The α-Reliable Mean-Excess Regret Model for Stochastic Facility Location Modeling

Gang Chen,[1] Mark S. Daskin,[2] Zuo-Jun Max Shen,[3] Stanislav Uryasev[4]

[1] *Bank of America, Optimization and Capacity Modeling, 1401 Elm, TX1-099-04-08, Dallas, Texas 75202*

[2] *Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208*

[3] *Department of Industrial Engineering and Operations Research, University of California, Berkeley, California 94720*

[4] *Department of Industrial and Systems Engineering, University of Florida, Gainesville, Florida 32611*

Received 13 August 2004; revised 25 April 2006; accepted 10 June 2006
DOI 10.1002/nav.20180
Published online 28 July 2006 in Wiley InterScience (www.interscience.wiley.com).

Abstract: In this paper, we study a strategic facility location problem under uncertainty. The uncertainty associated with future events is modeled by defining alternative future scenarios with probabilities. We present a new model called the α-reliable mean-excess model that minimizes the expected regret with respect to an endogenously selected subset of worst-case scenarios whose collective probability of occurrence is no more than $1-\alpha$. Our mean-excess risk measure is coherent and computationally efficient. Computational experiments also show that the α-reliable mean-excess criterion matches the α-reliable minimax criterion closely. © 2006 Wiley Periodicals, Inc. Naval Research Logistics 53: 617–626, 2006

Keywords: location model; p-median; stochastic; scenario modeling; risk management

1. INTRODUCTION

Supply chain network design decisions are usually strategic and once implemented they are difficult to reverse. During the time when the design decisions are in effect, many decision parameters—demands, costs—may change dramatically. This calls for models that address the inherent uncertainties of facility location problems. Unfortunately, most facility location models in the literature are static and deterministic. In the past few decades, researchers have dealt with the uncertainties by defining a number of possible future scenarios. Facility sites that either optimize the expected performance or optimize the worst-case performance over all the scenarios are recommended. However, such approaches may not be practical since, in real life, facilities are not typically designed for the average case or the worst-case scenario. For example, airports are never sized for an average day, since doing so would result in significant under-capacity much of the time. On the other hand, airports are never sized for the peak travel day, e.g., the Sunday of Thanksgiving weekend in the United States, since doing so would be prohibitively expensive.

Daskin, Hesse, and ReVelle [5] proposed a model called the α-reliable Minimax regret and applied it to the p-median problem. In this model, the maximum regret is computed over an endogenously selected subset of scenarios, called the reliability set, whose collective probability of occurrence is at least some user-defined value α. By minimizing the α-reliable maximum regret, the planner can be $100\alpha\%$ sure that the regret realized will be no more than that found by the model. Since the α-reliable maximum regret is essentially the α-quantile of the regrets and does not assess the magnitude of the regrets associated with the scenarios that are not included in the reliability set (the worst cases, or the tail), users of this model do not concern themselves with the possibility that the regrets in the tail are excessively higher than the α-reliable maximum regret. Although there are situations where it is appropriate to minimize

Correspondence to: Z.-J. Max Shen (shen@ieor.berkeley.edu).

the α-reliable maximum regret rather than the average or worst-case regret, computationally the α-reliable Minimax regret is very difficult to obtain, which limits its use in real life.

In this paper, we present a new model called the α-reliable Mean-excess regret model, or Mean-excess model for short. In contrast to the α-reliable Minimax model where the regret that defines the α-quantile of all regrets is minimized, in the new model we minimize the expectation of the regrets associated with the scenarios in the tail, which has a collective probability of 1-α. The α-reliable Mean-excess regret metric explicitly accounts for the magnitude of the regrets in the tail. Compared with the α-reliable Minimax model, the α-reliable Mean-excess regret model is computationally much easier to solve, making it easier to apply to practical situations.

The rest of this paper is organized as follows. In Section 2, we briefly review some of the literature on scenario modeling in the context of stochastic facility location. In Section 3, we formulate our new model and compare it with the α-reliable Minimax model. Computational results are presented in Section 4. Finally, we conclude and propose future directions of research in Section 5.

2. LITERATURE REVIEW

In the past few decades, researchers have used scenario planning to deal with the uncertainties in strategic facility location. In scenario planning, the decision maker identifies a number of future possible scenarios and estimates the likelihood of each scenario occurring. Scenario planning was chosen primarily because, as pointed out by Snyder, Daskin, and Teo [22], it allows the decision makers to model dependence among random parameters. For example, at a future time the demands at different locations may be correlated. Similarly, costs may be correlated. If a continuous approach is used to model such correlations, then the problem tends to become intractable.

Sheppard [19] was among the first to use scenario planning to model uncertainties in facility location. His model gives a siting plan that minimizes the expected cost over all scenarios. Schilling [16] proposes an approach in which the initial location decisions are those that are common across all (or most) scenarios' optimal plans. He suggests delaying other decisions until uncertainty is resolved. Daskin, Hopp, and Medina [6] demonstrate that this approach can lead to the adoption of the worst possible initial decision under conditions of future uncertainty. They propose a forecast horizon-based approach to facility planning over time.

