

Facility Location Under Uncertainty: A Review

**Lawrence V. Snyder
Lehigh University**

Report No. 04T-015

Facility Location Under Uncertainty: A Review

Lawrence V. Snyder

Dept. of Industrial and Systems Engineering

Lehigh University

Bethlehem, PA, USA

July, 2004

Abstract

Plants, distribution centers, and other facilities generally function for years or decades, during which time the environment in which they operate may change substantially. Costs, demands, travel times, and other inputs to classical facility location models may be highly uncertain. This has made the development of models for facility location under uncertainty a high priority for researchers in both the logistics and stochastic/robust optimization communities. Indeed, a large number of the approaches that have been proposed for optimization under uncertainty have been applied to facility location problems.

This paper reviews the literature on stochastic and robust facility location models. Our intent is to illustrate both the rich variety of approaches for optimization under uncertainty that have appeared in the literature and their application to facility location problems. In a few instances for which examples in facility location are not available, we provide examples from the more general logistics literature.

1 Introduction

Facility location decisions are costly and difficult to reverse, and their impact spans a long time horizon. During the time when design decisions are in effect, any of the parameters of the problem—costs, demands, distances—may fluctuate widely. Parameter estimates may also be inaccurate due to poor measurements or to tasks inherent

in the modeling process like aggregating demands points and choosing a distance norm. Recognizing this, researchers have been developing models for facility location under uncertainty for several decades.

The two-stage nature of facility location problems—choose locations now, before we know what the future holds, and react once the uncertainty has been resolved, say, by assigning customers to facilities—has made these problems very attractive to researchers exploring approaches to decision making under uncertainty. A large number of these approaches have been applied to facility location problems.

This paper aims to illustrate the rich variety of approaches for optimization under uncertainty by examining their application to facility location problems. To this end, we have chosen to categorize papers first by their approach to uncertainty, and then by the nature of the facility location problem they discuss. A few of the measures we discuss have not, to the best of our knowledge, been applied to facility location problems. We include these measures in our survey because of their relationship to measures that we do discuss, or simply because of the novelty of their approach. In these cases, we provide examples from the literature on capacity planning, network design, or other logistics problems that, like facility location, involve a strategic phase during which capital investments are made, followed by a tactical phase, with uncertainty being resolved between the two.

In the interest of brevity, we have opted not to discuss the large body of literature on facility location problems with congested facilities, which attempt to capture the possibility that a customer may need service from a facility that is occupied with another customer. Such models are commonly used in the siting of ambulances, fire stations, and other emergency services. They attempt to guarantee adequate service either by requiring redundant coverage (as in Daskin 1982, 1983, ReVelle and Hogan 1989, and Ball and Lin 1993) or by explicitly considering the queuing aspect of the problem (as in Larson 1974, 1975, Berman, Larson, and Chiu 1985, and Marianov and ReVelle 1996). The reader is referred to the surveys by Daskin, Hogan and ReVelle (1988) and Berman and Krass (2002) for thorough discussions of these models. A related branch of literature considers models in which the facilities may be unable to provide service due to facility disruptions (Bundschuh, Klabjan and Thurston 2003, Berman, Krass and Menezes 2004, Snyder and Daskin 2004) or link failures (Nel and Colbourn 1990, Eiselt,

Gendreau and Laporte 1992).

Throughout this paper we assume that the reader has some familiarity with deterministic facility location theory. For an introduction to this topic, the reader is referred to the texts by Daskin (1995), Drezner (1995), or Drezner and Hamacher (2002). Louveaux (1993) reviews models for stochastic (but not robust) facility location. See Owen and Daskin (1998) for a survey on strategic aspects of facility location, including both dynamic problems and problems under uncertainty. Brandeau and Chiu (1989), Louveaux (1993), Daskin and Owen (1999), and Current, Daskin and Schilling (2002) review both deterministic and stochastic facility location. See Birge and Louveaux (1997) for a textbook treatment of stochastic programming theory.

