Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Robust Linear Optimization: On the benefits of distributional information and applications in inventory control *

Ioannis Ch. Paschalidis†, *Member, IEEE*, Seong-Cheol Kang‡,

***Abstract*— Linear programming formulations cannot handle the presence of uncertainty in the problem data and even small variations in the data can render an optimal solution infeasible. A number of *robust linear optimization* techniques produce formulations (not necessarily linear) that guarantee the feasibility of the optimal solutions for all realizations of the uncertain data. A recent robust approach in [1] maintains the linearity of the formulation and is able to strike a balance between the conservatism and quality of a solution by allowing less robust solutions. In this work we demonstrate how to use distributional information on problem data in robust linear optimization. We adopt the robust model of [1] and present an approach that exploits distributional information on problem data to decide the level of robustness of the formulation, thus, leading to much more cost-effective solutions (by 50% or more in some instances). We apply our methodology to a stochastic inventory control problem with quality of service constraints.**

***Index Terms*— Robust optimization, Linear programming, Data uncertainty, Inventory Control, Quality-of-Service.**

I. Introduction

A *linear programming (LP)* problem is, perhaps, among the most fortunate outcomes when we formulate an optimization problem. Theory and methods are very mature, there are many excellent solvers to choose from, and the problem can be solved in polynomial time with interior-point methods. Alas, the world is neither (always) linear nor certain. In this paper we focus on the latter shortcoming of LP-based modeling, that is, the presence of uncertainty in the problem data.

A *certainty equivalence* approach offers a way to deal with uncertainty. For every uncertain data element, use a *nominal* value – usually its mean – and form a *nominal problem* which remains an LP. A solution obtained in this manner, however, is *non-robust* as small changes in the problem data can render the solution infeasible. In many realistic settings this implies that the solution becomes useless.

To prevent possible infeasibility of a solution and therefore to ensure its usability, one may construct a *robust problem* whose solution is guaranteed to be feasible. Soyster [2] appears to be the first who addressed this issue for LPs. He considers an LP problem with "column-wise" uncertainty:

$$\begin{aligned} \max\ & \mathbf{c}'\mathbf{x} \\ \text{s.t.}\ & \sum_{j=1}^{n} \mathbf{A}_j x_j \le \mathbf{b}, \quad \forall \mathbf{A}_j \in K_j,\ j=1,\dots,n \\ & \mathbf{x} \ge \mathbf{0}, \end{aligned}$$

where each column $\mathbf{A}_j$ of the constraint matrix $\mathbf{A}$ belongs to a given convex set K_j. Soyster [2] shows that the problem can be recast as a finite dimensional LP problem.

Ben-Tal and Nemirovski [3] point out that the case of "column-wise" uncertainty considered in [2] is extremely conservative. They instead consider "row-wise" uncertainty where the rows of the constraint matrix are known to belong to given convex sets. In this case, they show that the robust problem is typically not an LP problem; for example, when the uncertain sets for the rows of $\mathbf{A}$ are ellipsoids, the robust problem turns out to be a conic quadratic problem.

The robust models of Soyster [2] and Ben-Tal and Nemirovski [3] adopt a "worst-case" approach. Although the guaranteed feasibility is an attractive feature of those robust formulations, it comes with a price: a degradation of the objective value. In several applications, however, this price may be unacceptable, especially if the "worst-case" happens very rarely. Hence for these applications, it could be more desirable to obtain a less robust solution with a better objective value, which also admits a very low probability of being infeasible. This is the rationale for the robust optimization approach of Bertsimas and Sim [1].

Bertsimas and Sim [1] consider "element-wise" uncertainty: each uncertain element is modeled as an independent, symmetric, and bounded random variable whose range is known but the distribution is unknown. Their approach is flexible enough to encompass the nominal problem as well as the Soyster model as special cases. Like the Soyster model, the robust formulation in [1] remains an LP. The robustness of the formulation in [1] is controlled by a set of parameters which regulates the "degree of uncertainty" in the problem data. Bertsimas and Sim [1] provide bounds on the probability that an optimal solution of their robust formulation becomes infeasible due to data uncertainty and these bounds hold for all probability distributions of the problem data as long as they satisfy a symmetry assumption.

Another research direction in the literature to deal with uncertainty in the problem data is to use *chance constraints*:

$$P[\mathbf{a}_i'\mathbf{x} > b_i] \le \epsilon_i, \quad \forall i,$$

where $\mathbf{a}_i'$ is the ith row of $\mathbf{A}$. By adjusting the values of ϵ_i's, this approach also allows less robust solutions with better objective values. Two interesting approximation methods for more general forms of chance constraints are given in Nemirovski and Shapiro [4] and in Calafiore and Campi [5].

The starting point for our work is the robust formulation of [1]. We will quantify the benefit of having access to probability distributions of problem data. We will show

* Research partially supported by the NSF under a CAREER award ANI-9983221 and grants DMI-0330171, ECS-0426453, CNS-0435312, and DMI-0300359, and by the ARO under the ODDR&E MURI2001 Grant DAAD19-01-1-0465 to the Center for Networked Communicating Control Systems.

† Center for Information & Systems Eng., and Dept. of Manufacturing Eng., Boston University, 15 St. Mary's St., Brookline, MA 02446, e-mail: `yannisp@bu.edu`, url: `http://ionia.bu.edu/`.