Regret is a commonly used metric in decision-making with uncertainties. As defined by Zeelenberg [23], regret is a negative, cognitively based emotion experienced by individuals when they realize or imagine that their present situation could have been more positive if they had behaved differently. Psychological studies suggest that the anticipation of regret induces people to make more rational choices (e.g., [9]). If there is no knowledge about the probabilities of the possible outcomes, the Minimax regret principle can be useful to help people in making decisions. However, if there is knowledge about these probabilities, then the Minimax regret principle can be suboptimal. Recently, regret theory in the economics literature has taken the probability of regret into consideration (e.g., [2, 10]), and quantile-related objectives are widely used in areas such as the financial and insurance industries. For example, Value-at-Risk (VaR), defined as a quantile of potential losses, is by far the most popular and most accepted risk measure among financial institutions [7]. VaR provides information about the magnitude of losses that will not be exceeded with a certainty probability. In the same spirit, the quality of service in an inventory system can be measured as the probability of not stocking out in the service period [12]. However, VaR is not a coherent risk measure [1]. For example, the VaR of a portfolio of two stocks may end up being greater than the sum of the VaRs of the individual stocks. Furthermore, it is not easy to optimize VaR. In the next section we will introduce a coherent quantile-based risk measure, the conditional value-at-risk, and discuss how to use it to model facility location decisions.

In the context of stochastic facility location, the regret associated with each scenario under a given siting plan is usually defined as the difference between the objective function value when the siting plan is chosen to optimal solution for that scenario and the objective function value when the siting plan is chosen to be the given siting plan.

Ghosh and McLafferty [8] propose a model in which either the sum of the regrets or the sum of the squared regrets over all scenarios is minimized. Note that the objective of minimizing the sum of the regrets is equivalent to minimizing the expected regret with all scenarios having the same probability.

Serra and Marianov [17] look at a problem in which the parameters of the network, including travel times, demands, and distances, change over the course of a day. They model each period of the day as a scenario and identify either solutions that minimize the maximum travel time over the scenarios or solutions that minimize the total regret. The regret of each scenario is defined to be the difference between the objective function values given by the overall compromise solution and the optimal solution for that single scenario. Serra, Ratick, and ReVelle [18] study a maximum capture problem where the objective is to select the locations of servers for an entering firm that wishes to maximize its market share in a market where competitors are already in position. Their models either maximize the minimum capture associated with any scenario or minimize the expected regret over all the scenarios. Current, Ratick, and ReVelle [3] study problems where the total number of facilities to be located is uncertain and the objective is either to minimize the expected opportunity loss or to minimize the maximum opportunity loss. The opportunity loss is defined as the difference between the objective function value when the initial facility locations must be included in the final siting plan and the objective function value when there is no such constraint. The minimum expected opportunity loss criterion finds the initial set of facility locations that minimize the expected opportunity losses across all scenarios. The Minimax opportunity loss criterion finds the initial facility locations such that the maximum loss is minimized over all scenarios.

More recently, Snyder, Daskin, and Teo [22] study the stochastic location problem with risk pooling, which seeks to locate distribution centers to minimize the total fixed location costs, transportation costs, and inventory costs. They propose a model that minimizes the expected cost of the system across all scenarios and develop a Lagrangian-relaxation-based exact algorithm to solve the model. Snyder [20] extends this model to constrain the maximum relative regret in any scenario. For a more comprehensive review of recent dynamic and stochastic facility location problems, the reader is referred to Owen and Daskin [13] and Snyder [21].

3. THE α-RELIABLE MINIMAX AND MEAN-EXCESS MODELS

In this section we present and compare the α-reliable Minimax model and the α-reliable Mean-excess model in the context of the p-median problem. The classical deterministic p-median problem seeks to locate p facilities relative to a set of demand nodes such that the sum of the shortest demand weighted distance between demand nodes and facilities is minimized (see [4]). In the stochastic p-median problem the distances between demand nodes and facilities are stochastic, as are the demands at the demand nodes. The motivation for choosing the p-median problem is twofold. First, it is one of the key building block models of virtually all location models including the uncapacitated and capacitated fixed charge location models as well as many others. Second, it is the problem used by Daskin, Hesse, and ReVelle [5] in studying the the α-reliable Minimax regret and thus facilitates an "Apple-to-Apple" comparison between the two models. Before presenting the models, we define the following notation:

$i = 1, \ldots, m$: index of demand nodes
$j = 1, \ldots, n$: index of candidate locations
$k = 1, \ldots, K$: index of possible scenarios
h_{ik} : the demand at node i under scenario k
d_{ijk} : distance from node i to candidate site j under scenario k
p : number of facilities to locate
$\hat{V}_k$: best p-median value that can be obtained under scenario k, namely, the minimum demand-weighted total distance under scenario k
q_k : the probability that scenario k will occur
m_k : a large constant specific to scenario k such that $m_k \geq R_k$
α : desired reliability level

We define the following decision variables:

$$x_j = \begin{cases} 1: & \text{if we locate at candidate node } j \\ 0: & \text{otherwise} \end{cases}$$

$$y_{ijk} = \begin{cases} 1: & \text{if demand node } i \text{ is assigned to a facility at } j \text{ under scenario } k \\ 0: & \text{otherwise} \end{cases}$$

$$z_k = \begin{cases} 1: & \text{if scenario } k \text{ is included in the set over which the maximum regret is minimized} \\ 0: & \text{otherwise} \end{cases}$$

R_k : regret associated with scenario k and the current solution (x, y) (may not be optimal for scenario k). $R_k = \sum_i \sum_j h_{ik} d_{ijk} y_{ijk} - \hat{V}_k$

