their demands d_t . The retailer’s status at the beginning of the planning horizon is given through the parameters x_1 (initial inventory) and w_0 (a nominal value that might represent the last order prior to the planning horizon or some average of previous orders). Consequently, the following costs are incurred:

- holding cost = $h_t \max[0, x_t + q_t - d_t]$, where h_t are the unit holding costs,
- shortage cost = $p_t \max[0, d_t - x_t - q_t]$, where p_t are the unit shortage costs.

Moreover, due to the stipulations in the contract, the retailer incurs the following additional costs:

- penalty due to deviations between committed and actual orders (“forecast error”)

$$\alpha_t^+ \max[q_t - w_t, 0] + \alpha_t^- \max[w_t - q_t, 0],$$

where α_t^+ , α_t^- are the unit penalties for positive and negative deviations, respectively;

- penalty on deviations between successive commitments

$$\beta_t^+ \max[w_t - w_{t-1}, 0] + \beta_t^- \max[w_{t-1} - w_t, 0],$$

where β_t^+ , β_t^- are the associated unit penalties.

Inventory x_{T+1} left at the end of period T has a unit salvage value s .¹ To make sense in our context, s must be smaller than c_T . Also, as we show later, to maintain convexity of the objective function in our model, s must satisfy for the terminal period T the inequality

$$h_T - s \geq -p_T. \quad (1)$$

The constraints in the model include

- “balance” equations that link the inventory in each period to the inventory, order quantity, and demand in the preceding period;
- upper and lower bounds on the order quantities in each period, $L_t \leq q_t \leq U_t$; and

¹ For a thorough discussion of ending inventory valuation, the reader is referred to Fisher et al. (2001).

(c) upper and lower bounds on cumulative order quantities in each period, $\hat{L}_t \leq \sum_{\tau=1}^t q_\tau \leq \hat{U}_t$.

The choice of values to be assigned to the parameters in the last two types of constraints ($L_t, U_t, \hat{L}_t, \hat{U}_t$) offers a flexible modeling of various contractual agreements between the retailer and the supplier. For example, if the retailer only commits to a minimal cumulative quantity L for the entire horizon (as in Bassok and Anupindi 1997) we would set

$$L_t = \hat{L}_t = 0 \quad \forall t < T, \quad L_T = 0, \quad \hat{L}_T = L, \quad \text{and} \\ U_t = \hat{U}_t = \infty \quad \forall t.$$

If all the model's parameters are known, including the demand vector $d = (d_1, \dots, d_T)$, then the RSFC problem is represented by the following deterministic mathematical program (where $[\xi]^+ \equiv \max[\xi, 0]$):

$$\min_{x, q, w} \left\{ -s[x_{T+1}]^+ + \sum_{t=1}^T [c_t q_t + h_t [x_{t+1}]^+ + p_t [-x_{t+1}]^+ \right. \\ \left. + \alpha_t^+ [q_t - w_t]^+ + \alpha_t^- [w_t - q_t]^+ + \beta_t^+ [w_t - w_{t-1}]^+ \right. \\ \left. + \beta_t^- [w_{t-1} - w_t]^+ \right\} \quad (2)$$

$$\text{s.t. (a): } x_{t+1} = x_t + q_t - d_t, \quad t = 1, \dots, T,$$

$$\text{(b): } L_t \leq q_t \leq U_t, \quad t = 1, \dots, T,$$

$$\text{(c): } \hat{L}_t \leq \sum_{\tau=1}^t q_\tau \leq \hat{U}_t, \quad t = 1, \dots, T.$$

Problem (2) can be further reduced to a linear programming problem by adding auxiliary variables that eliminate piecewise linear terms in the original formulation. This results in the following LP model:

LP RSFC

$$\min_{w, x, q, y, z} \sum_{t=1}^T [c_t q_t + y_t + u_t + z_t] \quad (3)$$

$$\text{s.t. (a), (b), (c), and } \forall t = 1, \dots, T,$$

$$\text{(e}_1\text{): } y_t \geq \bar{h}_t x_{t+1},$$

$$\text{(e}_2\text{): } y_t \geq -p_t x_{t+1},$$

$$\text{(f}_1\text{): } u_t \geq \alpha_t^+ (q_t - w_t),$$

$$\text{(f}_2\text{): } u_t \geq -\alpha_t^- (q_t - w_t),$$

$$\text{(g}_1\text{): } z_t \geq \beta_t^+ (w_t - w_{t-1}),$$

$$\text{(g}_2\text{): } z_t \geq -\beta_t^- (w_t - w_{t-1}),$$

where $\bar{h}_t = h_t \quad \forall t = 1, \dots, T-1$, $\bar{h}_T = h_T - s$, and the new variables y_t, u_t, z_t are upper bounds on the piecewise linear components in the objective function of (2): y_t represents the holding and shortage costs as well as accounting for the salvage value in the last period; u_t represents the forecast error penalties; and z_t represents the commitments inconsistency penalties.

In reality the demand vector d is, of course, uncertain. Consequently, at any time t , the inventory in the system should be treated as a state variable x_t evolving according to the balance equation $x_{t+1} = x_t + q_t - d_t$, where d_t is the exogenous uncertain demand and the order q_t is chosen by the retailer and can, in principle, be an arbitrary function of the events preceding time t . For a given ordering policy, the events preceding time t are uniquely defined by the past demands $d^{t-1} = (d_1, \dots, d_{t-1})$, so that without loss of generality the retailer's policy can be thought of as a collection $q^t(d^{t-1})$ of decision rules specifying current orders as functions of the past demands. For such a policy, the states of the inventory and all other related variables in (3) (i.e., y_t, u_t) also become functions of demands linked with each other by the constraints

$$x_{t+1}(d^t) = x_t(d^{t-1}) + q_t(d^{t-1}) - d_t, \\ L_t \leq q_t(d^{t-1}) \leq U_t, \quad (4)$$

$$\hat{L}_t \leq \sum_{\tau=1}^t q_\tau(d^{\tau-1}) \leq \hat{U}_t,$$

$$\text{(e}_1\text{): } y_t(d^{t-1}) \geq \bar{h}_t x_{t+1}(d^t),$$

$$\text{(e}_2\text{): } y_t(d^{t-1}) \geq -p_t x_{t+1}(d^t),$$

$$\text{(f}_1\text{): } u_t(d^{t-1}) \geq \alpha_t^+ (q_t(d^{t-1}) - w_t),$$

$$\text{(f}_2\text{): } u_t(d^{t-1}) \geq -\alpha_t^- (q_t(d^{t-1}) - w_t),$$

$$\text{(g}_1\text{): } z_t \geq \beta_t^+ (w_t - w_{t-1}),$$

$$\text{(g}_2\text{): } z_t \geq -\beta_t^- (w_t - w_{t-1}),$$

which should be satisfied for all demand trajectories d^T from a given domain. Note that w_t and z_t relate to decisions that must be made before any realization of the data becomes known, and therefore they are not given as functions of the demand trajectories.

Our goal now is to choose all the dependencies $x_{t+1}(d^t), q_t(d^{t-1}), \dots$ in an optimal fashion under the

restriction that the above constraints (which now become inequalities/equalities between *functions* of the demands) are satisfied for any d^T from the set of all possible demand trajectories. In this context, two standard concepts of optimality are used. The first, used in stochastic programming, assumes that the set of demand trajectories is associated with a probability distribution, and the goal is to minimize the corresponding expected cost. The second concept—the min-max approach—requires the minimization of the maximal cost over all demand trajectories. The latter is the approach we adopt in this paper because (as explained in §1) we wish to address situations where the demand is only known to belong to a (convex and bounded) uncertainty set $\mathcal{U}^T = \mathcal{U}_1 \times \mathcal{U}_2 \times \dots \times \mathcal{U}_T$, where \mathcal{U}_t is the uncertainty of the demand d_t at period t . The min-max model can be written as

Min-Max RSFC Problem

$$\min_{0 \leq w_t, z_t, x_t(\cdot), q_t(\cdot), y_t(\cdot), u_t(\cdot) \in \mathcal{U}^T} \max_{\left\{ \begin{matrix} E = \sum_{t=1}^T [c_t q_t(d^{t-1}) + y_t(d^{t-1}) \\ + u_t(d^{t-1}) + z_t] \end{matrix} \right\}} \quad (5)$$

- s.t. $\forall d^t \in \mathcal{U}^t = \mathcal{U}_1 \times \mathcal{U}_2 \times \dots \times \mathcal{U}_t, \quad t = 1, 2, \dots, T:$
- $x_{t+1}(d^t) = x_t(d^{t-1}) + q_t(d^{t-1}) - d_t,$
- $L_t \leq q_t(d^{t-1}) \leq U_t,$
- $\hat{L}_t \leq \sum_{\tau=1}^t q_\tau(d^{\tau-1}) \leq \hat{U}_t,$
- (e₁): $y_t(d^{t-1}) \geq \bar{h}_t x_{t+1}(d^t),$
- (e₂): $y_t(d^{t-1}) \geq -p_t x_{t+1}(d^t),$
- (f₁): $u_t(d^{t-1}) \geq \alpha_t^+(q_t(d^{t-1}) - w_t),$
- (f₂): $u_t(d^{t-1}) \geq -\alpha_t^-(q_t(d^{t-1}) - w_t),$
- (g₁): $z_t \geq \beta_t^+(w_t - w_{t-1}),$
- (g₂): $z_t \geq -\beta_t^-(w_t - w_{t-1}).$

2.3. Methods for Solving the Min-Max RSFC Problem

A standard approach to solving multiperiod problems such as (5) is dynamic programming (DP) (see Iyengar 2005). Here the formulation of the DP

problem is as follows:

Min-Max DP

$$\min_{0 \leq w_1, w_2, \dots, w_T} \left\{ F(w_1, w_2, \dots, w_T) \right. \\ \left. = \sum_{t=1}^T [\beta_t^+(w_t - w_{t-1})^+ + \beta_t^-(w_{t-1} - w_t)^+] \right. \\ \left. + f_1(w_1, w_2, \dots, w_T, x_1, \hat{q}_1) \right\}, \quad (7)$$

where $f_1(w_1, w_2, \dots, w_T; x_1, \hat{q}_1)$ is computed recursively for fixed w_t s in terms of the state variables x_t = inventory level at start of period t (x_1 is a given data) and \hat{q}_t = cumulative order in periods $1, 2, \dots, t - 1$ ($\hat{q}_1 = 0$), as follows:

$$\left. \begin{aligned} & f_t(w_t, w_{t+1}, \dots, w_T, x_t, \hat{q}_t) \\ & = \min_{q_t \in \hat{Q}_t} \max_{d_t \in \mathcal{U}_t} \{ c_t q_t + \max(\bar{h}_t x_{t+1}, -p_t x_{t+1}) \} \\ & \quad + \alpha_t^+(q_t - w_t)^+ + \alpha_t^-(w_t - q_t)^+ \\ & \quad + f_{t+1}(w_{t+1}, w_{t+2}, \dots, w_T, x_{t+1}, \hat{q}_{t+1}), \end{aligned} \right\} \\ 1 \leq t \leq T - 1 \quad (8)$$

where $x_{t+1} = x_t + q_t - d_t$, $\hat{q}_{t+1} = \hat{q}_t + q_t$, and \hat{Q}_t is the set

$$\hat{Q}_t = \{q_t: L_t \leq q_t \leq U_t, \hat{L}_t - \hat{q}_t \leq q_t \leq \hat{U}_t - \hat{q}_t\}.$$

The end condition for the recursion is explicitly given by

$$f_T(w_T, x_T, \hat{q}_T) \\ = \min_{q_T \in \hat{Q}_T} \max_{d_T \in \mathcal{U}_T} \{ c_T q_T + \alpha_T^+(q_T - w_T)^+ + \alpha_T^-(w_T - q_T)^+ \\ + \underbrace{\max\{(h_T - s)(x_T + q_T - d_T), -p_T(x_T + q_T - d_T)\}}_{g(q_T)} \}.$$

Note that the piecewise linear function $g(\cdot)$ is convex if and only if $h_T - s \geq -p_T$, a condition we assumed a priori (see Equation (1)).