† Center for Information & Systems Eng., and Dept. of Manufacturing Eng., Boston University, e-mail: `jsckang@bu.edu`.

that if probability distributions are known one can obtain solutions that are much less conservative than the ones obtained in [1] (by 50% or more in several examples we present). The crux of the matter is that by exploiting distributional information we can obtain much tighter bounds on the probability that an optimal solution of the robust formulation becomes infeasible; this leads us to "injecting" less robustness into the formulation and to much more cost-effective solutions.

Our motivation comes from the (emerging) abundance of data for many real-world applications. Mining these data sets, one can obtain distributional information and, as we show, put it into good use. When data is not available, our work can help quantify the benefits that can result from data collection and from implementing estimation techniques for obtaining distributional information. We expect that in many settings, these benefits can exceed the associated costs.

The rest of the paper is organized as follows. In Section II, we outline the robust formulation for general LP problems proposed in [1] and present our new bounds on the constraint violation probability. In Section III, we consider as an application a discrete-time stochastic inventory control problem with *Quality of Service (QoS)* constraints. A related work can be found in Bertsimas and Thiele [6]. Our approach differs from [6] in two aspects: First we include the QoS constraints in an attempt to replace shortage costs that are hard to quantify. Second our robust formulation seems to be more natural and tighter than that of [6]. We will elaborate on this in Subsection III-A. Finally, concluding remarks are given in Section IV.

II. Robust Linear Optimization

In this section we consider general LPs and start by reviewing the robust formulation of [1]. Our new bounds on the constraint violation probability are presented in Subsection II-B. Subsection II-C reports some numerical results.

A. Data Uncertainty and Robust Problem

Consider the following LP problem

$$\begin{aligned} \max\ & \mathbf{c}'\mathbf{x} \\ \text{s.t.}\ & \mathbf{A}\mathbf{x} \le \mathbf{b} \\ & \mathbf{l} \le \mathbf{x} \le \mathbf{u}, \end{aligned} \tag{1}$$

where $\mathbf{c}, \mathbf{l}, \mathbf{u} \in \mathbb{R}^n$, $\mathbf{b} \in \mathbb{R}^m$, $\mathbf{A}$ is an $m \times n$ matrix, and $\mathbf{x} \in \mathbb{R}^n$ is the decision vector. We assume, without loss of generality, that only the elements of $\mathbf{A}$ are subject to uncertainty.[1] The elements of $\mathbf{A}$ are of two types: random and deterministic. In particular, consider the ith row of $\mathbf{A}$ and let J_i be the set of indices j such that the corresponding a_{ij} are subject to uncertainty. We assume that for all i each element a_{ij}, $j \in J_i$, is modeled as an *independent* and *bounded* random variable taking values in $[\overline{a}_{ij} - \hat{a}_{ij}, \overline{a}_{ij} + \hat{a}_{ij}]$, where $\overline{a}_{ij} = E[a_{ij}]$ and $\hat{a}_{ij} > 0$. Elements a_{ij}, $j \notin J_i$, on the other hand, are deterministic and have fixed values $\overline{a}_{ij}$. For the probability distribution of a_{ij}, for all i and $j \in J_i$, the following symmetry assumption will be in effect for most of the results we will present.

Assumption A *The probability distribution of a_{ij} is symmetric over $[\overline{a}_{ij} - \hat{a}_{ij}, \overline{a}_{ij} + \hat{a}_{ij}]$ for all i and $j \in J_i$.*

In [1], it is assumed that the probability distribution of a_{ij}, for all i and $j \in J_i$, is unknown other than its symmetry. In this paper we assume that these probability distributions are known.

Given this data uncertainty structure for $\mathbf{A}$, one may elect to solve the following *nominal problem* where each random element is replaced by its mean value

$$\begin{aligned} z_N = \max\ & \mathbf{c}'\mathbf{x} \\ \text{s.t.}\ & \sum_j \overline{a}_{ij} x_j \le b_i, \quad \forall i \\ & \mathbf{l} \le \mathbf{x} \le \mathbf{u}. \end{aligned} \tag{2}$$

One disadvantage of the nominal problem is that its optimal solution is highly likely to violate the original constraints $\mathbf{A}\mathbf{x} \le \mathbf{b}$. This leads us to consider a solution that is guaranteed to satisfy $\mathbf{A}\mathbf{x} \le \mathbf{b}$ for all realizations of a_{ij}, for all i and $j \in J_i$, and at the same time attains as maximal an objective value as possible. To obtain such a solution, we need to solve the following problem which, henceforth, will be referred to as the Soyster model or *fat problem*[2]

$$\begin{aligned} z_F = \max\ & \mathbf{c}'\mathbf{x} \\ \text{s.t.}\ & \sum_j \overline{a}_{ij} x_j + \sum_{j \in J_i} \hat{a}_{ij} y_j \le b_i, \quad \forall i \\ & -\mathbf{y} \le \mathbf{x} \le \mathbf{y} \\ & \mathbf{l} \le \mathbf{x} \le \mathbf{u} \\ & \mathbf{y} \ge \mathbf{0}. \end{aligned} \tag{3}$$

The following lemma is almost immediate.

Lemma II.1 *Let $(\mathbf{x}^*, \mathbf{y}^*)$ be an optimal solution of (3). Then $\mathbf{x}^*$ is a feasible solution of (1) for every possible realization of $\mathbf{A}$. Moreover $z_F \le z_N$.*

What Lemma II.1 states is that by solving the fat problem one may obtain a worse solution in exchange for *robustness* to data uncertainty.