The α-reliable Minimax model can be formulated as follows:

$$\text{Minimize} \quad W \tag{1}$$

$$\text{subject to :} \quad \sum_{j=1}^{n} x_j = p \tag{2}$$

$$\sum_{j=1}^{n} y_{ijk} = 1, \quad \forall\, i, k \tag{3}$$

$$y_{ijk} - x_j \leq 0, \quad \forall\, i, j, k \tag{4}$$

$$R_k - \left(\sum_{i=1}^{m} \sum_{j=1}^{n} h_{ik} d_{ijk} y_{ijk} - \hat{V}_k \right) = 0, \quad \forall\, k \tag{5}$$

$$\sum_{k=1}^{K} q_k z_k \geq \alpha \tag{6}$$

$$W - R_k + m_k(1 - z_k) \geq 0, \quad \forall\, k \tag{7}$$

$$x_j \in \{0, 1\}, \quad \forall j \tag{8}$$

$$y_{ijk} \in \{0, 1\}, \quad \forall\, i, j, k \tag{9}$$

$$z_k \in \{0, 1\}, \quad \forall\, k. \tag{10}$$

The objective function (1) minimizes the α-reliable maximum regret. Constraint (2) stipulates that exactly p facilities are to be located. Constraint (3) states that each demand node must be assigned to exactly one facility in each scenario. Constraint (4) states that demands at i cannot be assigned to a facility j under scenario k unless a facility is located at node j. Constraint (5) defines the regret associated with

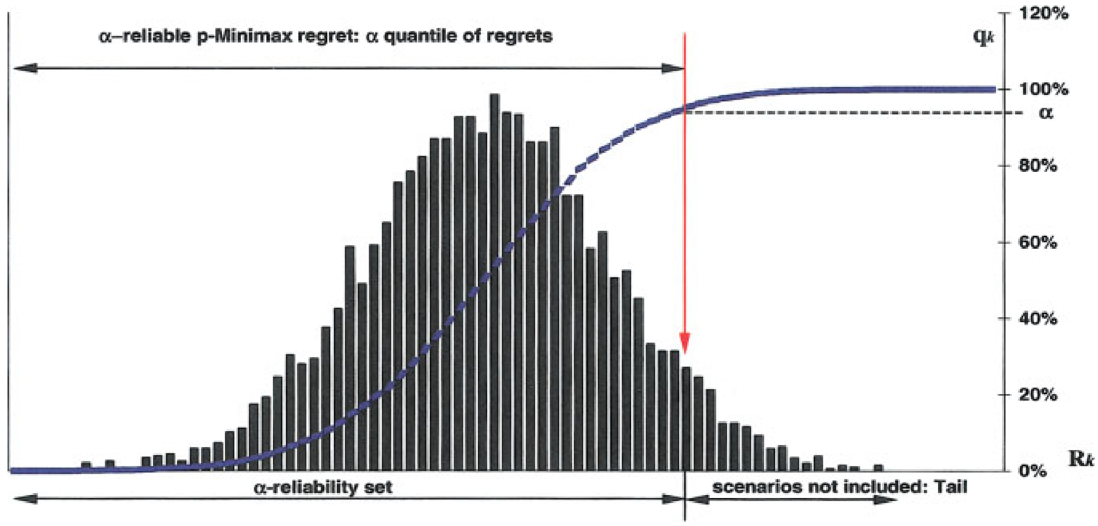

Figure 1. Illustration of the Minimax model. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

scenario k, as discussed previously. Constraint (6) stipulates that the probability associated with the set of scenarios over which the maximum regret is computed must be at least α. Constraint (7) defines the maximum regret in terms of the individual scenario regrets and the variables, z_k, which indicate which scenarios are to be included in the maximum regret computation. Thus, the Minimax model minimizes the maximum regret over a subset of the possible scenarios, with the added stipulation that the probability of realizing a scenario that is not included in the subset must be at most $1-\alpha$. In addition, by varying α over an appropriate range, the decision maker can identify a portfolio of siting plans.

Figure 1 illustrates the Minimax model. The black bars in the chart represent the K scenarios and are aligned from left to right in increasing order of their regrets (based on the current solution). The height of each black bar represents the probability of the corresponding scenario. The thick dotted line represents the cumulative probability of the scenarios. The α-reliable maximum regret (based on the current solution) is the regret of the scenario (highlighted by the vertical arrow) that corresponds to a cumulative probability of α. All the scenarios to the left of the vertical arrow form the α-reliability set, which has a collective probability of at least α. All the scenarios to the right of the vertical arrow form the tail, which has a collective probability of no more than $1-\alpha$. As the siting plan changes, the relative order of the scenarios changes and so does the α-reliable maximum regret. The Minimax model seeks a siting plan such that the α-reliable maximum regret is minimized.

Clearly, the α-reliable Minimax regret model does not assess the magnitude of the regrets associated with the scenarios that are not included in the α-reliable set and does not distinguish between situations where the regrets in the tail are only a little bit worse than the α-reliable maximum regret and those in which the regrets in the tail are overwhelmingly higher. In addition, mathematically the α-reliable Minimax model is difficult to solve. The α-reliable Minimax model has earlier been studied in the stochastic programming literature, albeit not in a facility location context. Mauser and Rosen [11] showed that the α-reliable maximum regret is a nonsmooth, nonconvex, and multiextreme function of the decision variables (x, y) over the feasible region. This is intuitive, since a small perturbation in the siting plan may cause abrupt jumps in the regrets associated with the K scenarios. Therefore, the α-reliable Minimax regret does not change smoothly as the siting plan changes and is a multiextreme and nonconvex function of (x, y).