An attempt to solve the min-max RSFC problem via the DP formulation (7) would encounter severe difficulties: the objective function $F(w_1, w_2, \dots, w_T)$ is generically nonsmooth. Each function evaluation needs the solution of the Bellman Equation (8), that (in the case of continuous demand) may entail the discretization of the state variables x_t, \hat{q}_t and hence produce only an approximate value of $F(w)$. Moreover,

the above computation does not produce derivative (or subdifferential) information. Hence, $F(w)$ has to be minimized by a zero-order optimization method (i.e., one that uses only function values), and such methods are notoriously slow. To make things worse, when extending the RSFC problem to the case of more than one product, even the computation of f_t via the Bellman recurrence, for a fixed $w = (w_1, \dots, w_T)$, becomes infeasible for realistic problem sizes.

There is, however, a case where problem (5) itself can be solved directly. This is the case where the uncertainty set is a T -dimensional box:

$$\mathcal{U}_{\text{box}} = \mathcal{U}_1 \times \mathcal{U}_2 \times \dots \times \mathcal{U}_T, \quad (9)$$

where

$$\mathcal{U}_t = [d_t^{\min}, d_t^{\max}], \quad t = 1, 2, \dots, T, \quad (10)$$

with $d_t^{\min} < d_t^{\max}$. In this case, it can be shown² that problem (5) is equivalent to a problem where \mathcal{U}_{box} is replaced by a finite set \mathcal{U}_{ext} consisting of the extreme points of \mathcal{U}_{box} . We call this problem P_{ext} . In this equivalent problem, the uncertainty set includes 2^T *extreme demand trajectories* that can be naturally identified with a path of length T in a binary tree.

Thus, the problem can be identified with a regular stochastic problem with a binary event tree and a min-max objective criterion. The so-called deterministic equivalent is a simple extension of the base deterministic problem, in which the temporal variables and constraints are made contingent to the nodes of the event tree. This is a large LP because for a problem with horizon T , the number of nodes is of order 2^T . When T is a moderate integer like 10–12, the resulting problem can be routinely solved to high accuracy by a standard LP solver. We have used this scheme in our computational study (see §4.2).

3. Robust Optimization (RO) Formulation of the RSFC Problem

3.1. The RO Methodology for Multiperiod Problems

The RSFC problem, like many other OM problems, is an instance of a generic uncertain T -period LP problem whose data $\{(A_t, c_t, b_t): t = 1, 2, \dots, T\}$ depend

affinely on a vector of uncertain parameters $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_T)$:

$$\min_{x=(x_1, \dots, x_T)} \left\{ \sum_{t=1}^T c_t(\lambda)x_t: \sum_{\tau=1}^t x_\tau A_\tau^r(\lambda) \leq b_t(\lambda), \right. \\ \left. t = 1, 2, \dots, T \right\}, \quad (11)$$

where A_t^r and b_t are vectors of the same dimension.

The RO approach was developed to deal with situations in which the vector of uncertain parameters λ is only known to reside within an uncertainty set Λ . Thus, formulation (11) in fact represents a *family* of LPs—one for each possible realization of the uncertain data.

Rather than adopting a feedback control policy (such as the one that is obtained by solving the min-max DP RSFC (7)), suppose that our “min-max decision maker” wants to optimally choose, a priori at the beginning of Period 1, a vector of *fixed* decisions through time. He then solves the following single deterministic problem (the so-called *robust counterpart* (RC) of (11)):

RC

$$E^* = \min_{x=(x_1, \dots, x_T)} \max_{\lambda \in \Lambda} \left\{ \sum_{t=1}^T c_t(\lambda)x_t: \sum_{\tau=1}^t x_\tau A_\tau^r(\lambda) \leq b_t(\lambda) \right. \\ \left. \forall \lambda \in \Lambda, t = 1, 2, \dots, T \right\}. \quad (12)$$

A solution x^* of (12) is called *robust*. Such a solution satisfies the constraints for *all* possible realizations of the data $\lambda \in \Lambda$, and *guarantees* an optimal objective function value not worse than E^* . Note that while the data (A, c, b) depend on the realization of λ , the decision vector x is not dependent on any particular scenario.

Problem (12) being a semi-infinite LP seems computationally intractable, yet it turns out that for a wide variety of compact, convex³ uncertainty sets Λ , the RC model is a tractable (polynomially solvable) convex mathematical problem, typically an LP or a conic-quadratic problem (see Ben-Tal and Nemirovski 2000, 2002). In particular, Bertsimas and Thiele (2004) have

² See the appendix.

³ Note that when Λ is composed of a finite number of points (scenarios), it can be replaced by its convex hull—a compact convex set.

used the RC approach (with polyhedral uncertainty sets, yielding an RC that is an LP) to study classical supply chain problems. They derived qualitative results on the structure of the optimal policy that parallels the classical (s, S) policy.⁴

For the general time dependent (multiperiod) problem (12), where the entire decision vector x is to be determined before actual realization of the uncertain data (A_t, c_t, b_t) occurs, one may encounter difficult cases where no feasible solution exists or where the objective function is grossly overestimated. For example, this is the case when (12) involves flow balance constraints and lost sales rather than backlogging. These difficulties can be remedied if we take into account the fact that in multiperiod settings part of the decisions do not have to be determined a priori. Instead, some of the decisions can be delayed to later periods when part of the uncertain parameters have already become known.

To be specific, let the vector x be partitioned to $x = (x^a, x^{na})$. The subvector x^{na} (nonadjustable variables) consists of “here and now” decisions, i.e., those that must be determined before any uncertain data is realized. (In the RSFC problem, these are: q_1 —the first period order—and the commitments w_1, w_2, \dots, w_T .) The subvector x^a (adjustable variables) consists of “wait and see” decisions that can be adjusted based on the revealed data by the time the decision must be made (in the RSFC problem $x^a = (q_2, \dots, q_T)$).

The partition of $x = (x^a, x^{na})$ induces a partition of the index set $\{1, 2, \dots, T\} = I_a \cup I_{na}$, in terms of which we can now model our time-dependent (multiperiod) uncertain LP as follows:

Min-Max Multiperiod Adjustable Robust Counterpart (ARC)

$$\min_{x^{na}, E} \left\{ E: \text{such that } \forall \lambda \in \Lambda \exists x_t = x_t(\lambda^{t-1}), t \in I_a: \right.$$

$$\sum_{\tau \in I_{na}} c_\tau(\lambda) x_\tau + \sum_{\tau \in I_a} c_\tau(\lambda) x_\tau(\lambda^{\tau-1}) \leq E,$$

$$\sum_{\tau \in I_{na}(t)} x_\tau A_\tau^T(\lambda) + \sum_{\tau \in I_a(t)} x_\tau(\lambda^{t-1}) A_\tau^T(\lambda) \leq b_t(\lambda),$$

$$t = 1, 2, \dots, T \left. \right\}, \tag{13}$$

⁴Note that in the RSFC problem, an (s, S) policy is not necessarily optimal even in the deterministic case due to the additional penalty terms in the objective function and the additional constraints.

where $\lambda^s = (\lambda_1, \dots, \lambda_s)$,

$$I_{na}(t) = I_{na} \cap \{1, \dots, t\}, \quad I_a(t) = I_a \cap \{1, \dots, t\}.$$

The min-max RSFC problem (5) is a particular instance of the general model (13). A special feature of (5) is that its constraints admit a simple form

$$f^\top x^a + g^\top x^{na} \geq l(d), \tag{14}$$

where the right-hand side $l(d)$ is an affine function of the uncertain demand vector d while the parameters on the left-hand side of (14) are certain (\top stands for transpose). As a demonstration, consider constraint $(3e_1)$ for $t < T$. By eliminating the x_t variables using the balance equation (2a), constraint $(3e_1)$ after rearrangement becomes

$$(e_1) \quad y_t - h_t \sum_{\tau=1}^t q_\tau \geq h_t x_1 - h_t \sum_{\tau=1}^t d_\tau.$$

Note that now we have on the left-hand side a linear term with certain coefficients involving the adjustable variables y_t and q_2, \dots, q_t and the nonadjustable variable q_1 . All the uncertain terms are in the right-hand side, which is indeed an affine function of d . In fact, all the constraints in the LP RSFC problem are of the same general form (14).

The added flexibility offered by the ARC model (13) is offset by the fact that it is often computationally intractable (NP-hard). This is the case even for simple uncertainty sets (e.g., a general polyhedral set); see Ben-Tal et al. (2004). The core of the difficulty in solving the ARC model is the unknown functional relations between x_t^a and the “history” λ^{t-1} . To overcome this difficulty, Ben-Tal et al. (2004) suggested to approximate the ARC solution by restricting these functional relations to be *affine*. Thus, each adjustable variable x_t^a in (13) is substituted by the following *linear decision rule* (LDR):

$$x_t^a(\lambda^{t-1}) = \mu_t^0 + \sum_{\tau=1}^{t-1} \mu_t^\tau \lambda_\tau, \tag{15}$$

where the coefficients μ_t^τ are the new decision variables. Note that the actual value of the *policy* $x_t^a(\lambda^{t-1})$ will be determined by (15) only at period t when the vector λ^{t-1} has been realized.

The idea of focusing on LDRs is quite a common heuristic in many branches of science and engineering (linear controllers or linear feedback rules in control theory, linear estimators/predictors in signal processing, etc.). In a different context, LDRs were suggested in the early 1960s by Holt et al. (1960, Ch. 6) to study the existence of a certainty equivalent in stochastic decision problems.

With the affine transformation (15), the ARC model is approximated by the affinely adjustable RC (AARC) model:

AARC

$$E^{**} = \min_{x^{na}, E, \mu_s^{\tau}} \left\{ E: \forall \lambda \in \Lambda: \sum_{\tau \in I_{na}} c_{\tau}(\lambda)x_{\tau} + \sum_{\tau \in I_a} c_{\tau}(\lambda) \right. \\ \cdot \left[\mu_{\tau}^0 + \sum_{s=1}^{\tau-1} \mu_{\tau}^s \lambda_s \right] \leq E, \sum_{\tau \in I_{na}(t)} x_{\tau} A_{\tau}^{\tau}(\lambda) \\ + \sum_{\tau \in I_a(t)} \left[\mu_{\tau}^0 + \sum_{s=1}^{\tau-1} \mu_{\tau}^s \lambda_s \right] A_{\tau}^{\tau}(\lambda) \leq b_t(\lambda), \\ \left. t = 1, 2, \dots, T \right\}. \quad (16)$$

The minimal cost E^{**} that the AARC model yields is “optimal” in the sense that no other solutions to (13), for which the adjustable vector x^a depends linearly on the uncertain parameters, can do better while satisfying the constraints for all possible realizations $\lambda \in \Lambda$. Note that the RC model (12) is a “degenerate” special case of the AARC model where all the variables μ_t^{τ} are forced to be 0 for $t = 1, \dots, T$ and $\tau > 0$.