Noting this trade-off between the conservatism of a solution and its objective value, Bertsimas and Sim [1] introduce a parameter Γ_i for each row i. Γ_i takes values in $[0, |J_i|]$ and is not necessarily an integer. Γ_i is interpreted as the number of a_{ij}, $j \in J_i$ that are allowed to be random. For instance if $\Gamma_i = 0$, then all a_{ij}, $j \in J_i$, are forced to be deterministic and take values $\overline{a}_{ij}$; i.e., there is no randomness in row i. If $\Gamma_i = |J_i|$, then all a_{ij}, $j \in J_i$, are random and take values from their respective ranges $[\overline{a}_{ij} - \hat{a}_{ij}, \overline{a}_{ij} + \hat{a}_{ij}]$. In general, only $\lfloor \Gamma_i \rfloor$ elements among a_{ij}, $j \in J_i$, are allowed to take values from their respective ranges, and one other element, say a_{it_i}, takes values from its truncated range $[\overline{a}_{it_i} - (\Gamma_i - \lfloor \Gamma_i \rfloor)\hat{a}_{it_i}, \overline{a}_{it_i} + (\Gamma_i - \lfloor \Gamma_i \rfloor)\hat{a}_{it_i}]$, while the remaining $|J_i| - \lceil \Gamma_i \rceil$ random elements are forced to be deterministic, taking their respective mean values.

Given Γ_i for all i, we seek to obtain a *robust solution* with maximal objective value, which is guaranteed to be feasible

[1] If $\mathbf{c}$ and $\mathbf{b}$ are also subject to uncertainty, we can reformulate the problem so that all uncertain elements are contained in $\mathbf{A}$ of the new formulation.

[2] In Section I this problem was called the robust problem. We reserve the term, *robust problem*, for a new robust formulation to be introduced later. This should not cause any confusion.

to (1) as long as up to Γ_i elements of the ith row of $\mathbf{A}$ are random. Let $\mathbf{\Gamma} = (\Gamma_1, \ldots, \Gamma_m)$. Let S_i be a subset of J_i such that $|S_i| = \lfloor \Gamma_i \rfloor$, and let $t_i \in J_i \setminus S_i$. Such a robust solution is obtained from the following *robust problem*

$$
\begin{aligned}
z_R(\mathbf{\Gamma}) = \max \;\; & \mathbf{c}'\mathbf{x} \qquad (4)\\
\text{s.t.} \;\; & \sum_j \overline{a}_{ij} x_j + \\
& \max_{S_i \cup \{t_i\}} \left\{ \sum_{j \in S_i} \hat{a}_{ij} y_j + (\Gamma_i - \lfloor \Gamma_i \rfloor) \hat{a}_{it_i} y_{t_i} \right\} \le b_i, \quad \forall i \\
& -\mathbf{y} \le \mathbf{x} \le \mathbf{y} \\
& \mathbf{l} \le \mathbf{x} \le \mathbf{u} \\
& \mathbf{y} \ge \mathbf{0}.
\end{aligned}
$$

We have the following lemma; the proof is omitted.

Lemma II.2 *$z_R(\mathbf{\Gamma})$ is a non-increasing function of $\mathbf{\Gamma}$, and $z_F \le z_R(\mathbf{\Gamma}) \le z_N$.*

By varying Γ_i between 0 and $|J_i|$ for all i, one is able to strike a balance between the conservatism of a solution and its objective value. Bertsimas and Sim [1] show that the nonlinear formulation (4) can be reformulated as an equivalent LP problem.

Theorem II.3 ([1]) *The following LP problem is equivalent to the robust problem (4):*

$$
\begin{aligned}
\max \;\; & \mathbf{c}'\mathbf{x} \qquad (5)\\
\text{s.t.} \;\; & \sum_j \overline{a}_{ij} x_j + z_i \Gamma_i + \sum_{j \in J_i} p_{ij} \le b_i, \quad \forall i \\
& z_i + p_{ij} \ge \hat{a}_{ij} y_j, \quad \forall j \in J_i, \quad \forall i \\
& p_{ij} \ge 0, \quad \forall j \in J_i, \quad \forall i \\
& -\mathbf{y} \le \mathbf{x} \le \mathbf{y} \\
& \mathbf{l} \le \mathbf{x} \le \mathbf{u} \\
& \mathbf{y}, \; \mathbf{z} \ge \mathbf{0}.
\end{aligned}
$$

B. Bounds on the Constraint Violation Probability

Assume $\Gamma_i > 0$ for all i. Let $\mathbf{x}^*$ be an optimal solution of (4), which can be obtained by solving (5). To avoid degenerate cases and without loss of generality, we will assume $|\mathbf{x}^*| > \mathbf{0}$.[3] Unless $\Gamma_i = |J_i|$ for all i, $\mathbf{x}^*$ may violate constraints $\mathbf{Ax} \le \mathbf{b}$. Bertsimas and Sim [1] derive the following distribution-free upper bound on the probability that the ith constraint is violated at $\mathbf{x}^*$.

Theorem II.4 ([1]) *Under Assumption A,*

$$
P\left[\sum_j a_{ij} x_j^* > b_i\right] \le \exp\left[-\frac{\Gamma_i^2}{2|J_i|}\right]. \qquad (6)
$$

Henceforth, we will use the terminology bound (6) to refer to the right hand side of (6) and a similar terminology for such bounds. Bound (6) is an *a priori* bound (i.e., solution-independent bound) in the sense that its computation requires only problem parameters and not an optimal solution to the robust problem. However, as we will see, bound (6) is typically weak because it utilizes neither the probability distributions of a_{ij}'s nor the optimal solution $\mathbf{x}^*$.