In this paper we present a different model, the α-reliable Mean-excess regret, which minimizes the expected regret (the probability-weighted regret) with respect to an endogenously selected subset of worst-case scenarios whose collective probability of occurrence is no more than $1-\alpha$. To present the new model, we need the following additional definitions:

Ψ : the p-median feasibility constraint set, namely, the set defined by constraints (2), (3), (4), (5), (8), and (9)

X : decision variable (x, y)

$R(X, k)$: regret as a function of X and scenario index k

$f(X, \zeta) = P\{k \mid R(X, k) \le \zeta\}$: with X fixed, the collective probability of those scenarios in which the regret does not exceed ζ

$\zeta_\alpha(X) = \min\{\zeta \in \mathbf{R} : f(X,\zeta) \geq \alpha\}$: with X fixed, the minimum value ζ such that $f(X,\zeta) \geq \alpha$, namely, the α-quantile of the regrets of the K scenarios

$\phi_\alpha(X) = \dfrac{\sum_{k:R(X,k)>\zeta_\alpha(X)} q_k \cdot R(X,k)}{\sum_{k:R(X,k)>\zeta_\alpha(X)} q_k}$: with X fixed, the conditional probability-weighted average of the regrets strictly exceeding $\zeta_\alpha(X)$

k_α : index of the scenario such that $R(X,k_\alpha) = \zeta_\alpha(X)$

Using this notation, it is easy to see that the α-reliable Minimax regret model can be rewritten as follows:

$$\begin{aligned} &\text{Minimize} \quad \zeta_\alpha(X) \qquad (11)\\ &\text{subject to :} \quad X \in \Psi. \end{aligned}$$

In contrast, the α-reliable Mean-excess regret model can be formulated as follows:

$$\begin{aligned} &\text{Minimize} \quad \lambda\zeta_\alpha(X) + (1-\lambda)\phi_\alpha(X) \qquad (12)\\ &\text{subject to :} \quad X \in \Psi, \end{aligned}$$

where $\lambda = \frac{[f(X,\zeta_\alpha(X))-\alpha]}{1-\alpha} \in [0,1]$. Figure 2 illustrates the α-reliable Mean-excess regret model.

Since both the regret function $R(X,k)$ and the feasibility set Ψ are convex with respect to X, it can be shown that (see [14, 15]) formulation (12) can be reduced to the following problem:

$$\begin{aligned} &\text{Minimize} \quad F_\alpha(X,\zeta) = \zeta + \frac{1}{1-\alpha} \times \sum_{k=1}^{K} q_k \cdot \text{Max}\{[R(X,k)-\zeta],0\} \qquad (13)\\ &\text{subject to :} \quad X \in \Psi, \end{aligned}$$

where ζ is a free variable. Hence, the α-reliable Mean-excess regret model for the p-median problem can be formulated as the following mixed integer problem:

$$\text{Minimize} \quad F_\alpha((x,y),\zeta) = \zeta + \frac{1}{1-\alpha}\sum_{k=1}^{K} q_k U_k \qquad (14)$$

$$\begin{aligned} \text{subject to :} \quad & \sum_{j=1}^{n} x_j = p \\ & \sum_{j=1}^{n} y_{ijk} = 1, \quad \forall\, i,k \\ & y_{ijk} - x_j \leq 0, \quad \forall\, i,j,k \\ & R_k - \left(\sum_{i=1}^{m}\sum_{j=1}^{n} h_{ik} d_{ijk} y_{ijk} - \hat{V}_k\right) = 0, \quad \forall\, k \\ & U_k \geq R_k - \zeta, \quad \forall\, k \qquad (15)\\ & x_j \in \{0,1\}, \quad \forall\, j \\ & y_{ijk} \in \{0,1\}, \quad \forall\, i,j,k \\ & U_k \geq 0, \quad \forall\, k. \qquad (16) \end{aligned}$$