An important class of AARC models (referred to as “fixed recourse”) corresponds to a situation where the parameters (c_t, A_t^{τ}) , associated with the adjustable variables in the uncertain LP (13), are *not* uncertain. In particular, the RSFC problem belongs to this class (see (14) and the discussion therein). For problems in this class, AARC depends affinely on all uncertain parameters, and its mathematical structure is similar to that of the RC model (12). As a result, AARCs with fixed recourse are computationally tractable for a wide spectrum of uncertainty sets (see Ben-Tal et al. 2004).⁵

⁵ For the nonfixed recourse case, good approximations of AARC (with guaranteed level of proximity to the true optimal solutions) are available (see Ben-Tal et al. 2004, §4).

3.2. 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 (for $t > 1$). Here we set

$$(i) \quad q_t = q_t^0 + \sum_{\tau=1}^{t-1} q_{\tau}^{\tau} d_{\tau}, \quad (17)$$

$$(ii) \quad q_{\tau}^{\tau} = 0 \quad \text{if } (\tau, t) \in J_t,$$

where J_t 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_t may include periods τ that are too distant in the past, or it may include periods that are too close to the present. For example, if there exists a two-period delay in reporting past demands, then $J_t = \{t-1, t-2\}$.

With q_t being affine functions of the demand d_{τ} , $\tau < t$, the balance equations (2a) enforce the variables x_{t+1} to also become affine functions of the d_{τ} s:

$$x_{t+1}(d^t) = x_{t+1}^0 + \sum_{\tau=1}^t x_{t+1}^{\tau} d_{\tau}. \quad (18)$$

With this definition of x_{t+1} , the variables y_t (see (3e)) should become maxima of affine functions of the d_{τ} s. However, working with variables that are piecewise linear functions of the data would lead to a very complicated robust counterpart. To overcome this difficulty and keep the problem tractable, we take a somewhat conservative approach by making the y_t s affine functions of the d_{τ} s. With this approach (which works both for the y_t s and for the u_t s, see (3e) and (3f)), we arrive at

$$y_t = y_t^0 + \sum_{\tau=1}^{t-1} y_{\tau}^{\tau} d_{\tau}, \\ u_t = u_t^0 + \sum_{\tau=1}^{t-1} u_{\tau}^{\tau} d_{\tau}. \quad (19)$$

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 the objective function and each of the constraints of problem (3) involving the uncertain demands. A guiding principle is to choose uncertainty sets for which the resultant AARC can be solved efficiently. From the theory

in Ben-Tal and Nemirovski (2000), we know that this will be the case if we choose the uncertainty set to be either a polyhedral or an ellipsoidal set. The resultant AARC will either be an LP or a conic-quadratic program. Both can be solved very efficiently using readily available software even for large-scale problem instances. To be concrete in what follows, we select a particular configuration of these two types of uncertainty sets.

We assume that the only data we have on each uncertain demand d_t is that it resides within a certain interval around a “nominal” demand value \bar{d}_t . To ensure that the “physical” constraints (like balance equations (3a) and bounds on the orders (3b–c)) will hold for any demand realization, we choose for all of the constraints a “box uncertainty” set

$$\mathcal{U}_{\text{box}} = \{d \in \mathbb{R}^T : |d_t - \bar{d}_t| \leq \rho G_t, t = 1, 2, \dots, T\}, \quad (20)$$

where the positive numbers G_t represent “uncertainty scale” and $\rho > 0$ is the “uncertainty level.” A particular case of interest is $G_t = \bar{d}_t$, which corresponds to a simple case where \mathcal{U}_{box} contains demands whose relative deviation from the nominal demand is of size up to ρ . An uncertainty set of type \mathcal{U}_{box} may be used to model situations of independent demand.

In contrast, when treating the objective function, it is not compulsory to provide an absolute guarantee of its optimal value for *all* possible demands in \mathcal{U}_{box} . For example, when the demands are independent random variables, the probability that they simultaneously occur at the “corners” of \mathcal{U}_{box} is very small; hence such rare events can be ignored quite safely. The ellipsoidal uncertainty set \mathcal{U}_{ell} can indeed cut the corners of \mathcal{U}_{box} . The general form of \mathcal{U}_{ell} is

$$\mathcal{U}_{\text{ell}} = \{d \in \mathbb{R}^T : (d - \bar{d})S^{-1}(d - \bar{d}) \leq \Omega^2\}, \quad (21)$$

where S is a $T \times T$ symmetric positive definite matrix and $\Omega \geq 0$ is a “safety parameter.”⁶

With the above choices of the linear decision rules (Equations (17)–(19)) and the uncertainty sets

⁶ Natural choices of \bar{d} and S when the demand vector is stochastic are $\bar{d} = E(d)$ and $S = \text{cov}(d)$ (see Ben-Tal and Nemirovski 2000, 2002).

(21)–(20), the AARC formulation corresponding to the LP RSFC problem (3) is finally the following:

$$\begin{aligned} & \min_{q_t^r, x_{t+1}^r, y_t^r, u_t^r, w_t, z_t, E} E \\ \text{s.t. } & E \geq \sum_{t=1}^T \left[c_t q_t^0 + y_t^0 + u_t^0 + z_t \right. \\ & \quad \left. + \sum_{\tau=1}^{t-1} [c_t q_t^\tau + y_t^\tau + u_t^\tau] d_\tau \right] \\ & \quad \forall d \in \mathcal{U}_{\text{ell}} \text{ and } \forall t = 1, \dots, T, \\ (*) & x_{t+1}^0 + \sum_{\tau=1}^t x_{t+1}^\tau d_\tau = x_t^0 + \sum_{\tau=1}^{t-1} x_t^\tau d_\tau + q_t^0 \\ & \quad + \sum_{\tau=1}^{t-1} q_t^\tau d_\tau - d_t, \\ & y_t^0 + \sum_{\tau=1}^{t-1} y_t^\tau d_\tau \geq \bar{h}_t \left[x_{t+1}^0 + \sum_{\tau=1}^{t-1} x_{t+1}^\tau d_\tau \right] \\ & \quad \forall d \in \mathcal{U}_{\text{box}}, \\ & y_t^0 + \sum_{\tau=1}^{t-1} y_t^\tau d_\tau \geq -p_t \left[x_{t+1}^0 + \sum_{\tau=1}^{t-1} x_{t+1}^\tau d_\tau \right] \\ & \quad \forall d \in \mathcal{U}_{\text{box}}, \\ & u_t^0 + \sum_{\tau=1}^{t-1} u_t^\tau d_\tau \geq \alpha_t^+ \left[q_t^0 + \sum_{\tau=1}^{t-1} q_t^\tau d_\tau - w_t \right] \\ & \quad \forall d \in \mathcal{U}_{\text{box}}, \\ & u_t^0 + \sum_{\tau=1}^{t-1} u_t^\tau d_\tau \geq -\alpha_t^- \left[q_t^0 + \sum_{\tau=1}^{t-1} q_t^\tau d_\tau - w_t \right] \\ & \quad \forall d \in \mathcal{U}_{\text{box}}, \\ & z_t \geq \beta_t^+ (w_t - w_{t-1}), \\ & z_t \geq -\beta_t^- (w_t - w_{t-1}), \\ & L_t \leq q_t^0 + \sum_{\tau=1}^{t-1} q_t^\tau d_\tau \leq U_t \quad \forall d \in \mathcal{U}_{\text{box}}, \\ & \hat{L}_t \leq \sum_{\tau=1}^t \left(q_\tau^0 + \sum_{\theta=1}^{\tau-1} q_\tau^\theta d_\theta \right) \leq \hat{U}_t \quad \forall d \in \mathcal{U}_{\text{box}}, \\ & q_t^\tau = 0, \quad (t, \tau) \in J_t. \end{aligned} \quad (22)$$

The equality constraints (*) above are used to eliminate the variables x_{t+1}^0 that are then substituted in the next two inequalities. Note that after this substitution, the problem retains its affine structure and consists only of inequalities.

REMARK. The first constraint in (22), which relates to the objective function of the AARC model, is of the form

$$(\diamond) \quad d^\top x \leq \alpha \quad \forall d \in \mathcal{U}_{\text{ell}}.$$

Because we assume that actually $d \in \mathcal{U}_{\text{box}}$, we may violate the actual robust constraint

$$(\diamond\diamond) \quad d^\top x \leq \alpha \quad \forall d \in \mathcal{U}_{\text{box}}.$$

Let x_{ell} be the solution of (\diamond) . Then, we would like to estimate the probability

$$p = \text{pr}\{d^\top x_{\text{ell}} \leq \alpha \quad \forall d \in \mathcal{U}_{\text{box}}\}.$$

As shown in Ben-Tal and Nemirovski (2001), when d_1, d_2, \dots, d_T are independent random variables with support \mathcal{U}_{box} and mean \bar{d}_t , then by choosing the diagonal matrix S with diagonal entries $\rho^2 G_1^2, \rho^2 G_2^2, \dots, \rho^2 G_T^2$, the above probability satisfies

$$p \geq 1 - \exp(-\Omega^2/2).$$

In particular, by taking $\Omega = 3$, the constraint (\diamond) is assured with probability $p \geq 0.989$.

3.3. Deriving a Computationally Tractable Optimization Problem Equivalent to the AARC Problem (22)

Problem (22) is intractable in the sense that it contains a continuum of constraints. In this section we show how to convert each such continuum into a few deterministic constraints that are either linear or conic quadratic. Let v denote the vector consisting of all the decision variables $q_t^\tau, x_{t+1}^\tau, y_t^\tau, u_t^\tau, w_t, z_t, E$ ($1 \leq t \leq T, 1 \leq \tau < t$). To denote the fact that a function f depends affinely on v , we write $f[v]$.

Clearly, each of the constraints in (22) has the general form

$$\lambda_0[v] + \sum_{t=1}^T \lambda_t[v] d_t \leq 0, \quad \text{either } \forall d \in \mathcal{U}_{\text{ell}} \text{ or } \forall d \in \mathcal{U}_{\text{box}}. \quad (23)$$

For the box uncertainty set \mathcal{U}_{box} , inequality (23) is equivalent to (we temporarily write $\lambda_t = \lambda_t[v]$)

$$\max_{d \in \mathbb{R}^T} \left\{ \lambda_0 + \sum_{t=1}^T \lambda_t d_t : |d_t - \bar{d}_t| \leq \rho G_t, \quad t = 1, \dots, T \right\} \leq 0. \quad (24)$$

The optimal solution of problem (24) is $d_t = \bar{d}_t + \text{sign}(\lambda_t) \rho G_t$, and so (24) becomes

$$\lambda_0 + \sum_{t=1}^T (\lambda_t \bar{d}_t + \rho G_t |\lambda_t|) \leq 0,$$

which can be further expressed by adding new variables η_1, \dots, η_T as in the following linear inequalities:

$$\begin{cases} \lambda_0 + \sum_{t=1}^T (\lambda_t \bar{d}_t + \rho G_t \eta_t) \leq 0, \\ -\eta_t \leq \lambda_t \leq \eta_t, \quad t = 1, \dots, T. \end{cases} \quad (25)$$

Recalling that each $\lambda_t = \lambda_t[v]$ is an affine function of v , it follows that (25) remains a linear system of inequalities in the original vector of variables v .