When the probability distributions of a_{ij}'s are available, we can derive a tighter *a priori* bound as the following theorem shows (due to space limitations we omit the proof). Let $\eta_{ij} = (a_{ij} - \overline{a}_{ij})/\hat{a}_{ij}$, $\forall j \in J_i$. Define the logarithmic moment generating function of η_{ij} as $\Lambda_{\eta_{ij}}(\theta) \triangleq \log E[e^{\theta \eta_{ij}}]$.

Theorem II.5 *Let Assumption A be in effect.*
(a) The ith constraint violation probability at $\mathbf{x}^*$ satisfies

$$
P\left[\sum_j a_{ij} x_j^* > b_i\right] \le \exp\left[-\sup_{\theta \ge 0}\left[\theta \Gamma_i - \sum_{j \in J_i} \Lambda_{\eta_{ij}}(\theta)\right]\right]. \qquad (7)
$$

(b) bound (7) $\le$ bound (6).

Observe that both bounds (6) and (7) are decreasing in Γ_i. This monotonicity and Theorem II.5(b) have the following important implication: to ensure the same constraint violation probability, bound (7) requires smaller Γ_i than bound (6) does. Since $z_R(\mathbf{\Gamma})$ is a non-increasing function of $\mathbf{\Gamma}$, we can achieve a higher $z_R(\mathbf{\Gamma})$ by using Γ_i required by bound (7), while maintaining the same constraint violation probability. In other words, bound (7) enables us to obtain an equally robust solution with a better objective value.

Once an optimal solution $\mathbf{x}^*$ of (4) is available, we can also compute the following *a posteriori* bound (i.e., solution-dependent bound). As the following theorem shows, this bound is tighter than bound (7). A host of numerical examples we present later establish that the difference can be dramatic. Due to space limitations, we omit the proof of the theorem.

Theorem II.6 *Let Assumption A be in effect. Let $C_i(\mathbf{x}^*) = b_i - \sum_j \overline{a}_{ij} x_j^*$ and $\beta_{ij} = \hat{a}_{ij}|x_j^*|$ for $j \in J_i$.*
(a) The ith constraint violation probability at $\mathbf{x}^*$ satisfies

$$
P\left[\sum_j a_{ij} x_j^* > b_i\right] \le \exp\left[-\sup_{\theta \ge 0}\left[\theta C_i(\mathbf{x}^*) - \sum_{j \in J_i} \Lambda_{\eta_{ij}}(\theta \beta_{ij})\right]\right]. \qquad (8)
$$

(b) bound (8) $\le$ bound (7).

If Assumption A is not in effect, a slightly different bound (Corollary II.7) can be obtained. This can be quite useful as in many instances the symmetry assumption can be quite restrictive.

Corollary II.7 *Let $C_i(\mathbf{x}^*) = b_i - \sum_j \overline{a}_{ij} x_j^*$ and $\kappa_{ij} = \hat{a}_{ij} x_j^*$ for $j \in J_i$. The ith constraint violation probability at $\mathbf{x}^*$ satisfies*

$$
P\left[\sum_j a_{ij} x_j^* > b_i\right] \le \exp\left[-\sup_{\theta \ge 0}\left[\theta C_i(\mathbf{x}^*) - \sum_{j \in J_i} \Lambda_{\eta_{ij}}(\theta \kappa_{ij})\right]\right]. \qquad (9)
$$

[3] If not, we can always fix to zero all zero coordinates of $\mathbf{x}^*$ and recast the LP in a lower-dimensional space.

TABLE I

ROBUST PROBLEM OBJECTIVE VALUES AND SOLUTION-DEPENDENT BOUND

No.	$z_R(\mathbf{\Gamma}^1)$	$z_R(\mathbf{\Gamma}^2)$	$\frac{z_R(\mathbf{\Gamma}^2)}{z_R(\mathbf{\Gamma}^1)}$	Bound (8)
1	25.71	34.58	1.345	0.001609
2	31.36	45.58	1.454	0.000600
3	387.36	464.09	1.198	0.000081
4	48.71	59.85	1.229	0.001783
5	96.72	148.01	1.530	0.000223
Average	89.64	110.06	1.228	-

C. Numerical Tests

To assess how much the objective value improves when smaller Γ_i are used and to see how tight bound (8) is, we generate problem instances as follows: $\mathbf{A}$ is a 10×10 matrix; all elements of the matrix $\mathbf{A}$ are random, i.e., $|J_i| = 10$ for $i = 1, \ldots, 10$; $\overline{a}_{ij}$ is randomly drawn from the range $[-100, 100]$, $\forall i, j$; $\hat{a}_{ij}$ is randomly drawn from $[1, 50]$, $\forall i, j$; c_j is randomly selected from $[-50, 50]$, $\forall j$; b_i is randomly selected from $[0, 100]$, $\forall i$; the lower bound l_j is randomly chosen from $[-20, 0]$, $\forall j$; the upper bound u_j is randomly chosen from $[0, 20]$, $\forall j$. We further assume that all a_{ij} are uniformly distributed and that $P[\sum_j a_{ij}x_j^* > b_i] \leq 0.05$ is required for all i.