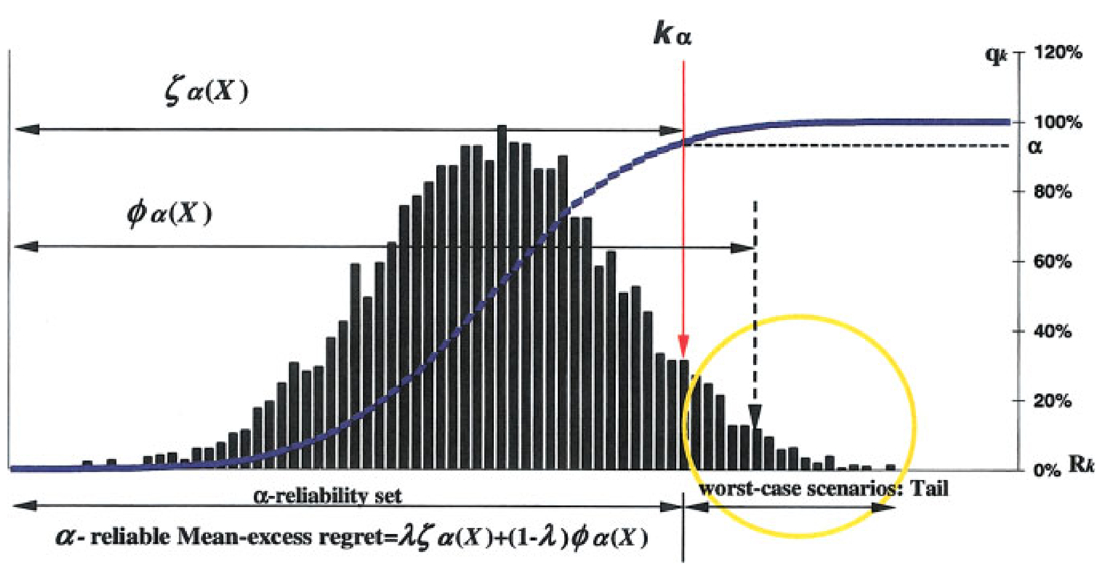

Figure 2. Illustration of the Mean-Excess model. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

$F_\alpha((x, y), \zeta)$ is a convex function of both ζ and (x, y), which makes the α-reliable Mean-excess model easier to solve than the α-reliable regret model (see [14, 15]).

It is important to note that the α-reliable Mean-excess regret is closely related to the α-reliable maximum regret. Minimizing $F_\alpha((x, y), \zeta)$ gives both the optimal α-reliable Mean-excess regret and the corresponding (non-optimal) α-reliable Minimax regret. Specifically, when the α-reliable Mean-excess regret is minimized, the value of ζ^* gives the corresponding (non-optimal) α-reliable Minimax regret. In addition, by definition, the α-reliable Mean-excess regret is an upper bound of the corresponding α-reliable maximum regret. Therefore, minimizing the α-reliable Mean-excess regret will also lead to a low α-reliable maximum regret.

4. COMPUTATIONAL RESULTS

In this section, we summarize our computational results with both the α-reliable Minimax regret model and the α-reliable Mean-excess regret model outlined above. For the sake of abbreviation, we will call these two models Minimax and Mean-excess, respectively. All of our computational experiments are based on the data found in [4] and [5] for 88 major U.S. cities. Specifically, we use the nine scenarios of the 88-city problem found in [5] to generate more scenarios to be used in our computational experiments. For example, to have K ($K \geq 9$) alternative scenarios, we generate $\lfloor K/9 \rfloor$ scenarios from each of the original 9 scenarios in [5], using a normal distribution in which each city has a mean demand equal to the demand of that city in the original scenario and a standard deviation equal to 1/10 of the demand of that city in the original scenario. The remaining $K - 9\lfloor \frac{K}{9} \rfloor$ scenarios are then generated from the original scenario No. 5. The probability of occurrence associated with each scenario is first generated with a uniform distribution Uniform(0,1] and then normalized such that the total probability of all the scenarios is equal to 1.

All of our tests involve siting five facilities. In all runs, all 88 demand nodes are also eligible candidate facility sites. In addition, only the demands differ between any two scenarios. The distance between any two cities is scenario-independent. This allows us to test problems with a larger number of scenarios. Therefore, in both the Minimax model and the Mean-excess model we replace y_{ijk} by y_{ij}, thereby significantly reducing the number of decision variables.

Note that, in both the Minimax model and the Mean-excess model, the $\hat{V}_k$ values, i.e., the minimum sum of the demand-weighted distances attainable for scenario k ($k = 1, 2, \ldots, K$), are required as inputs parameters. We obtain these values by optimally solving the p-median problem for each of the K scenarios.

In addition, in the Minimax model, the m_k values—the upper bound of the largest possible regret for each scenario k ($k = 1, 2, \ldots, K$)—are also required as input parameters. We use three different procedures to obtain these values. In the first procedure (see [5]), for each scenario k, we compute the regret associated with the optimal locations found for each of the other $K - 1$ scenarios; m_k is then taken as the maximum of these quantities. This procedure requires solving $K(K-1) + 1$ sub-problems. In the second procedure, each m_k ($k = 1, 2, \ldots, K$) is set to a constant value that is large enough to be used as the upper bound of the R_k's (in our test we used the constant value of 3.0×10^{11}). In the third procedure, we compute the maximum possible regret associated with each scenario k by solving the p-median problem associated with it, with the objective function changed to maximize R_k. This procedure requires evaluating K solutions. To facilitate our discussion, we denote the Minimax models with the first, second, and third procedures of solving for the m_k's by MinimaxI, MinimaxII, and MinimaxIII, respectively.

Both models as well as their sub-problems are coded with C++ and run on a workstation with an Intel 3.06GHz Xeon CPU and 3.25GB of RAM. The problems are solved exactly by CPLEX version 8.1 and the operating system is Microsoft Windows XP Professional Edition. All the computational times presented are in seconds.

Table 1 presents the computational times for optimally solving the Mean-excess, MinimaxI, MinimaxII, and MinimaxIII models, with the reliability level α fixed at 95%. All computational times in Table 1 exclude the input/output times, as well as the time needed to solve for the $\hat{V}_k$'s and m_k's. Note that MinimaxII could not solve problems with more than 126 scenarios while MinimaxIII could not solve problems with more than 72 scenarios.