For the ellipsoidal uncertainty set \mathcal{U}_{ell} , inequality (23) is equivalent to

$$\max_{d \in \mathbb{R}^T} \left\{ \lambda_0 + \sum_{t=1}^T \lambda_t d_t : (d - \bar{d})^\top S^{-1} (d - \bar{d}) \leq \Omega^2 \right\} \leq 0. \quad (26)$$

Problem (26) can be solved easily using the Karush-Kuhn-Tucker conditions; the optimal d is given by

$$d = \bar{d} + \frac{\Omega}{\sqrt{\lambda^\top S \lambda}} S \lambda,$$

and so (26) becomes

$$\lambda_0 + \sum \lambda_t \bar{d}_t + \Omega \sqrt{\lambda^\top S \lambda} \leq 0; \quad (27)$$

this is a conic-quadratic constraint in the λ_t s. When making an affine substitution of variables $\lambda = \lambda[v]$, (27) remains a conic-quadratic inequality in the original decision vector v .

When the demand vector d is a random vector with mean \bar{d} and covariance matrix S , inequality (27) has an intuitively appealing explanation. Note that the right-hand side of (23) is a random variable $\Phi[v] \equiv \lambda_0[v] + \sum_{t=1}^T \lambda_t[v] d_t$ with mean $E[\Phi[v]] = \lambda_0 + \sum \lambda_t \bar{d}_t$ and standard deviation $SD[\Phi[v]] = \sqrt{\lambda^\top S \lambda}$. Thus, the robust counterpart (27) is

$$E[\Phi[v]] + \Omega SD[\Phi[v]] \leq 0. \quad (28)$$

The latter corresponds to a common engineering practice, where an uncertain inequality $\Phi[v] \leq 0$ is replaced by its “safe” version (28) (typically, with $\Omega \approx 3$).

The analysis we have carried out shows that each of the semi-infinite constraints in the AARC problem (22) 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., MOSEK⁷) are available to find their solution. In particular, if *all* the uncertainty sets are of type \mathcal{U}_{box} (or, more generally, a polyhedral set), then the AARC is an LP.

4. Analysis of the RSFC Model

4.1. Preliminaries

In this section, we provide a summary of a large set of experiments that were generated to test the performance of the AARC method, as applied to the RSFC problem, on simulated data. We benchmark the AARC model against the RC solution (12), the optimal min-max solution (5), and a “perfect hindsight” (PH) solution⁸—the solution that would have been obtained to the LP RSFC model (3) if it was possible to “reverse the time” and know the realizations of demand at the beginning of the horizon. The PH solution is clearly the *ultimate* lower bound on the minimal cost of the AARC method (or any other solution method for that matter). We first compare the AARC and RC solutions to the true *optimal min-max RSFC solution* (opt(min-max) for short). Then, we test the performance of the AARC method by comparing its mean cost to the mean cost of the PH solution where both means are computed over the *same* simulated realizations of the demands. Then, we discuss the resultant LDRs in the AARC solutions and analyze other aspects of the AARC methodology.

The most interesting question is, of course, how far is the AARC solution from the opt(min-max) solution? This comparison is possible in the case where the uncertainty set \mathcal{U}^T is equal to $\mathcal{U}_{\text{box}} = \mathcal{U}_1 \times \mathcal{U}_2 \times \dots \times \mathcal{U}_T$ with $\mathcal{U}_t = [(1 - \theta)\bar{d}_t, (1 + \theta)\bar{d}_t]$, \bar{d}_t being the nominal demand at period t . Indeed, as explained in §2.3, in this case opt(min-max) is obtained by

solving a large LP associated with the 2^T extreme demand trajectories. The AARC solution, on the other hand, is obtained by solving problem (22) (but with \mathcal{U}_{ell} replaced by \mathcal{U}_{box} in the first constraint), which is an LP as well, but of a much smaller size⁹ (for $T = 12$ it has 417 variables, compared to 16,405 variables in the LP associated with opt(min-max)). We also obtained the solution of the RC problem (12) to test the added advantage of using the *dynamic* AARC solution compared to the *static* RC solution.

Our experiments assume that actual demand fluctuations are restricted by a box uncertainty set. In each experiment, we fixed the parameters associated with the costs, penalties, upper and lower bounds on order quantities, the horizon T , the initial inventory x_1 , the initial commitment w_0 , the parameters describing the uncertainty sets \mathcal{U}_{box} and \mathcal{U}_{ell} , and the set J of indices corresponding to time periods that are not used in the linear decision rule (see (17)). We generated more than 300 data sets; most of them were randomly generated, while some (e.g., D2, W12) were determined manually. Data set W12 was created to illustrate a case where there is a marked difference between the outcomes of the RC and the AARC models (see §4.2). Data sets A12 and A10 were chosen from the randomly generated data sets to demonstrate instances in which the optimal LDRs rely on the entire demand history (see §4.4) and to compare an optimal base stock policy that is an LDR with an optimal AARC solution (see §4.5), respectively. These data sets are given in Table 1.

The first step in each experiment is to solve the opt(min-max), RC, and AARC models. For the latter we obtain the coefficients of the various LDRs. Then, we simulate demand scenarios, consistent with our box uncertainty assumption $d \in \mathcal{U}_{\text{box}}$, using either a uniform or a truncated normal distribution with support in \mathcal{U}_{box} . For each simulation s we record the realized demand vector $d^s = (d_1^s, \dots, d_T^s)$, employ the optimal LDRs, and compute the resultant cost. Finally, we compute the mean and standard deviation of these costs over all the simulated scenarios. Then, we solve the LP model (3) with the demand realizations d^s , obtain the PH solution ((w^s, q^s) pairs and the optimal

⁷ See <http://www.mosek.com/>.

⁸ See Talluri and van Ryzin (1998, p. 1,584).

⁹ The number of variables in (LP) AARC is $2T(T + 1) + 2T + 1$, and the number of variables in (LP) opt(min-max) is $4(2^T - 1) + 2T + 1$.

Table 1 A Sample of Data Sets


Parameter	A12	D2	W12	A10
T	12	12	12	10
x_1	57	100	0	18.895
w_0	12	100	100	19.9613
c_t	1.01	40	10	0.1818
h_t	0.3	2	2	4.6194
p_t	1.0	5	10	1.193
s_{T+1}	1.13	7	0	0.2195
α_t^+	0.43	2	10	1.204
α_t^-	0.58	3	10	2.11
β_t^+	0.37	1	10	0.4308
β_t^-	0.04	2	10	0.6278
L_t	44	50	0	14.9948
U_t	76, 54, 66, 88, 68, 60, 82, 53, 53, 78, 72, 63	200	200	50.7185
\hat{L}_t	0	0	0	0
\hat{U}_t	814	∞	$200t$	∞
\bar{d}_t	64	100	100	46

cost), and compute the mean and standard deviation over all the simulated demands.

Solving the AARC problem once with \mathcal{U}_{ell} for the objective function and \mathcal{U}_{box} for all other uncertain constraints over a horizon of 12 or 24 periods, where in each period the LDRs capture the entire past demands up to that period, leads to a mathematical program (conic quadratic or LP) with several hundred to a few thousand variables/constraints. To run such programs efficiently thousands of times so as to compare them in different data settings, we developed the robust commitment optimization (RobCop) package—a dedicated software that employs the Matlab and MOSEK software packages at its core. RobCop is composed of two components. The first component employs a user-friendly interface to assist users in constructing data sets with various cost parameters, demand uncertainty sets, etc. The second component solves the AARC, runs the simulations, computes the mean AARC and PH solutions, and reports the results. A typical experiment with 100 simulation replications was executed by RobCop on a Pentium 4 laptop computer in less than 20 seconds.

A typical outcome of the RSFC model is depicted in Figure 1 (such figures are automatically generated by our software). The particular outcome depicted in Figure 1 corresponds to a specific experiment with

Figure 1 Outcome of a Typical AARC Model Run (• Bounds on Demand, □ Realized Demand, * Orders, × Commitments, ○ Inventory, * Total Costs Per Period)



a 24-period horizon, with nominal demand that follows an annual cyclical pattern with peak demand at the middle of each year. The demand uncertainty starts at $\pm 20\%$ and increases by 10% every six months. We see that the realized demand fluctuates considerably within its limiting boundaries. Consequently, the retailer’s inventory, whose initial value was $x_1 = 100$, fluctuates quite significantly (dropping to about 35 in Months 3–4 and climbing to about 250 in Month 20). The commitments were nearly always at the upper bound on possible demand realizations. Toward the end of the 24-period horizon, they sharply drop (to avoid unnecessary surplus at the end of the horizon). Although the optimal orders generally follow the changes in the realized demands, they do so in a more moderate fashion (due to the penalties imposed by the contract). This demonstrates that the flexible commitment contract lives up to its expectations—although the problem was solved from the point of view of the retailer, the supplier enjoys a more stable sequence of commitments and actual orders with relatively small fluctuations.

4.2. Comparing the AARC and RC Solutions to the Opt(Min-Max) Solution

The comparison of the three alternative solutions is demonstrated in Table 2 for data sets D2, A12, and W12 over various levels of uncertainty. In data

Table 2 Opt(Min-Max), AARC, and RC Solutions for Data Sets *A12*, *D2*, and *W12* (in Parentheses: Excess Over the Opt(Min-Max) Solution)

Data	Uncertainty		AARC	RC
	(in %)	Opt(min-max)		
<i>D2</i>	10	40,750.0	40,750.0 (+0.0%)	40,750.0 (+0.0%)
	20	44,150.0	44,150.0 (+0.0%)	44,150.0 (+0.0%)
	30	47,550.0	47,550.0 (+0.0%)	47,550.0 (+0.0%)
	40	50,950.0	50,950.0 (+0.0%)	50,950.0 (+0.0%)
	50	54,350.0	54,350.0 (+0.0%)	54,350.0 (+0.0%)
	60	57,760.0	57,760.0 (+0.0%)	57,760.0 (+0.0%)
	70	61,170.0	61,170.0 (+0.0%)	61,170.0 (+0.0%)
<i>A12</i>	10	913.128	913.128 (+0.0%)	1,002.941 (+9.8%)
	20	1,397.440	1,397.440 (+0.0%)	1,397.440 (+0.0%)
	30	2,190.620	2,190.620 (+0.0%)	2,190.620 (+0.0%)
	40	3,087.540	3,087.540 (+0.0%)	3,087.540 (+0.0%)
	50	4,006.040	4,006.040 (+0.0%)	4,006.040 (+0.0%)
	60	4,934.680	4,934.680 (+0.0%)	4,934.680 (+0.0%)
	70	5,863.320	5,863.320 (+0.0%)	5,863.320 (+0.0%)
<i>W12</i>	10	13,531.8	13,531.8 (+0.0%)	15,033.4 (+11.1%)
	20	15,063.5	15,063.5 (+0.0%)	18,066.7 (+19.9%)
	30	16,595.3	16,595.3 (+0.0%)	21,100.0 (+27.1%)
	40	18,127.0	18,127.0 (+0.0%)	24,300.0 (+34.1%)
	50	19,658.7	19,658.7 (+0.0%)	27,500.0 (+39.9%)
	60	21,190.5	21,190.5 (+0.0%)	30,700.0 (+44.9%)
	70	22,722.2	22,722.2 (+0.0%)	33,960.0 (+49.5%)

sets *D2* and *A12* there was little or no change among the three solutions. This phenomenon was quite common to many of the *randomly* generated data sets. However, there are data sets such as *W12* in which the nonadaptive RC solution yields a much larger cost estimate than the corresponding AARC solution. In this data set, the RC's deviation from the optimal min-max solution (which appears in parentheses in the table) reaches nearly 50% when the uncertainty fluctuates within 70% above or below nominal demand.