Bounds (6) and (7) require $\Gamma_i = 7.76$ and $\Gamma_i = 4.34$ respectively for all i to guarantee at most 5% constraint violation probability. (We can use a binary search to determine such Γ_i's because bounds (6) and (7) are monotone in Γ_i.) Let $\mathbf{\Gamma}^1 = 7.76\mathbf{e}$ and $\mathbf{\Gamma}^2 = 4.34\mathbf{e}$, where $\mathbf{e}$ is the vector of all ones. We generated 20 problem instances and in Table I report $z_R(\mathbf{\Gamma}^1)$ and $z_R(\mathbf{\Gamma}^2)$ for some of these instances as well as the averages over the 20 instances. Using the optimal solution $\mathbf{x}^*$ associated with $\mathbf{\Gamma}^2$, we compute bound (8) for all i. Among these bound values, the maximum one is reported in Table I. This experiment demonstrates that by using the distribution-dependent bound (7), we can improve the objective value by more than 50% in some instances and 22% on average. It also shows that bound (8) yields significantly smaller values than 5%.

III. An Inventory Control Problem with QoS Constraints

In this section we apply the robust optimization approach of Section II to a single-station single-item stochastic inventory control problem over N discrete time periods. Let x_k denote the inventory at the beginning of the kth period; w_k the demand during the kth period; u_k the stock ordered at the beginning of the kth period (decision variables); c the ordering cost per unit stock; and h the holding cost per unit stock per period. We assume that each demand w_k is an *independent* and *bounded* random variable with range $[\overline{w}_k - \hat{w}_k, \overline{w}_k + \hat{w}_k]$, where $\overline{w}_k$ and $\hat{w}_k$ are positive constants with $\overline{w}_k > \hat{w}_k$ and $\overline{w}_k = E[w_k]$. The following symmetry assumption on the probability distribution of w_k will be required for some results later on.

Assumption B *The probability distribution of w_k is symmetric over $[\overline{w}_k - \hat{w}_k, \overline{w}_k + \hat{w}_k]$ for all k.*

We also assume that the stock ordered at the beginning of the kth period is delivered instantly, i.e., zero lead time. If we further assume that excess demand is backlogged and filled as soon as additional stock becomes available, inventory evolves according to $x_{k+1} = x_k + u_k - w_k,\ \forall k$.

We consider the *Quality of Service (QoS)* constraints

$$x_k + u_k \geq w_k, \quad k = 1, \ldots, N. \tag{10}$$

We then formulate the following inventory control problem that minimizes total ordering and holding costs, while enforcing the QoS constraints:

$$\begin{aligned} \min\ & \sum_{k=1}^{N}(cu_k + hx_{k+1}) \\ \text{s.t.}\ & x_{k+1} = x_k + u_k - w_k, \quad k = 1, \ldots, N \\ & x_k + u_k \geq w_k, \quad k = 1, \ldots, N \\ & u_k \geq 0, \quad k = 1, \ldots, N, \end{aligned} \tag{11}$$

where x_1 is given. Using the inventory evolution equations, we get $x_{k+1} = x_1 + \sum_{j=1}^{k}(u_j - w_j)$ and we can eliminate x_k, $k = 2, \ldots, N$ from (11). By defining $a_k \triangleq N - k + 1$, $k = 1, \ldots, N$ and introducing an auxiliary variable z for the objective function, (11) can be written as

$$\begin{aligned} \min\ & z \\ \text{s.t.}\ & \sum_{j=1}^{k} u_j \geq -x_1 + \sum_{j=1}^{k} w_j, \quad k = 1, \ldots, N \\ & z - \sum_{k=1}^{N}(c + ha_k)u_k \geq Nhx_1 - h\sum_{k=1}^{N} a_k w_k \\ & u_k \geq 0, \quad k = 1, \ldots, N. \end{aligned} \tag{12}$$

Notice that in (12) the elements w_k in the right hand side of the constraints are random variables, i.e., the vector $\mathbf{b}$ of the generic LP problem (1) corresponding to (12) is subject to uncertainty. The coefficients of the decision variables $\mathbf{u} = (u_1, \ldots, u_N)$ and z are not subject to uncertainty, i.e., the matrix $\mathbf{A}$ is deterministic. As mentioned in Subsection II-A, we may transform (12) in such a way that only the elements of matrix $\mathbf{A}$ are random. In this particular problem, however, it will be more convenient not to do so.

A. Robust Inventory Control Problem

When we introduced the robust formulation for general LPs in Section II, Γ_i was defined for each row i of matrix $\mathbf{A}$, i.e., for each constraint of an LP. It was appropriate because each constraint had a distinct set of random elements. This is not the case for the inventory control problem under consideration. The problem has N random elements, i.e., N random demands w_k, and all of them appear in multiple constraints. For instance w_1 is involved in all of the constraints. Therefore it only makes sense to define a single (global) Γ for the problem. (We point out that Bertsimas and Thiele [6] define different Γ_k for each k.)