2 Decision Making Under Uncertainty

Many authors divide decision-making environments into three categories: certainty, risk, and uncertainty (Rosenhead, Elton and Gupta 1972). In *certainty* situations, all parameters are deterministic and known, whereas risk and uncertainty situations both involve randomness. In *risk* situations, there are uncertain parameters whose values are governed by probability distributions that are known by the decision maker. In *uncertainty* situations, parameters are uncertain, and furthermore, no information about probabilities is known. Problems in risk situations are known as *stochastic optimization problems*; a common goal is to optimize the expected value of some objective function. Problems under uncertainty are known as *robust optimization problems* and often attempt to optimize the worst-case performance of the system.

The goal of both stochastic and robust optimization is to find a solution that will perform well under any possible realization of the random parameters. The definition of “performing well” varies from application to application, and choosing an appropriate performance measure is part of the modeling process. The random parameters can be either continuous or described by discrete scenarios. If probability information is known, uncertainty is described using a (continuous or discrete) probability distribution on the parameters. If no probability information is known, continuous parameters are generally restricted to lie in some pre-specified intervals.

The scenario approach has two main drawbacks. One is that identifying scenarios

(let alone assigning probabilities to them) is a daunting and difficult task; indeed, it is the focus of a large body of stochastic programming literature. The second disadvantage is that one generally wants to identify a relatively small number of scenarios for computational reasons, but this limits the range of future states under which decisions are evaluated. But the scenario approach generally results in more tractable models, and furthermore, it has the advantage of allowing parameters to be statistically dependent, which is often not practical when parameters are described by continuous probability distributions (though there are exceptions). Dependence is often necessary to model reality, since, for example, demands are often correlated across time periods or geographical regions and costs are often correlated among suppliers.

Most of the stochastic and robust facility location problems discussed in this paper are NP-hard since they often have as special cases classical facility location problems, which are themselves NP-hard. Min-expected-cost extensions of “minisum” models like the P -median problem (PMP; Hakimi 1964) and the uncapacitated fixed-charge location problem (UFLP; Balinski 1965)—for example, those discussed in Section 3.2.1—are relatively easy to solve since they can often be treated as larger instances of deterministic problems, for which good algorithms exist despite their being NP-hard. For example, a problem with 100 nodes and 10 scenarios can be solved in approximately the time required to solve a deterministic problem with 1000 nodes (a minute or two using today’s state-of-the-art algorithms on a desktop computer). On the other hand, problems with a minimax structure like those described in Section 4.1 are more difficult to solve to optimality; today’s best algorithms are able to solve problems perhaps an order of magnitude smaller than corresponding stochastic problems in the same amount of time. This discrepancy parallels the difference in difficulty between deterministic minisum and minimax problems. For example, relatively large instances of the UFLP and PMP may be solved quickly, but the P -center problem, which has a minimax structure, is generally solved by embedding a set-covering problem (which is itself NP-hard) inside a binary search routine.

We next discuss stochastic location problems, then turn our attention to robust location problems in Section 4.

3 Stochastic Location Problems

In this section, we discuss stochastic models for facility location. Many of these models have as an objective to minimize the expected cost or maximize the expected profit of the system. Others take a probabilistic approach—for example, maximizing the probability that the solution is in some sense “good.” Some models are solved using algorithms designed specifically for the problem, while others are solved using more general stochastic programming (SP) techniques. In any stochastic programming problem, one must decide which decision variables are first stage and which are second stage; that is, which variables must be set now and which may be set after the uncertainty has been resolved. In stochastic location modeling, locations are generally first-stage decisions while assignments of customers to facilities are second-stage, i.e., recourse, decisions. (If both decisions occur in the first stage, most problems can be reduced easily to deterministic problems in which uncertain parameters are replaced by their means.)