The most interesting part in Table 2 is, no doubt, the comparison between AARC and opt(min-max) solutions revealing that the AARC solution is identical to the opt(min-max) solution across all uncertainty levels in all three data sets. In fact, in only 4 out of 300 randomly generated data sets did we find any deviations between these two solutions (the largest deviation being 4%).

Although there are no analytical results on the structure of the optimal solution to the RSFC problem, and particularly there is no theoretical foundation for the optimal solution being an LDR, our empirical results strongly support the fact that

(at least for box uncertainty) the optimal solution is well-approximated by the LDR found by our AARC method.¹⁰

4.3. Comparing Mean AARC and PH Performance

The AARC solution is in fact a conservative cost estimator. To evaluate the *actual* outcomes that might result from employing this model, we ran hundreds of simulations with different data sets and compared the mean performance vis-à-vis the mean PH outcome. Table 3 reports such results for data set *W12*. The mean AARC cost (over the simulation runs) is clearly lower than the AARC solution itself. Moreover, the deviation increases with the uncertainty levels. The PH solution yields the lowest mean cost, which is significantly lower than the AARC mean costs because it solves the problem to optimality under conditions of perfect information. Each row in Table 3 corresponds to 100 simulations of a scenario in which demand was allowed to fluctuate around the nominal demand $\bar{d} = 100$ according to the ρ value given in the left-most column. The second column from the left reports the AARC solutions (where only q_1 and the commitments w_i are predetermined while the solution for the other variables is expressed via optimal LDRs). The next column gives the mean and standard deviation of the AARC solutions (i.e., where the policies determined by the AARC solutions were actually employed according to the realized demand), and the fourth column gives the corresponding PH solutions over the same 100 simulations. As expected, because the mean demand stays fixed, the PH solutions change very little with the variability of demand. Note that the difference between the mean simulation result and the corresponding AARC solution grows from 1.2% to 4.8% as the size of demand fluctuations grows from 10% to 70%. That is, one can expect larger relative "savings" (actual AARC cost versus the original AARC solution) as the uncertainty becomes larger. Note also that the gap between the conservative estimate of the AARC solution (which protects against the *worst* case) and the PH solution increases significantly with demand variability.

¹⁰ Admittedly, this may not necessarily be the case for criteria other than the min-max, such as minimizing mean cost.

Table 3 Simulated AARC and PH Costs with Data $W12$ (% in Parentheses—Excess Over Opt(Min-Max))

Uncertainty (in %)	Opt(AARC) = opt(min-max)	Simulated AARC = simulated(min-max) mean(%): STD	Simulated PH mean(%): STD
10	13,531.8	13,375.4 (−1.2%): 41.1	12,421.1 (−8.2%): 1,406.2
20	15,063.5	14,745.4 (−2.1%): 85.9	12,725.4 (−15.5%): 1,816.9
30	16,595.3	16,122.8 (−2.8%): 124.2	13,283.8 (−20.0%): 2,421.1
40	18,127.0	17,477.7 (−3.6%): 170.0	13,626.8 (−24.8%): 2,624.9
50	19,658.7	18,858.2 (−4.1%): 206.7	13,420.6 (−31.7%): 2,447.4
60	21,190.5	20,267.3 (−4.4%): 235.6	13,653.5 (−35.6%): 2,363.1
70	22,722.2	21,642.3 (−4.8%): 286.8	14,079.4 (−38.0%): 2,892.2

4.4. Realized Patterns in the LDRs

The AARC formulation gives the user the flexibility to determine which part of the “history” of realizations of the uncertain data will be included in the LDRs (as given in (17)). The extreme approach is to allow *all* the uncertain data that were realized at the time of decision to be included, but this does not necessarily mean that the optimal LDRs will use this entire

history. So, one of the issues we wanted to explore in these experiments is the question of whether the resultant LDRs follow some consistent patterns. In particular, we were curious to see whether the resultant LDRs exhibit a “Markovian” behavior, i.e., the LDR for a variable in period t (say q_t) depends *only* on d_{t-1} —the demand in the previous period. Table 4 demonstrates two opposite outcomes for the LDRs.

Table 4 Optimal LDRs for Data Sets $W12$ and $A12$

$\begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \\ q_5 \\ q_6 \\ q_7 \\ q_8 \\ q_9 \\ q_{10} \\ q_{11} \\ q_{12} \end{bmatrix} =$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.95 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.92 \end{bmatrix} \cdot$	$\begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \\ d_9 \\ d_{10} \\ d_{11} \\ d_{12} \end{bmatrix} +$	$\begin{bmatrix} 120 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ -1.59 \\ 1.60 \\ 3.56 \end{bmatrix}$
(a)			
$\begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \\ q_5 \\ q_6 \\ q_7 \\ q_8 \\ q_9 \\ q_{10} \\ q_{11} \\ q_{12} \end{bmatrix} =$	$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.25 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.18 & 0.39 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.14 & 0.27 & 0.73 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.02 & 0.03 & 0.03 & 0.55 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.02 & 0.02 & 0.02 & 0.04 & 0.31 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.02 & 0.03 & 0.02 & 0.05 & 0.22 & 0.64 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.01 & 0.01 & 0.02 & 0.02 & 0.03 & 0.03 & 0.12 & 0 & 0 & 0 & 0 & 0 \\ 0.01 & 0.01 & 0.02 & 0.03 & 0.02 & 0.02 & 0.05 & 0.07 & 0 & 0 & 0 & 0 \\ 0.03 & 0.03 & 0.03 & 0.04 & 0.06 & 0.05 & 0.16 & 0.20 & 0.28 & 0 & 0 & 0 \\ 0.02 & 0.02 & 0.03 & 0.04 & 0.06 & 0.06 & 0.10 & 0.11 & 0.12 & 0.14 & 0 & 0 \\ 0.03 & 0.03 & 0.03 & 0.04 & 0.05 & 0.05 & 0.06 & 0.04 & 0.04 & 0.05 & 0.07 & 0 \end{bmatrix} \cdot$	$\begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \\ d_9 \\ d_{10} \\ d_{11} \\ d_{12} \end{bmatrix} +$	$\begin{bmatrix} 76.02 \\ 32.95 \\ 19.04 \\ -7.16 \\ 16.16 \\ 26.31 \\ -0.07 \\ 33.69 \\ 33.69 \\ 5.50 \\ 12.46 \\ 22.89 \end{bmatrix}$
(b)			

The table shows optimal LDRs of the orders q_1, \dots, q_T generated under 30% uncertainty for two data sets. The top half of the table (Part a), which corresponds to data set W12, indeed exhibits a Markovian behavior as defined above, while the bottom half of the table (Part b), which corresponds to data set A12, reveals a set of LDRs that are determined by *all* the demand realizations prior to the time of decision. Note, however, that the weights are decreasing in time (i.e., recent periods have a greater impact on q_t than more distant ones).

4.5. Base Stock Policies and LDR

The optimal LDR solution to data set W12 in Table 4 bears some resemblance to a base stock (BS) solution: In the first 10 periods $q_t = d_{t-1}$ ($d_0 = 120$) that implies the BS policy $q_t = 120 - x_t$. However, we note that due to end-of-horizon effects, which characterize all finite horizon problems, the optimal LDR found for data set W12 deviates from the BS pattern in the last three periods. Moreover, there is no theoretical foundation to assume that a BS policy is an LDR. Indeed, let S_t denote the BS level for period t . Then, the order quantity in period t is given by $q_t = (S_t - x_t)^+$ —a nonlinear equation. Further, substituting q_t in the inventory balance equations $x_{t+1} = x_t + q_t - d_t$, we get $x_{t+1} = \max\{S_t, x_t\} - d_t$. Thus, even in the deterministic case (let alone in the stochastic case) a BS policy is not necessarily an LDR. In any event, if it happens to be an LDR, its optimal BS levels will be found by the AARC method because the latter recovers the optimal LDR policy.

For a BS policy to be *linear*, it must have the form $q_t = S_t - x_t$. In this case it is an LDR with a very specific form as described in the following lemma.

LEMMA 1. *A BS policy is an LDR iff*

$$q_t = q_0^t + d_{t-1}, \tag{29}$$

and in this case the BS levels are given by

$$S_t = x_0 - q_0^t + \sum_{\tau=1}^{t-1} q_0^\tau.$$

PROOF. Consider the BS policy

$$q_t = S_t - x_t. \tag{30}$$

An LDR for the orders was given in (17) by $q_t = q_0^t + \sum_{\tau=1}^{t-1} q_\tau^t d_\tau$.

Now, recall the inventory balance equation (2a), which can be rewritten as

$$x_t = x_0 + \sum_{\tau=1}^{t-1} q_\tau - \sum_{\tau=1}^{t-1} d_\tau. \tag{31}$$

Substituting (31) in (30) and equating the outcome to (17) yields

$$q_0^t + \sum_{\tau=1}^{t-1} q_\tau^t d_\tau = S_t - x_0 + \sum_{\tau=1}^{t-1} d_\tau - \sum_{\tau=1}^{t-1} \left[q_0^\tau + \sum_{r=1}^{\tau-1} q_r^\tau d_r \right]. \tag{32}$$

Equating the coefficients of d_t on both sides of (32), we get the system of equalities

$$\begin{aligned} \forall t \geq 1, \quad q_{t-1}^t &= 1, \\ 1 \leq r \leq t-2, \quad q_r^t &= 1 - \sum_{\tau=r+1}^{t-1} q_\tau^r. \end{aligned} \tag{33}$$

It can be easily seen that the solution of the system (33) is

$$q_\tau^t = \begin{cases} 1 & \text{if } \tau = t-1, \\ 0 & \text{if } 1 \leq \tau \leq t-2, \end{cases}$$

which proves the first part of the lemma.

Also, by equating the constant terms in (32), $q_0^t = S_t - x_0 - \sum_{\tau=1}^{t-1} q_0^\tau$, we can express S_t as a function of the constants x_0 and q_0^τ that yields the second part of the lemma. □

Now, in our setting, the constraint $L_t \leq q_t \leq U_t$ is expressed by

$$\begin{aligned} L_t &\leq q_0^t + d_{t-1} \leq U_t \\ \forall d_{t-1} \in [d_{\min}, d_{\max}] &\equiv \mathcal{U}_{\text{box}}(\rho), \end{aligned} \tag{34}$$

where $d_{\min} = (1 - \rho) \cdot \bar{d}$, $d_{\max} = (1 + \rho) \bar{d}$, and \bar{d} is the nominal demand. Clearly, (34) is equivalent to

$$L_t - (1 - \rho) \bar{d} \leq q_0^t \leq U_t - (1 + \rho) \bar{d}, \tag{35}$$

which implies

$$\rho \leq \frac{U_t - L_t}{2\bar{d}} \equiv \rho_{\text{critical}}.$$

Hence, a BS policy will be feasible iff the parameter ρ for $\mathcal{U}_{\text{box}}(\rho)$ is smaller than ρ_{critical} .

Numerical Example. The following example demonstrates that the solution obtained from a BS policy, which is LDR could be substantially worse than the *optimal* LDR solution obtained through the AARC method. Specifically, for data set A10, we obtain $\rho_{\text{critical}} = 0.388$. Running this data with $\rho = 0.35$, we get a feasible BS solution, 1,022, while the optimal AARC solution is 880.