The parameter Γ takes values in $[0, N]$ and need not be an integer. We give the same interpretation to Γ as before: $\lfloor\Gamma\rfloor$ out of N demands take values from their respective ranges, and one other demand, say w_t, takes values from its truncated range $[\overline{w}_t - (\Gamma - \lfloor\Gamma\rfloor)\hat{w}_t,\ \overline{w}_t + (\Gamma - \lfloor\Gamma\rfloor)\hat{w}_t]$. The remaining $N - \lceil\Gamma\rceil$ demands are forced to be deterministic and are set equal to their respective mean values. Given Γ, there are many different ways to choose the demands that are subject to uncertainty. Let $\Omega(\Gamma)$ denote the set of all demand vectors $\mathbf{w} = (w_1, \ldots, w_N)$ satisfying the (randomness) model we put forth, that is, all demand vectors with up to Γ random elements. Our goal is to construct an optimization problem such that its optimal solution is guaranteed to be feasible to (12) for all $\mathbf{w} \in \Omega(\Gamma)$. Such an optimization problem is formulated as

$$\begin{aligned} \min\ & z \qquad\qquad (13)\\ \text{s.t.}\ & \sum_{j=1}^{k} u_j \geq \max_{\mathbf{w}\in\Omega(\Gamma)} \left\{ -x_1 + \sum_{j=1}^{k} w_j \right\}, \quad k = 1, \ldots, N \\ & z - \sum_{k=1}^{N} (c + ha_k) u_k \geq \max_{\mathbf{w}\in\Omega(\Gamma)} \left\{ Nhx_1 - h\sum_{k=1}^{N} a_k w_k \right\} \\ & u_k \geq 0, \quad k = 1, \ldots, N. \end{aligned}$$

For $k = \lfloor\Gamma\rfloor + 1, \ldots, N$, let S_k be a subset of the index set $\{1, \ldots, k\}$ such that $|S_k| = \lfloor\Gamma\rfloor$. Furthermore let t_k be an index such that $t_k \in \{1, \ldots, k\} \setminus S_k$. It can be seen that (13) is equivalent to

$$\begin{aligned} z_R(\Gamma) = \min\ & z \qquad\qquad (14)\\ \text{s.t.}\ & \sum_{j=1}^{k} u_j \geq -x_1 + \sum_{j=1}^{k} (\overline{w}_j + \hat{w}_j), \quad k = 1, \ldots, \lfloor\Gamma\rfloor \\ & \sum_{j=1}^{k} u_j \geq -x_1 + \sum_{j=1}^{k} \overline{w}_j + \max_{S_k \cup \{t_k\}} \Bigg\{ \sum_{j\in S_k} \hat{w}_j \\ & \qquad + (\Gamma - \lfloor\Gamma\rfloor)\hat{w}_{t_k} \Bigg\}, \quad k = \lfloor\Gamma\rfloor + 1, \ldots, N \\ & z - \sum_{k=1}^{N} (c + ha_k) u_k \geq Nhx_1 - h\sum_{k=1}^{N} a_k \overline{w}_k \\ & \qquad + h \max_{S_N \cup \{t_N\}} \Bigg\{ \sum_{j\in S_N} a_j \hat{w}_j + (\Gamma - \lfloor\Gamma\rfloor) a_{t_N} \hat{w}_{t_N} \Bigg\} \\ & u_k \geq 0, \quad k = 1, \ldots, N. \end{aligned}$$

We will be referring to (14) as the *robust inventory control problem*.

Let $M(k) \triangleq \max_{S_k\cup\{t_k\}}\{\sum_{j\in S_k} \hat{w}_j + (\Gamma - \lfloor\Gamma\rfloor)\hat{w}_{t_k}\}$, $k = \lfloor\Gamma\rfloor + 1, \ldots, N$ and $A(N) \triangleq \max_{S_N\cup\{t_N\}}\{\sum_{j\in S_N} a_j\hat{w}_j + (\Gamma - \lfloor\Gamma\rfloor)a_{t_N}\hat{w}_{t_N}\}$. Define constant demands $\widetilde{w}_k$ as

$$\widetilde{w}_k = \begin{cases} \overline{w}_k + \hat{w}_k, \ k = 1, \ldots, \lfloor\Gamma\rfloor, \\ \overline{w}_{\lfloor\Gamma\rfloor+1} + M(\lfloor\Gamma\rfloor + 1) - \sum_{j=1}^{\lfloor\Gamma\rfloor} \hat{w}_j, \ k = \lfloor\Gamma\rfloor + 1, \\ \overline{w}_k + M(k) - M(k-1), \ k = \lfloor\Gamma\rfloor + 2, \ldots, N. \end{cases}$$

The following proposition shows that the robust inventory control problem (14) admits an optimal base-stock policy; the proof is left out due to space limitations.

Proposition III.1 *Assume $x_1 < \sum_{k=1}^{N} \widetilde{w}_k$. The following base-stock policy is optimal to (14):*

$$u_k^* = \begin{cases} \widetilde{w}_k - x_k, & \text{if } x_k < \widetilde{w}_k, \\ 0, & \text{otherwise.} \end{cases} \qquad (15)$$

The cost of the optimal policy is $c(\sum_{k=1}^{N} \widetilde{w}_k - x_1) + h\sum_{k=1}^{L}(x_1 - \sum_{j=1}^{k} \widetilde{w}_j) + h\sum_{k=1}^{\lfloor\Gamma\rfloor}(\lfloor\Gamma\rfloor - k + 1)\hat{w}_k + h\sum_{k=\lfloor\Gamma\rfloor+1}^{N-1} M(k) + a_N hM(N) + hA(N)$, where $L = \max\{k \mid x_1 - \sum_{j=1}^{k} \widetilde{w}_j \geq 0\}$.