Figure 3 illustrates the computational times needed to optimally solve the Mean-excess, MinimaxI, MinimaxII, and MinimaxIII models as the number of scenarios increases. Note that, as the number of scenarios increases, the computational times of MinimaxI, MinimaxII, and MinimaxIII appear to increase exponentially while the computational time of Mean-excess appears to increase only linearly.

Regressing the Mean-excess time against the number of scenarios, we obtain the following linear equation to predict the Mean-excess solution time:

$$\text{Mean-excess Time} = \underset{(17.691)}{74.595} + \underset{(0.1011)}{0.8295}\ (\text{No. of Scenarios})$$

$$R^2 = 0.882,$$

where the numbers in parentheses under the estimated coefficients are the standard errors of the estimates. For the MinimaxI time, we obtain the following equation

$$ln(\text{MinimaxI Time}) = \underset{(0.2267)}{4.380} + \underset{(0.00149)}{0.01955}\ (\text{No. of Scenarios})$$

$$R^2 = 0.956$$

Table 1. Solution times for MinimaxI, MinimaxII, MinimaxIII, and Mean-excess.

	Solution time			
No. of scenarios	MinimaxI	MinimaxII	MinimaxIII	Mean-excess
9	72	90	95	85
27	82	92	114	48
45	349	619	722	137
72	288	1409	79343[a]	130
99	709	8738	188111[a]	194
126	1356	198298[a]	129305[a]	203
162	1317	[b]	119023[a]	212
198	4706	[b]	[b]	239
234	10225	[b]	[b]	218
279	12135	[b]	[b]	276
324	[b]	[b]	[b]	385

[a]Amount of time spent before CPLEX stopped without finding any solution.
[b]Not tested.

or

$$\text{MinimaxI Time} = 79.8214e^{0.01955\,(\text{No. of Scenarios})}.$$

Table 2 presents the total computational times needed for optimally solving the Mean-excess, MinimaxI, MinimaxII, and MinimaxIII models, with the reliability level α fixed at 95%. Specifically, the total computational times of the Mean-excess model include the time needed to solve for the $\hat{V}_k$'s and the time needed to solve the Mean-excess model. The total computational times of the MinimaxI, MinimaxII, and MinimaxIII models include the time needed to solve for the $\hat{V}_k$'s and m_k's as well as the time needed to solve the Minimax model. All the input/output times are excluded.

The total computational times in Table 2 are illustrated in Figure 4. The total solution time for the Mean-excess model again increases linearly in the number of scenarios as shown in the following regression (which has the highest R^2 value when comparing a linear time model with a polynomial time model and an exponential time model):

$$\text{Total MeanExcess Time} = \underset{(69.787)}{-134.695} + \underset{(0.340)}{6.379} \times (\text{No. of Scenarios}) \quad R^2 = 0.966.$$

The total time for the MinimaxI algorithm, however, increases very dramatically with the size of the problem. Of the three models tested, the polynomial time model gave the

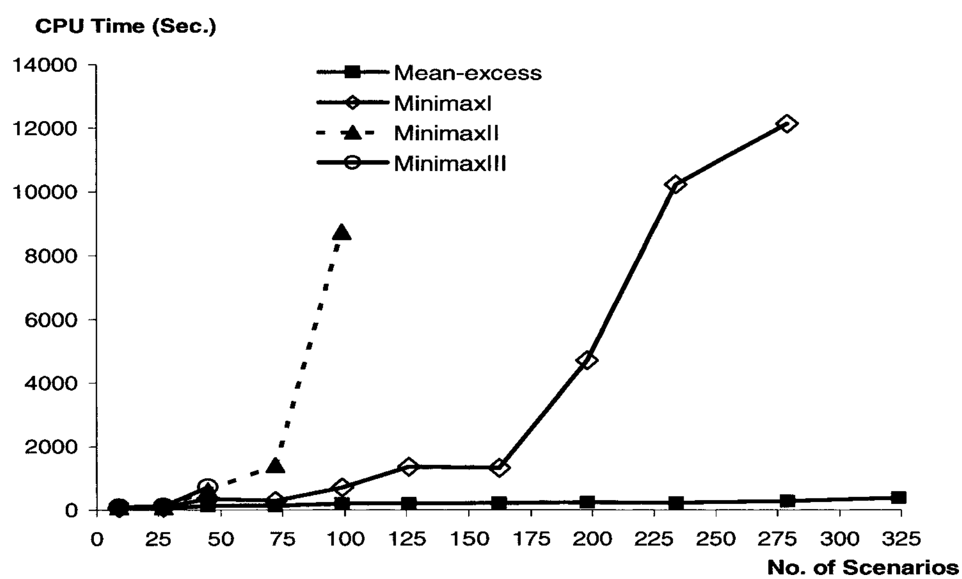

Figure 3. Solution times for MinimaxI, MinimaxII, MinimaxIII, and Mean-excess with $\alpha = 0.95$.

Table 2. Total solution times for MinimaxI, MinimaxII, MinimaxIII, and Mean-excess.