Furthermore, running the same problem with $\rho = 0.4 > \rho_{\text{critical}}$ indeed yields no feasible solution for a BS policy, while the AARC solution is 1,018.

4.6. Box vs. Ellipsoidal Uncertainty Sets

All the scenarios we analyzed in Table 2 assumed \mathcal{U}_{box} for both the objective function and the constraints. In the next set of experiments we wanted to test the effect of replacing the uncertainty set for the objective function with \mathcal{U}_{ell} . This means that when demand is indeed generated from a box uncertainty set, the objective function may occasionally underestimate the real cost. Table 5 compares the results achieved with ellipsoidal uncertainty sets (with various values of Ω as explained in (26)) to the corresponding results for a box uncertainty set for data set A12. The interesting column in this table is the one on the right (where $\Omega = 3$ implies that the probability of underestimating the worst-case cost is about 1%). There we see that for large uncertainty sets (exceeding 50%) the cost estimate given under the assumption of ellipsoidal uncertainty is smaller by about 1.3% from that provided under the assumption of box uncertainty.

4.7. Effects of Information Gaps

One of the advantages of our proposed approach is the flexibility it offers in constructing the LDRs from

historical data. As explained in (17), the set J_t indicates past periods whose data is *not* included in the LDR. In this section we investigate the impact of various configurations of J_t on the outcomes of the AARC method. Clearly, we expect the performance of the AARC method to deteriorate as more periods are included in J_t . In the extreme case, when J_t covers all past periods, the AARC reduces to the RC model (12). By analyzing the changes in the “price of uncertainty” (relative difference between the mean PH and mean AARC solutions), we give management a useful tool to evaluate the value of information from various periods. In other words, we provide management with an upper bound on the amount it would be willing to pay to obtain additional information when some information elements are missing.

First, we test what happens when at any period t we are unable to use data from the periods that immediately preceded it (say, periods $t - 1$ and $t - 2$). Such situations may occur when information is delayed due to data verification procedures, regulatory requirements, or other organizational constraints. We ran two of the uncertainty scenarios ($\pm 30\%$ and $\pm 70\%$) that are shown in Table 2 with $J_t = \{t - 1, t - 2\}$ over 30 simulations and report the results in Table 6. Comparing the relevant values in Tables 2 and 6, we observe that the price of uncertainty nearly doubled in these two scenarios. Second, we ran two experiments, each with 30 simulations, corresponding to the same two uncertainty scenarios, in which any information that is older than three periods is not used (that is, $J_t = \emptyset$ for $t = 1, 2, 3, 4$ and $J_t = \{1, 2, \dots, t - 4\}$ for $t > 4$). Such situations may occur due to short product life cycles, highly seasonal demand, or other organizational constraints. Our findings, summarized in Table 7, indicate that although there was an increase in the price of uncertainty as compared to the case where $J_t = \emptyset$, the effect is smaller than the one that occurred when the most recent two periods were excluded. This result suggests that the value of information for the AARC

Table 5 Ellipsoidal vs. Box Uncertainty Sets with Data A12
 (in Parentheses: Excess Over Box)

Uncertainty (in %)	Opt(AARC), box uncertainty	Opt(AARC), ellipsoidal uncertainty	
		$\Omega = 1$	$\Omega = 3$
0	840.732	840.732 (+0.0%)	840.732 (+0.0%)
10	1,014.143	983.942 (−3.0%)	1,014.143 (−0.0%)
20	1,300.212	1,224.414 (−5.8%)	1,300.212 (−0.0%)
30	1,720.621	1,697.849 (−1.3%)	1,718.693 (−0.1%)
40	2,326.479	2,245.172 (−3.5%)	2,313.209 (−0.6%)
50	2,969.785	2,813.708 (−5.3%)	2,942.431 (−0.9%)
60	3,629.710	3,392.953 (−6.5%)	3,587.161 (−1.2%)
70	4,289.636	3,972.198 (−7.4%)	4,231.891 (−1.3%)

Table 6 Effects of Delayed Information—Data Set W12

ρ ($\pm\%$)	AARC solution	Mean AARC results	Mean PH results	Price of uncertainty (%)
30	17,984	16,731	12,587	32.9
70	26,044	22,858	14,135	61.7

Table 7 Effects of Discarding Old Information—Data Set W12

ρ ($\pm\%$)	AARC solution	Mean AARC results	Mean PH results	Price of uncertainty (%)
30	16,595	16,108	12,712	26.7
70	22,722	21,582	14,484	49.0

method is period-dependent with higher values associated with more recent periods.

4.8. Trade-offs Analysis

The RSFC model we have developed might be used during *contract negotiations* by either the retailer or the supplier in an attempt to identify fair combinations of discretionary data parameters that will best serve their wishes vis-à-vis the handling of demand variability. To demonstrate how such trade-offs 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 α_t while the overall cost to the retailer remained stable. Table 8 presents the results of these runs for data set D2 with $\pm 20\%$ uncertainty. Each run consisted of 10 simulations, each with a cycle of 12 periods. The AARC solution in each one of the runs reported in this table was $39,620 \pm 20$. The penalties for both upper and lower deviations were kept identical throughout these runs, i.e., $\alpha_t = \alpha_t^+ = \alpha_t^- \forall t$. The table starts with total mean deviations per cycle $(q_t - w_t)_+ + (w_t - q_t)_+ = 92.32$ (equivalent to a mean deviation of 7.7 per period). To cut this value by nearly a half (to 41.81 per cycle) we needed to double the penalty on these deviations. To maintain the same total cost, the purchase cost was reduced from 40 to 39.8. This pattern continued until the penalty costs were set at $\alpha_t^+ = \alpha_t^- = 18$ where the deviations became zero.

A similar trade-off may exist between purchase cost and minimum order quantities. The lower bounds L_t in the RSFC model (3) play a crucial role in the

Table 8 Trade-offs Between Purchase Costs and Forecast Error Penalties—Data Set D2

c_t	α_t	Mean $(q_t - w_t)_+$	Mean $(w_t - q_t)_+$
40.0	2	21.07	71.25
39.9	3	44.84	12.23
39.8	4	19.36	22.45

supplier-retailer relations. From the supplier’s perspective, it is desirable to see large values of these minimum order quantities. Conversely, from the retailer’s point of view, the larger these bounds are, the less flexibility he has and the larger the probability that he will be stuck with unnecessary inventories. We demonstrate this trade-off by searching for 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 expense). To make the comparison valid, we generated for the same data set a single realization of demand for a path of 10 cycles with 12 periods in each cycle and used it throughout this experiment. The trade-off is shown in Table 9. Initially, with a purchase cost of 40 and a minimum order quantity of 50, the retailer’s mean total order per cycle is 1,130 and his mean total cost is 47,837. 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 1,150 items per cycle (thus increasing his revenue from 45,233 to 45,457) while causing a slight decrease in the retailer’s total cost (to 47,312). The same pattern continues until we reach a near perfect equilibrium in which the total cost to the retailer is nearly equal to the supplier’s revenues. Thus, the supplier is able in this case to generate a “win-win” situation in which both he and the retailer will gain by the simultaneous change in the parameters.

4.9. A Folding Horizon Approach

The AARC method considered thus far can be classified as “offline” in the sense that the AARC problem is solved once at the beginning of the planning horizon (i.e., the optimal commitments and the optimal coefficients of the LDRs are found a priori).

Table 9 Trade-offs Between Purchase Costs and Minimum Order Quantities

c_t	L_t	$\sum q_t$	Purchase cost (supplier’s revenue)	Total cost (retailer’s expenses)
40	50	1,130	45,233	47,837
39.5	60	1,150	45,457	47,312
39	70	1,180	46,052	46,826
38.8	80	1,191	46,204	46,608

In this section, we describe an “online” version of the AARC method that we coin the *folding horizon AARC model*. A retailer using this method re-solves in each period t an AARC problem for the remaining periods $t, t + 1, \dots, T$, starting from initial inventories that are the actual ones, i.e., those resulting from his earlier decisions and the demand realizations that occurred in periods $1, 2, \dots, t - 1$. The commitments are still determined once in the first period for the entire horizon.

Folding horizons similar to the one we explore here can be found in the *revenue management* literature (e.g., Talluri and van Ryzin 1998, Feng and Gallego 2000) dealing with demand for seats on flights, hotel rooms, theater tickets, etc. In these settings, there exists a fixed deadline after which the “goods” are lost. Pricing is then done in a folding horizon manner—an initial price is set at the beginning of the horizon and then, as demand realizations are observed, new pricing policies are generated. Folding horizons are also 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.

Implementation of the folding horizon approach 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 shorten the horizon by one period, fix the commitments according to their optimal values in the first run, and fix the starting inventory according to the realization of demand and order replenishment from the previous period. This procedure is repeated until the entire original horizon is exhausted.

The folding horizon approach allows the retailer to incorporate more accurate information each time he solves the problem. Hence, we can only expect it to improve the results obtained through the fixed horizon approach. To compare the performance of these two model variations, we ran the RobCop software with data set W12 with box uncertainty set for both the objective function and the constraints over various uncertainty levels with 100 simulations for each

Table 10 Comparison Between Fixed and Folding Horizon Models with Data Set W12 (in Parentheses—% Excess Over AARC)

Uncertainty (in %)	Opt(AARC)	Simulated AARC fixed	Simulated AARC folding
10	13,531.8	13,375.4 (−1.2%): 41.1	13,372.5 (−1.2%): 40.7
20	15,063.5	14,745.4 (−2.1%): 85.9	14,742.5 (−2.1%): 85.5
30	16,595.3	16,122.8 (−2.8%): 124.2	16,115.2 (−2.9%): 126.6
40	18,127.0	17,477.7 (−3.6%): 170.0	17,463.7 (−3.7%): 173.5
50	19,658.7	18,858.2 (−4.1%): 206.7	18,847.8 (−4.1%): 209.2
60	21,190.5	20,267.3 (−4.4%): 235.6	20,261.2 (−4.4%): 228.7
70	22,722.2	21,642.3 (−4.8%): 286.8	21,632.8 (−4.8%): 280.4

uncertainty value. Table 10 reports the results of these runs. To ensure fair comparisons, the realizations of demand that were generated for each line of the fixed horizon problems in Table 3 were stored in memory and used for the corresponding simulation of the folding horizon model. Thus, the demand data is identical within each line of these two tables, but it is different across different lines. For each run, Table 10 reports the mean and standard deviation values for the folding horizon and the fixed horizon models over the 100 simulations. As expected, the mean folding horizon solution is better than the mean fixed horizon solution. However, at least for data set W12, the difference between them is almost negligible.