The robust inventory control problem (14) is not an LP problem because of the $\max\{\cdot\}$ functions. However, using LP duality it is possible to construct an equivalent LP problem from which we can obtain an optimal $\mathbf{u}^*$ efficiently. The following proposition states the result; the proof is omitted due to space limitations.

Proposition III.2 *The robust inventory control problem (14) is equivalent to the following LP problem:*

$$\begin{aligned} \min\ & z \qquad\qquad (16)\\ \text{s.t.}\ & \sum_{j=1}^{k} u_j \geq -x_1 + \sum_{j=1}^{k} (\overline{w}_j + \hat{w}_j), \quad k = 1, \ldots, \lfloor\Gamma\rfloor \\ & \sum_{j=1}^{k} u_j \geq -x_1 + \sum_{j=1}^{k} \overline{w}_j + \Gamma p_k + \sum_{j=1}^{k} q_{kj}, \\ & \qquad\qquad k = \lfloor\Gamma\rfloor + 1, \ldots, N \\ & p_k + q_{kj} \geq \hat{w}_j, \quad j = 1, \ldots, k, \quad k = \lfloor\Gamma\rfloor + 1, \ldots, N \\ & z - \sum_{k=1}^{N} (c + ha_k) u_k \geq Nhx_1 - h\sum_{k=1}^{N} a_k \overline{w}_k \\ & \qquad + h\left(\Gamma r + \sum_{k=1}^{N} s_k \right) \\ & r + s_k \geq a_k \hat{w}_k, \quad k = 1, \ldots, N \\ & p_k,\ q_{kj} \geq 0, \quad j = 1, \ldots, k, \quad k = \lfloor\Gamma\rfloor + 1, \ldots, N \\ & r,\ s_k,\ u_k \geq 0, \quad k = 1, \ldots, N. \end{aligned}$$

B. Bounds on the QoS Constraint Violation Probability

Assume $\Gamma > 0$. Let $\mathbf{u}^*$ be an optimal solution to (14), which can be obtained by solving (16). We examine the probabilities that $\mathbf{u}^*$ violates the QoS constraints $\sum_{j=1}^{k} u_j \geq -x_1 + \sum_{j=1}^{k} w_j$, $k = 1, \ldots, N$. Clearly none of the constraints for $k = 1, \ldots, \lfloor\Gamma\rfloor$ are violated at $\mathbf{u}^*$ because Γ provides full protection for those constraints. For $k = \lfloor\Gamma\rfloor + 1, \ldots, N$, the probability that the kth period QoS constraint is violated at $\mathbf{u}^*$ can be upper bounded as the following theorem describes; the result is analogous to Theorem. II.5(a), Corollary II.7, and Theorem. II.6(b), and the proof is omitted. Let $z_k = (w_k - \overline{w}_k)/\hat{w}_k$, $k = 1, \ldots, N$, and let $\Lambda_{z_k}(\theta) \triangleq \log E[e^{\theta z_k}]$ denote the logarithmic moment generating function of z_k.

Theorem III.3 *(a) Suppose that demands w_k satisfy Assumption B. Then for $k = \lfloor\Gamma\rfloor + 1, \ldots, N$*

$$P\left[\sum_{j=1}^{k} u_j^* < -x_1 + \sum_{j=1}^{k} w_j \right] \leq \exp\left[-\sup_{\theta\geq 0} \left[\theta\Gamma - \sum_{j=1}^{k} \Lambda_{z_j}(\theta) \right] \right]. \quad (17)$$

(b) Let $C_k(\mathbf{u}^*) = \sum_{j=1}^{k} u_j^* + x_1 - \sum_{j=1}^{k} \overline{w}_j$. Then for $k = \lfloor\Gamma\rfloor + 1, \ldots, N$

$$P\left[\sum_{j=1}^{k} u_j^* < -x_1 + \sum_{j=1}^{k} w_j \right] \leq \exp\left[-\sup_{\theta\geq 0} \left[\theta C_k(\mathbf{u}^*) - \sum_{j=1}^{k} \Lambda_{z_j}(\theta\hat{w}_j) \right] \right]. \quad (18)$$

(c) bound (18) $\leq$ bound (17).

Under Assumption B, bound (6) developed in [1] can be tailored to the inventory control problem as follows.

$$P\left[\sum_{j=1}^{k} u_j^* < -x_1 + \sum_{j=1}^{k} w_j\right] \leq \exp\left[-\frac{\Gamma^2}{2k}\right]. \qquad (19)$$

By applying Theorem II.5(b), we can show bound (17) $\leq$ bound (19). Bound (17) depends on the probability distributions of w_j's, but not on the optimal solution $\mathbf{u}^*$. Bound (18), on the other hand, relies on both the probability distributions and $\mathbf{u}^*$. Moreover, bound (18) does not require the symmetry assumption (Assumption B). We can also show the following monotonicity property.

Lemma III.4 *Given k, bounds (17) and (18) are non-increasing functions of Γ.*

Bounds (17) and (18) can be used together in the following scheme. Assume that the probability distributions of w_k's are symmetric. For simplicity, further assume $x_1 = 0$. Suppose there is a period, say the kth period, for which we want the QoS constraint violation probability to be kept below $p\%$. We need to determine $\mathbf{u}^*$ that satisfies this requirement. One option is to use any $\Gamma \geq k$ and solve (16) because the resulting $\mathbf{u}^*$ guarantees a zero violation probability for the kth period. However, this may be unnecessarily conservative resulting in a high cost. To obtain a better solution, we follow the procedure outlined below.