No. of scenarios	Total solution time			
	MinimaxI	MinimaxII	MinimaxIII	Mean-excess
9	97	101	110	96
27	435	131	170	87
45	1841	695	841	213
72	5912	1559	79586[a]	280
99	15657	8990	188528[a]	445
126	30716	198671[a]	129928[a]	577
162	64317	[b]	119988[a]	773
198	118394	[b]	[b]	1019
234	193056	[b]	[b]	1235
279	318242	[b]	[b]	1650
324	[b]	[b]	[b]	2191

[a]Amount of time spent before CPLEX stopped without finding any solution.
[b]Not tested.

highest R^2 resulting in an equation of

$$\underset{(0.5410)}{ln(\text{Total MinimaxI Time}) = -1.55871} + \underset{(0.1185)}{2.4742} \times ln(\text{No. of Scenarios}) \quad R^2 = 0.982$$

or

$$\text{Total MinimaxI Time} = 0.2104\ (\text{No. of Scenarios})^{2.742}$$

While this is not an exponential increase in time with problem size, the time does clearly increase very rapidly as the number of scenarios increases. A problem with 100 scenarios would take almost 300 times as much time to solve as would a problem with only 10 scenarios. With 279 scenarios, it takes more than 88 hours to solve MinimaxI optimally. In light of the rapidly increasing computational time of MinimaxI, it seems impractical to solve MinimaxI optimally for more than 279 scenarios. For MinimaxII and MinimaxIII, CPLEX was unable to find any solution for more than 126 and 72 scenarios, respectively. It is obvious that the input parameters m_k's have a significant impact on the solution time of the Minimax model. In particular, tighter upper bounds on the R_k's dramatically reduce the solution time of the Minimax model but obtaining these tighter bounds requires a significant amount of time.

Table 3 presents the objective function values for the Mean-excess model and the Minimax model. Note that the solution to the Minimax model is independent of which procedure is used to solve for the m_k's. Therefore, in Table 3 we only

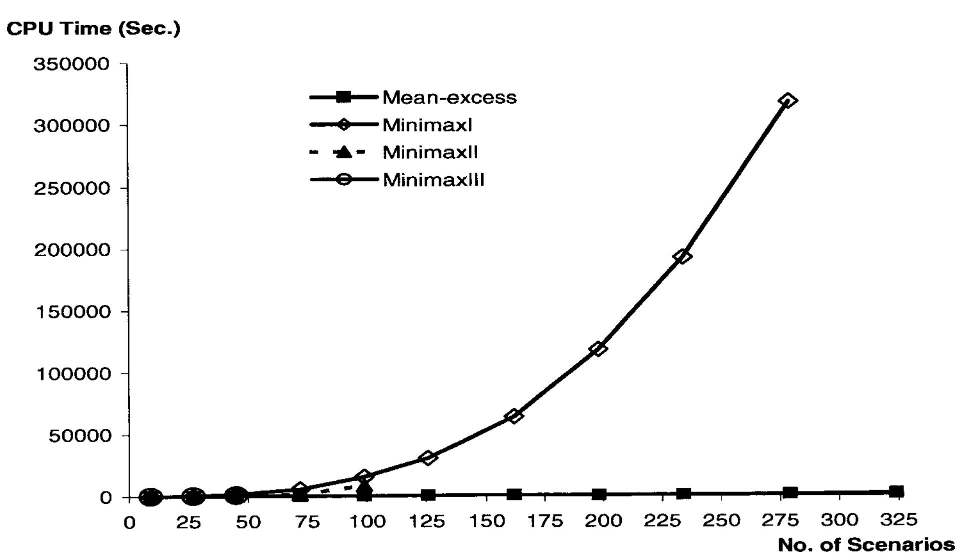

Figure 4. Total solution times for MinimaxI, MinimaxII, MinimaxIII, and Mean-excess with $\alpha = 0.95$.

Table 3. Solutions of Minimax and Mean-excess.

	Mean-excess model			Minimax model	
No. of scenarios	α-Reliable Mean-excess regret	α-Reliable Minimax regret	Worst-case regret	α-Reliable Minimax regret	Worst-case regret
9	2247428811	2247428811	2247428811	2247428811	2247428811
27	2258738311	2051996284	2321909291	2051996284	2321909291
45	2302417752	2200321339	2468654038	2200321339	2468654038
72	2405582073	2246449186	2532201475	2246449186	2532201475
99	2432596111	2270260483	2644634954	2270260483	2644634954
126	2346786023	2228805388	2648265518	2228805388	2648265518
162	2353539209	2214090364	2496477548	2199348673	2552036761
198	2495007119	2278354467	2827041127	2278354467	2827041127
234	2409139093	2243677946	2980885404	2243677946	2980885404
279	2455512957	2251328482	2869632915	2251328482	2869632915

present the solutions of MinimaxI, since it was able to solve problems with up to 279 scenarios. In addition, in Table 3 we also present the worst-case scenario regrets associated with the two models and the α-reliable Minimax regret associated with minimizing the α-reliable Mean-excess regret.

Table 3 shows that, in 9 of the 10 instances, minimizing the α-reliable Mean-excess regret also led to the optimal value of the α-reliable Minimax regret. In the instance with 162 scenarios, minimizing the α-reliable Mean-excess regret gives an α-reliable Minimax regret that is 0.67% higher than the optimal α-reliable Minimax regret. On the other hand, in each of the 10 instances tested, the Mean-excess model gives a worst-case scenario regret no bigger than the worst-case scenario regret given by the Minimax model. In the 162-scenario instance, the worst-case scenario regret given by the Mean-excess model is 2.23% lower than the worst-case scenario regret given by the Minimax model.