5. Extensions and Future Research

The models we have presented in the preceding sections can be extended in several ways. First, as in Anupindi and Bassok (1998), the model can be extended to a multiproduct setting. To do so, suppose we now have N products and add the index i to designate product i ($i = 1, \dots, N$) to the relevant parameters and decision variables. Then, the AARC formulation that replaces (2) is

$$\min_{x, q, w} \left\{ - \sum_{i=1}^N s_i \max[x_{i, T+1}, 0] + \sum_{i=1}^N \sum_{t=1}^T [c_{it} q_{it} + h_{it} \max[x_{i, t+1}, 0] + p_{it} \max[-x_{i, t+1}, 0] + \alpha_{it}^+ \max[q_{it} - w_{it}, 0] + \alpha_{it}^- \max[w_{it} - q_{it}, 0] + \beta_{it}^+ \max[w_{it} - w_{i, t-1}, 0] + \beta_{it}^- \max[w_{i, t-1} - w_{it}, 0]] \right\}$$

$$\begin{aligned}
 \text{s.t. (a): } & x_{i,t+1} = x_{it} + q_{it} - d_{it}, \\
 & i = 1, \dots, N; t = 1, \dots, T, \\
 \text{(b): } & L_{it} \leq q_{it} \leq U_{it}, \quad i = 1, \dots, N; t = 1, \dots, T, \\
 \text{(c): } & \hat{L}_{it} \leq \sum_{\tau=1}^t q_{i\tau} \leq \hat{U}_{it}, \\
 & i = 1, \dots, N; t = 1, \dots, T. \quad (36)
 \end{aligned}$$

To convert (36) into a linear programming formulation, we again replace the piecewise linear terms in the objective by auxiliary variables and obtain

$$\begin{aligned}
 \min_{x, q, w} \max_{y, z} & \left\{ \sum_{i=1}^N \sum_{t=1}^T [c_{it}q_{it} + y_{it} + u_{it} + z_{it}] \right\} \\
 \text{s.t. (a), (b), (c), and } & \forall i = 1, \dots, N, t = 1, \dots, T, \\
 \text{(e}_1\text{): } & y_{it} \geq h_{it}x_{it+1}, \\
 \text{(e}_2\text{): } & y_{it} \geq -p_{it}x_{it+1}, \\
 \text{(f}_1\text{): } & u_{it} \geq \alpha_{it}^+(q_{it} - w_{it}), \\
 \text{(f}_2\text{): } & u_{it} \geq -\alpha_{it}^-(q_{it} - w_{it}), \\
 \text{(g}_1\text{): } & z_{it} \geq \beta_{it}^+(w_{it} - w_{it-1}), \\
 \text{(g}_2\text{): } & z_{it} \geq -\beta_{it}^-(w_{it} - w_{it-1}), \\
 & \text{for } t = T, \text{ (e}_1\text{) becomes} \\
 & y_{it} \geq (h_{it} - s_i)x_{it+1}, \quad (37)
 \end{aligned}$$

where, as before, y_{it} represents the holding and shortage costs, u_{it} represents the forecast error penalties, and z_{it} represents the commitments' inconsistency penalties.

The AARC model for problem (37), using LDRs over the entire history of realizations of the adjustable variables $\{x_{it}, y_{it}, u_{it}, q_{it}, t = 1, \dots, T, i = 1, \dots, N\}$, will have $2NT(T + 1) + 2NT$ variables so it can handle many products with long planning horizons. In contrast, solving such problems with DP or with the "large LP" method we developed in §2.3 will quickly become computationally intractable as N and T grow.

Second, new and innovative formulations of the relations between suppliers and retailers can be tested. These might include larger chains (perhaps with multiple suppliers and/or retailers), new types of flexible arrangements to govern the relationships between the parties (e.g., different kinds of penalties

on deviations from announced commitments), and more. For example, an interesting variation of the RSFC model is one 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 (3). The terms that involve β_t^+, β_t^- disappear, and the terms that involve α_t^+, α_t^- are taken out of the summation and replaced by $\alpha^+[\sum_{t=1}^T q_t - W]_+ + \alpha^- [W - \sum_{t=1}^T q_t]_+$. The rest of the steps in transforming this adapted formulation to an AARC model are identical to the ones shown above.

Third, additional uncertain parameters may be considered. For example, it is reasonable to assume that future purchase costs are uncertain or that the size of actual shipments may be different than the size of the corresponding orders. This possibility was already observed by Bassok et al. (1997) who list *quality* among the key elements in a typical supply contract. They define quality as "the upper limit on the percentage of defective units in the supply that is acceptable to the buyer." However, 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 done by replacing the orders q_t with the quantities $\epsilon_t q_t$ that the retailer would actually receive where ϵ_t is a random variable defined over the interval $[0, 1]$. This formulation introduces a difficulty we have not encountered so far—namely, the model is no longer a fixed recourse model (see §3.1). In this case, the coefficients associated with the adjustable variables q_2, \dots, q_T are uncertain. As mentioned in §3, for such cases the corresponding

AARC model may be intractable (NP-hard). However, it can be tightly approximated by a tractable convex program (conic-quadratic or semi-definite)—see §4 in Ben-Tal et al. (2004). Additional analysis of such models will require further development of the computational aspects of the RO methodology.

6. Summary

The main purpose of this paper is to present the potential benefits of applying the robust optimization methodology, in particular the new AARC heuristic, to the modeling and analysis of complex operations management problems. To do so, we focused on the topic of flexible supplier-retailer contracts that have attracted significant attention in recent years. We presented a “prototype” formulation (the RSFC model) and developed its robust counterpart. Through the analysis of this model and its extensions we have shown that RO is a powerful technique that is capable of treating large-scale problems that could not have been solved through the earlier solution methodologies (mainly DP and SPR) that were proposed for such problems. We have also demonstrated the flexibility of the methodology and its ability to easily adapt itself to different versions of the RSFC model.

The low computational complexity of the AARC heuristic makes it possible to use it “offline” to assess the value of a contract during the negotiations stages and search for the best-fitted arrangement among parties. However, it can also be used “online,” in a rolling or folding horizon mode, where the current period decisions are taken by solving an updated version of the RO model to account for the actual realization of the demand. This approach can be interpreted as a kind of an open loop control, where the control is not given in an analytical form, but as the optimal solution of a computationally simple optimization problem.

In much the same way, RO can be employed to analyze a large collection of OM problems. In fact, any such problem with uncertain parameters, whose deterministic version can be modeled as a multi-stage linear program, can be treated with our affinely adjustable robust optimization methodology. One such problem is the transshipment problem discussed by Robinson (1990). That problem, which may be defined over finite, rolling, or folding horizons,

involves multiple retailers who must make nonadjustable decisions at the beginning of each period on the replenishment quantities they need to order in that period. Then, after demand is realized, they can adjust the decision variables that determine how much they transship (or receive) to (from) other retailers.

Acknowledgments

This work was partially supported by the Bi-National Science Foundation (BSF) under Grant 2002038 and by the Minerva Foundation.

Appendix

Let (P) be problem (5) with the box demand uncertainty (9)–(10), and let (P_+) be the problem obtained from (5) by replacing the original uncertainty set with the one given by

$$\text{ext}(\mathcal{U}_{\text{box}}) = \text{ext}(\mathcal{U}_1) \times \cdots \times \text{ext}(\mathcal{U}_T),$$

where $\text{ext}(\mathcal{U}_t) = \{d_t^{\min}, d_t^{\max}\}$ is the set of extreme points of the segment \mathcal{U}_t . Our goal is to prove the following statement:

PROPOSITION 1. *The optimal values in (P) and (P_+) are equal to each other.*

PROOF OF PROPOSITION 1. We start with the following simple fact about worst-case oriented dynamic programming:

LEMMA 2. *Let $D_\tau \subset \mathbf{R}^{n_\tau}$ be convex polytopes, $\tau = 0, 1, \dots, T$, where D_0 is a singleton, and let $\mathcal{D} = D_0 \times D_1 \times \cdots \times D_T$. For $d = (d_0, d_1, \dots, d_T) \in \mathcal{D}$ and $0 \leq \tau \leq T$, let $d^\tau = (d_1, \dots, d_\tau)$. Further, given sets F_0, \dots, F_T , $\emptyset \neq F_\tau \subset D_\tau$, consider the optimization problem*

$$(P[F_0, \dots, F_T]):$$

$$\begin{aligned} \min_{S_1(\cdot), \dots, S_{T+1}(\cdot), E} \{ & E \geq f_1(S_1(d^0)) + f_2(S_2(d^1)) \\ & + \cdots + f_{T+1}(S_{T+1}(d^T)); A_1 S_1(d^0) \geq b_1; \\ & A_{t+1} S_{t+1}(d^t) \geq B_{t+1} d_t + C_{t+1} S_t(d^{t-1}) + b_{t+1}, \\ & t = 1, \dots, T; \|S_t(d^{t-1})\|_\infty \leq R, \\ & t = 1, \dots, T + 1; \forall d \in \mathcal{F}_T \}, \end{aligned}$$

where $\mathcal{F}_T = F_0 \times \cdots \times F_T$, $S_t(\cdot)$ are allowed to be arbitrary functions of $d^{t-1} \in F_0 \times \cdots \times F_{t-1}$ taking values in \mathbf{R}^{m_t} , A_t , B_t , and C_t are fixed matrices of appropriate sizes, b_1, \dots, b_{T+1} are fixed vectors, and $\|u\|_\infty$ is the maximum of modulae of coordinates of a vector u .

Assume that for every $t \leq T + 1$, $f_t(\cdot)$ is a convex polyhedral function (c.p.f.) on \mathbf{R}^{m_t} , that is, $f_t(\cdot)$ takes real values and the value $+\infty$, the domain $\text{Dom} f_t \equiv \{s_t: f_t(s_t) < \infty\}$ is a convex set given by finitely many nonstrict linear inequalities, and f_t is convex and lower semicontinuous on $\text{Dom} f_t$. Let $\text{ext}(D_t)$ denote the set of extreme points of D_t . Then, the optimal value in $(P[D_0, \dots, D_T])$ is the same as the optimal value in $(P[\text{ext}(D_0), \dots, \text{ext}(D_T)])$.

PROOF OF LEMMA 2. The proof is obtained by induction in T . Base $T = 0$ is trivial because D_0 is a singleton and therefore $\text{ext}(D_0) = D_0$. Let us justify the inductive step. Consider the optimization problem (depending on $s_T \in \mathbf{R}^{m_T}$ and $d_T \in \mathbf{R}^{n_T}$ as on parameters)

$$\begin{aligned}\Phi_T(s_T, d_T) &\equiv \min_{s_{T+1}} \{f_{T+1}(s_{T+1}): s_{T+1} \in \mathcal{S}_{T+1}(s_T, d_T)\} \\ &= \min_{s_{T+1}} \{f_{T+1}(s_{T+1}): A_{T+1}s_{T+1} \\ &\quad \geq B_{T+1}d_T + C_{T+1}s_T + b_{T+1}, \|s_{T+1}\|_\infty \leq R\},\end{aligned}$$

and let

$$\phi(s_T) = \max_{d_T \in F_T} \Phi_T(s_T, d_T).$$

Applying the Bellman equation (cf. (8)), we see that the optimal value in $(P[F_0, \dots, F_T])$ is exactly the optimal value in the problem

$$\begin{aligned}(P_+[F_0, \dots, F_{T-1}]) \\ \min_{s_1(\cdot), \dots, s_{T-1}(\cdot), E} \{E \geq f_1(S_1(d^0)) + \dots + f_{T-1}(S_{T-1}(d^{T-2})) \\ + \tilde{f}_T(S_T(d^{T-1})); A_1 S_1(d^0) \geq b_1; \\ A_{t+1} S_{t+1}(d^t) \geq B_{t+1} d_t + C_{t+1} S_t(d^{t-1}) + b_{t+1}, \\ t = 1, \dots, T-1; \|S_t(d^{t-1})\|_\infty \leq R, \\ t = 1, \dots, T; \forall d \in \mathcal{F}_{T-1}\},\end{aligned}$$

where

$$\tilde{f}_T(s_T) = f_T(s_T) + \phi(s_T).$$

Let us specify F_T as D_T . The following result is well known. Consider the optimization program:

$$\text{val}(b) = \min_s \{f(s): As \geq b, \|s\|_\infty \leq R\}$$

and assume that the objective function f is a convex polyhedral function. $\text{val}(\cdot)$ is a c.p.f. of b . Moreover, an affine transformation of the variables \cdot in $\text{val}(\cdot)$ retain its c.p.f. property. In our situation, the function Φ_T is of this form, i.e.,

$$\Phi_T(s_T, d_T) = \text{val}(B_{T+1}d_T + C_{T+1}s_T),$$

where $\text{val}(\cdot)$ is a c.p.f. By convexity of $\text{val}(\cdot)$, we have

$$\begin{aligned}\max_{d_T \in D_T} \text{val}(B_{T+1}d_T + C_{T+1}s_T) \\ = \max_{d_T \in \text{ext}(D_T)} \text{val}(B_{T+1}d_T + C_{T+1}s_T); \quad (38)\end{aligned}$$

that is,

$$\phi(s_T) = \max_{d_T \in \text{ext}(D_T)} \text{val}(B_{T+1}d_T + C_{T+1}s_T).$$