First using (17), we determine the minimum Γ for which the requirement is achieved. Since bound (17) is monotone in Γ, we can use a binary search. Let $\overline{\Gamma}$ be the result. We then solve (16) with $\overline{\Gamma}$ and obtain $\mathbf{u}^*$. Using $\mathbf{u}^*$ we compute bound (18) for the kth period. Due to Theorem III.3(c), it will be no greater than $p\%$. So we choose a new Γ in the interval $[0, \overline{\Gamma}]$, solve (16), and compute bound (18). If the bound is larger than $p\%$, we increase Γ. Otherwise we decrease Γ. We iterate in this fashion until we find a solution $\mathbf{u}^*$ (and optimal Γ) for which bound (18) yields a value smaller but sufficiently close to $p\%$. Note that a binary search can also be employed in choosing Γ during these iterations because bound (18) is also monotone in Γ.

C. Numerical Tests

We set $N = 30$, $c = 2$, $h = 1$, and $x_1 = 0$. Let $\overline{\mathbf{w}} = (\overline{w}_1, \ldots, \overline{w}_N)$ and $\hat{\mathbf{w}} = (\hat{w}_1, \ldots, \hat{w}_N)$. The value of $\overline{w}_k$ is randomly drawn from the range $[N, 10N]$ for all k. The value of $\hat{w}_k$ is randomly chosen from the range $[1, \overline{w}_k - 1]$ for all k. In this manner, we generate 20 instances of $(\overline{\mathbf{w}}, \hat{\mathbf{w}})$. For simplicity, we assume that all w_k have the same type of probability distribution, i.e., all z_k are identically distributed. We consider three types of symmetric probability distribution: triangle (T), uniform (U), and reverse-triangle (R). For each period $k = 1, \ldots, N$, the QoS constraint violation probability is required to be no more than 5%. Let Γ_D, Γ_S, Γ_B be the minimum values of Γ for which bounds (17), (18), (19) are less than 5% for all k, respectively. Since bounds (17) and (19) are non-decreasing functions of k for given Γ, it suffices to consider the Nth period to determine Γ_D and Γ_B. A simple calculation yields $\Gamma_B = 13.42$. Γ_D depends on the probability distributions of w_k's, and it is equal to 5.45, 7.68, and 9.38 for the triangle, uniform, and reverse-triangle distributions, respectively. We cannot determine Γ_S *a priori* because it depends on $\mathbf{u}^*$. It should be determined *on-line* as we described in the previous subsection.

TABLE II

ROBUST INVENTORY CONTROL PROBLEM OBJECTIVE VALUES AND Γ_S.

Dist.	$z_R(\Gamma_B)$	$z_R(\Gamma_D)$	$z_R(\Gamma_S)$	Γ_S	$\frac{z_R(\Gamma_D)}{z_R(\Gamma_B)}$	$\frac{z_R(\Gamma_S)}{z_R(\Gamma_B)}$
T	75848	51868	35192	2.55	0.684	0.464
U	75848	60976	43322	3.81	0.804	0.571
R	75848	66467	49025	4.85	0.876	0.646

For each instance of $(\overline{\mathbf{w}}, \hat{\mathbf{w}})$, we solve (16) with Γ_B and Γ_D; we also determine Γ_S and solve (16) with Γ_S. In Table II, we list the objective values of (16) corresponding to Γ_B, Γ_D, and Γ_S, averaged over the 20 instances of $(\overline{\mathbf{w}}, \hat{\mathbf{w}})$. The average values of Γ_S are also shown in the table. It can be seen that the distributional information we use and the tighter bounds we provide can lead to significant cost savings (up to 54% in these examples) over the distribution-free robust approach of [1].

IV. CONCLUSIONS

Using the probability distributions of uncertain data, we showed that we can (drastically) improve the quality of a solution to the robust problem without compromising its robustness (quantified as the probability that the solution becomes infeasible). To that end, we derived a new bound on the constraint violation probability. This bound is distribution-dependent, but is independent of the solution. We showed that the bound is tighter than a distribution-free bound given in [1]. We also derived another bound that relies on the solution. We showed that this solution-dependent bound is tighter (and in our numerical tests significantly so) than all other bounds.

As an application, we considered a discrete-time stochastic inventory control problem with QoS constraints. We constructed a robust formulation for the inventory control problem and showed that its optimal ordering policy is a base-stock policy. We derived two bounds on the probability that the QoS constraints are violated. We explained how these two bounds can be used systematically to obtain a better solution of the inventory control problem. In some of the examples we provided, cost savings amount to 54%.

REFERENCES

[1] D. Bertsimas and M. Sim, "The price of robustness," *Operations Research*, vol. 52, no. 1, pp. 35–53, 2004.

[2] A. Soyster, "Convex programming with set-inclusive constraints and applications to inexact linear programming," *Operations Research*, vol. 21, pp. 1154–1157, 1973.

[3] A. Ben-Tal and A. Nemirovski, "Robust solutions of uncertain linear programs," *Operations Research Letters*, vol. 25, pp. 1–13, 1999.

[4] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," Preprint, 2004.

[5] G. Calafiore and M. Campi, "Uncertain convex programs: randomized solutions and confidence levels," *Mathematical Programming*, vol. 102, pp. 25–46, 2005.

[6] D. Bertsimas and A. Thiele, "A robust optimization approach to supply chain management," Preprint, 2003.