In most of the test problems, the solutions to the α-reliable Mean-excess model are identical to the solutions to the α-reliable Minimax model. This is not surprising, as in Section **3** we have observed that minimizing the α-reliable Mean-excess regret will also lead to a low α-reliable maximum regret. In addition, when the number of scenarios is small (less than 500), the α-reliable maximum regret and the α-reliable Mean-excess regret should be close. To see this, consider a problem with 10 equally likely scenarios, the 0.95-reliable maximum regret coincides with the 0.95-reliable Mean-excess regret.

5. DISCUSSIONS AND CONCLUSION

In this paper we present the α-reliable Mean-excess model for strategic facility location planning. In this framework, decision makers identify future scenarios and estimate the likelihood of each scenario occurring. The model then finds a solution that minimizes the expected regret with respect to an endogenously selected subset of worst-case scenarios whose collective probability of occurrence is no more than $1 - \alpha$. The α-reliable Mean-excess model is different from the previously proposed α-reliable Minimax model in that it explicitly account for the magnitude of the regrets in the the worst-case scenarios, making it appropriate for problems where the objective function value is not subject to any threshold. Moreover, the computational efficiency of the α-reliable Mean-excess model and its close relationship to the α-reliable Minimax model make it a useful tool for obtaining an approximate solution to the α-reliable Minimax model.

Coherence has become an important risk measure in recent years. Our mean-excess measure is coherent, and it can be applied to a wide range of applications in supply chain management, capacity planning, and financial engineering, etc. For instance, it can be used to design robust supply chains to hedge against uncertainties in demand, costs, or other parameters. In the area of hydro-electric power generation, the model can be used to minimize the cost associated with the costly startups and shutdowns of back-up thermal generating units. In the near future, we plan to extend this model to multi-stage and multi-dimensional environments and to include capacity constraints, fixed charges, scenario-dependent distances, and travel costs in our computational experiments.

ACKNOWLEDGMENTS

The authors thank Professor John Birge and two anonymous referees for their helpful comments, which led to the improved version. The work of Chen and Shen was partially supported by NSF Grants DMI-0348209. This support is gratefully acknowledged.

REFERENCES

[1] P. Artzner, F. Delbaen, J.M. Eber, and D. Heath, Coherent measures of risk, Math Finance 9 (1999), 203–228.
[2] D. Bell, Regret in decision making under uncertainty, Oper Res 30 (1982), 961–981.
[3] J.R. Current, S. Ratick, and C.S. ReVelle, Dynamic facility location when the total number of facilities is uncertain: A decision analysis approach, Eur J Oper Res 110 (1998), 597–609.
[4] M.S. Daskin, Network and discrete location: Models, algorithms and applications, Wiley, NY, 1995.
[5] M.S. Daskin, S.M. Hesse, and C.S. ReVelle, α-reliable P-Minimax regret: A new model for strategic facility location modeling, Locat Sci 5 (1997), 227–246.
[6] M.S. Daskin, W.J. Hopp, and B. Medina, Forecast horizons and dynamic facility location, Ann Oper Res 40 (1992), 125–151.
[7] D. Duffie and J. Pan, An overview of value-at-risk, J Deriv 4 (1997), 7–49.
[8] A. Ghosh and S.L. McLafferty, Locating stores in uncertain environments: A scenario planning approach, J Retail 58(4) (1982), 5–22.
[9] I.L. Janis and L. Mann, Decision making, The Free Press, New York, 1977.
[10] G. Loomes and R. Sugden, Regret theory: An alternative theory of rational choice under uncertainty, Econ J 92 (1982), 805–824.
[11] H. Mauser and D. Rosen, Beyond CVaR: From measuring risk to managing risk, Algor Res Q 1(2) (1998), 5–20.
[12] S. Nahmias, Production and operations analysis, 4th ed., Irwin/McGraw Hill, Burr Ridge, IL, 2001.
[13] S.H. Owen and M. S. Daskin, Strategic facility location: A review, Eur J Oper Res 111 (1998), 423–447.
[14] R.T. Rockafellar and S. Uryasev, Optimization of conditional value-at-risk, J Risk 2(3) (2000), 21–41.
[15] R.T. Rockafellar and S. Uryasev, Conditional value-at-risk for general loss distributions, J Bank Finance 26(7) (2002), 1443–1471.
[16] D.A. Schilling, Strategic facility planning: The analysis of options, Decision Sci 13 (1982), 1–14.
[17] D. Serra and V. Marianov, The p-Median problem in a changing network: The case of Barcelona, Locat Sci 6 (1998), 383–394.
[18] D. Serra, S. Ratick, and C. ReVelle, The maximum capture problem with uncertainty, Environ Plan B 26 (1996), 49–59.
[19] E.S. Sheppard, A conceptual framework for dynamic location–allocation analysis, Environ Plan A 6 (1974), 547–564.
[20] L.V. Snyder, Supply chain robustness and reliability: Models and algorithms, Ph.D. Dissertation, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, 2003.
[21] L.V. Snyder, Facility location under uncertainty: A review, Lehigh University Department of Industrial and Systems Engineering Technical Report 04T-015. Forthcoming, 2005.
[22] L.V. Snyder, M.S. Daskin, and C.-P. Teo, The stochastic location model with risk pooling, Lehigh University Department of Industrial and Systems Engineering Technical Report 04T-017, Forthcoming, 2002.
[23] M. Zeelenberg, Anticipated regret, expected feedback and behavioral decision making, J Behav Decision Making 12(2) (1999), 93–106.