This observation implies that

(a) $\phi(\cdot)$ is a cumulative probability function (as a maximum of finitely many functions of this type), so that $\tilde{f}_T(s_T)$ also is a c.p.f., and

(b) we have

$$\begin{aligned}\text{opt}(P[F_0, \dots, F_{T-1}, D_T]) \\ = \text{opt}(P_+[F_0, \dots, F_{T-1}]) \\ = \text{opt}(P[F_0, \dots, F_{T-1}, \text{ext}(D_T)]); \quad (39)\end{aligned}$$

indeed, both equalities are readily given by (38) combined with Bellman equations as applied to problems $(P[F_0, \dots, F_{T-1}, D_T])$ and $(P[F_0, \dots, F_{T-1}, \text{ext}(D_T)])$, respectively.

From (a) it follows that problem $(P_+[D_0, \dots, D_{T-1}])$ satisfies the premise of our lemma, so that by the inductive hypothesis we have

$$\begin{aligned}\text{opt}(P_+[D_0, \dots, D_{T-1}]) \\ = \text{opt}(P_+[\text{ext}(D_0), \dots, \text{ext}(D_{T-1})]). \quad (40)\end{aligned}$$

It follows that

$$\begin{aligned}\text{opt}(P[\text{ext}(D_0), \dots, \text{ext}(D_T)]) \\ = \text{opt}(P_+[\text{ext}(D_0), \dots, \text{ext}(D_{T-1})]) \\ = \text{opt}(P_+[D_0, \dots, D_{T-1}]) = \text{opt}(P[D_0, \dots, D_T])\end{aligned}$$

(the first equality is the right equality in (39) when $F_t = \text{ext}(D_t)$, $t \leq T-1$, the second is (40), and the third is the left equality in (39) when $F_t = D_t$, $t \leq T-1$); the inductive step is justified. \square

PROOF OF PROPOSITION 1 (CONTINUED). Now we are in a position to prove the proposition. To this end, note that (P) , (P_+) can be rewritten as $(P[D_0, \dots, D_T])$, respectively, $(P[\text{ext}(D_0), \dots, \text{ext}(D_T)])$ with the data defined as follows:

- D_t s and s_t s are given by

$$D_t = [d_t^* - \delta d_t, d_t^* + \delta d_t], \quad t = 1, \dots, T,$$

$$s_1 = \begin{bmatrix} q_1 \\ \hat{q}_1 \\ u_1 \\ w^1 \\ z^1 \end{bmatrix}, \quad s_t = \begin{bmatrix} x_t \\ y_t \\ q_t \\ \hat{q}_t \\ u_t \\ w^t \\ z^t \end{bmatrix}, \quad t = 2, \dots, T, \quad s_{T+1} = \begin{bmatrix} x_{T+1} \\ y_{T+1} \end{bmatrix}.$$

- Systems of inequalities $A_{t+1}s_{t+1} \geq B_{t+1}d_t + C_{t+1}s_t + b_{t+1}$ are equivalent to the following relations:

$t = 0$: $x_1 = x_{\text{ini}}$, $L_1 \leq q_1 \leq U_1$, $\hat{q}_1 = q_1$, $\hat{L}_1 \leq \hat{q}_1 \leq \hat{U}_1$, $w^1 \geq 0$, $u_1 - \alpha_1^+(q_1 - w_1^1) \geq 0$, $u_1 - \alpha_1^-(w_1^1 - q_1) \geq 0$, $z_\tau^1 \geq \beta_\tau^+(w_\tau^1 - w_{\tau-1}^1)$, $z_\tau^1 \geq \beta_\tau^-(w_{\tau-1}^1 - w_\tau^1)$, $\tau = 1, \dots, T$ ($w_0^1 \equiv w_0$ and x_1 are given by the data).

$0 < t < T$: $x_{t+1} = x_t + q_t - d_t$, $y_{t+1} - h_t x_{t+1} \geq 0$, $y_{t+1} + p_t x_{t+1} \geq 0$, $L_{t+1} \leq q_{t+1} \leq U_{t+1}$, $\hat{q}_{t+1} = q_{t+1} + \hat{q}_t$, $\hat{L}_{t+1} \leq \hat{q}_{t+1} \leq \hat{U}_{t+1}$, $u_{t+1} - \alpha_{t+1}^+(q_{t+1} - w_{t+1}^1) \geq 0$, $u_{t+1} - \alpha_{t+1}^-(w_{t+1}^1 - q_{t+1}) \geq 0$, $w^{t+1} = w^t$, $z^{t+1} = z^t$.

$t = T$: $x_{T+1} = x_T + q_T - d_T$, $y_{T+1} - (h_T - s)x_{T+1} \geq 0$, $y_{T+1} + p_T x_{T+1} \geq 0$.

- The functions $f_t(s_t)$ are given by

$$t = 1: f_1(s_1) = c_1 q_1 + \sum_{\tau=1}^T z_\tau^1 + u_1; \quad 1 < t \leq T: f_t(s_t) = y_t + u_t + c_t q_t;$$

$$t = T + 1: f_{T+1}(s_{T+1}) = y_{T+1}.$$

With this setup, the only difference between problems (P) and $(P[D_0, \dots, D_T])$ is that the latter problem has additional box constraints. Clearly, for an R large enough this restriction does not affect the optimal value, and it is similar for the pair of problems (P_+) , $(P[\text{ext}(D_0), \dots, \text{ext}(D_T)])$. Assuming R large enough and applying Lemma 2, we arrive at

$$\begin{aligned} \text{opt}(P) &= \text{opt}(P[D_0, \dots, D_T]) \\ &= \text{opt}(P[\text{ext}(D_0), \dots, \text{ext}(D_T)]) = \text{opt}(P_+). \quad \square \end{aligned}$$

References

- Anupindi, R., Y. Bassok. 1998. Approximation for multiproduct contracts with stochastic demands and business volume discounts: Single supplier case. *IIE Trans.* **30** 723–734.
- Anupindi, R., Y. Bassok. 1999. Supply contracts with quantity commitments and stochastic demand. S. Tayur, M. Magazine, R. Ganeshan, eds. *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, Norwell, MA.
- Bassok, Y., R. Anupindi. 1997. Analysis of supply contracts with total minimum commitment. *IIE Trans.* **29** 373–381.
- Bassok, Y., R. Anupindi. 1998. Analysis of supply contracts with commitments and flexibility. Working paper, J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, IL.
- Bassok, Y., A. Bixby, R. Srinivasan, H. Z. Wiesel. 1997. Design of component-supply contracts with commitment-revision flexibility. *IBM J. Res. Development* **41**(6) 1–14.
- Ben-Tal, A., A. Nemirovski. 1997. Robust truss topology design via semidefinite programming. *SIAM J. Optim.* **7** 991–1016.
- Ben-Tal, A., A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming (Series B)* **88** 411–424.
- Ben-Tal, A., A. Nemirovski. 2001. *Lectures on Modern Convex Optimization: Analysis, Algorithms and Engineering Applications*. SIAM-MPS Series in Optimization. SIAM, Philadelphia, PA.
- Ben-Tal, A., A. Nemirovski. 2002. A robust optimization—Methodology and applications. *Math. Programming (Series B)* **92** 453–480.
- Ben-Tal, A., A. Goryashko, E. Guslitzer, A. Nemirovski. 2004. Adjusting robust solutions of uncertain linear programs. *Math. Programming* **99**(2) 351–376.
- Ben-Tal, A., M. Kocvara, A. Nemirovski, J. Zowe. 1999. Free material design via semidefinite programming, the multi load case with contact conditions. *SIAM J. Optim.* **9** 813–832.
- Bertsekas, D., J. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA.
- Bertsimas, D., A. Thiele. 2004. A robust optimization approach to supply chain management. *Lecture Notes Comput. Sci.* **3064** 86–100.
- Chen, F. Y., D. Krass. 2001. Analysis of supply contracts with minimum total commitment and non-stationary demands. *Eur. J. Oper. Res.* **131** 309–323.
- Feng, Y. Y., C. Gallego. 2000. Perishable asset revenue management with Markovian time dependent demand intensities. *Management Sci.* **46** 941–956.
- Fisher, M., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* **44** 87–99.
- Fisher, M., K. Ramdas, Y. S. Zheng. 2001. Ending inventory valuation in a supply chain: The bullwhip effect. *Management Sci.* **45** 679–692.
- Holt, C. C., F. Modigliani, J. F. Muth, H. Simon. 1960. *Planning Production, Inventories, and Work Force*. Prentice-Hall, Englewood Cliffs, NJ.
- Iyengar, G. 2005. Robust dynamic programming. *Math. Oper. Res.* **30** 257–280.
- Lee, H., P. Padmanabhan, S. Whang. 1996. Information distortion in a supply chain: The bullwhip effect. *Management Sci.* **43** 546–558.
- Monahan, J. P. 1984. Quantity discount pricing models to increase vendor profits. *Management Sci.* **30** 720–726.
- Robinson, L. W. 1990. Optimal and approximate policies in multiperiod, multilocation inventory models with transshipment. *Oper. Res.* **38** 278–295.
- Scarf, H. 1960. The optimality of (S, s) policies in the dynamic inventory problems, Ch. 13. K. J. Arrow, S. Karlin, P. Suppes, eds. *Mathematical Methods in the Social Sciences*. Stanford University Press, Stanford, CA.
- Schweitzer, P. J., A. Seidmann. 1985. Generalized polynomial approximations in Markovian decision-processes. *J. Math. Anal. Appl.* **110** 568–582.
- Talluri, K., G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Sci.* **44** 1577–1593.
- Tsay, A. A. 1995. *Supply Chain Control with Quantity Flexibility*. Ph.D. dissertation, Graduate School of Business, Stanford University, Stanford, CA.
- Tsay, A. A., W. S. Lovejoy. 1999. Quantity flexible contracts and supply chain performance. *Manufacturing Service Oper. Management* **1** 89–111.
- Urban, T. L. 2000. Supply contracts with periodic stationary commitments. *Production Oper. Management* **9** 400–413.
- van Delft, C., J. P. Vial. 2004. A practical implementation of stochastic programming: An applications to the evaluation of option contracts in supply chains. *Automatica* **40** 743–756.
- Veinott, A. F. 1966. State of mathematical inventory theory. *Management Sci.* **12** 747–777.
- Wang, Y. Z., Y. Gerchak. 1996. Periodic review production models with variable capacity, random yield, and uncertain demand. *Management Sci.* **42**(1) 130–137.