3.1 The Hakimi Property

Several early papers on stochastic location were devoted to establishing whether the Hakimi property holds. The Hakimi property (Hakimi 1964, 1965) states that there exists an optimal solution to a network location problem in which the facilities are located on the nodes of the network, not along the edges; it holds for minimization problems like the PMP and UFLP. Mirchandani and Odoni (1979) prove that the Hakimi property holds for a PMP on a network with shortest-path travel costs in which the cost of a path may be any concave, non-decreasing function of its length. In their problem, both demands and transportation costs may be uncertain. As a result, the shortest path between two points may change depending on the scenario, as may the optimal assignments of customers to facilities. Mirchandani (1980) uses similar analysis to determine whether the Hakimi property applies to stochastic versions of the PMP and UFLP under a variety of assumptions. Louveaux and Thisse (1985) maximize expected utility of profit in a production–distribution system in which they locate a single facility and set production levels in the first stage and make distribution decisions in the second. They show that the Hakimi property applies when the firm is risk neutral (i.e., the utility function is linear) but not when it is risk averse.

3.2 Mean Outcome Models

3.2.1 Minisum Location Problems

The most common objective in stochastic programming is to optimize the mean outcome of the system; e.g., minimize expected cost or maximize expected profit. Cooper (1974) considers the Weber problem in which the locations of the demand points may be random. He assumes a bivariate normal distribution for these locations. The objective is to choose a point for the single facility location to minimize the expected demand-weighted distance to the customers. Cooper proves that the objective function is convex with respect to the location chosen and develops an iterative algorithm that solves the first-order conditions; the algorithm is shown to be globally convergent by Katz and Cooper (1974). A robust version of this problem, in which the demand points are restricted to lie in pre-specified regions, is considered by Cooper (1978) and Juel (1980, 1981); see Section 4.1.1.

Sheppard (1974) was one of the first authors to propose a scenario approach to facility location. He suggests selecting facility locations to minimize expected cost, though he does not discuss the issue at length. The first rigorous attempt to choose facility locations to minimize expected cost under scenario-based uncertainty was offered by Mirchandani and Oudjit (1980), who discuss the 2-median problem on a tree with stochastic edge lengths described by discrete scenarios. The objective is to minimize the expected demand-weighted distance. The authors first show that the stochastic 1-median problem is equivalent to the deterministic problem, which can be solved in linear time using Goldman's (1971) algorithm. They then derive analytical results for the 2-median problem indicating, for example, that the stochastic 1-median lies on the path connecting the two 2-medians if and only if the demands at the 1-median are served by the same 2-median in all scenarios. They suggest an algorithm based on partial enumeration.

Weaver and Church (1983) present a Lagrangian relaxation algorithm for the stochastic PMP on a general network discussed by Mirchandani and Odoni (1979). Their approach illustrates how some scenario-based stochastic problems can be treated simply as larger versions of the deterministic problem. We are given a set of scenarios, each of which specifies a realization of the demands and travel costs and has a fixed prob-

ability of occurrence. The objective is to minimize the expected travel cost, subject to the standard PMP constraints. Customers may be assigned to different facilities in different scenarios. (If customers must be assigned to the same facility in every scenario, the problem reduces to a deterministic PMP in which the uncertain parameters are replaced by their means.) Weaver and Church implicitly treat this problem with n customers and s scenarios as though it were a deterministic problem with ns customers. They solve the model using the standard Lagrangian relaxation method for the PMP (Cornuejols, Fisher and Nemhauser 1977), relaxing the assignment constraints. The Lagrangian subproblem decouples by facility in the same way that the subproblem for the classical PMP does.

Mirchandani, Oudjit and Wong (1985) begin with the same formulation as Weaver and Church, explicitly reformulating it as a deterministic PMP with ns customers, each corresponding to a customer-scenario pair in the original problem. Like Weaver and Church, Mirchandani et al. also suggest a Lagrangian relaxation method, but instead of relaxing the assignment constraints, they relax the single constraint requiring P facilities to be opened. The resulting subproblem is equivalent to the UFLP in which all facilities have the same fixed cost, equal to the single Lagrange multiplier. The authors solve this subproblem using Erlenkotter's (1978) DUALOC algorithm and update the multiplier using a subgradient method. Since the UFLP does not have the integrality property, Mirchandani et al.'s Lagrangian bound is at least as strong as Weaver and Church's. Of course, this comes at a cost since the subproblem is more difficult to solve, but the authors prove that the procedure is guaranteed to converge to the optimal multiplier in no more than $n - 1$ iterations. This model can be used to model multi-commodity or multi-objective problems instead of multi-scenario ones; Mirchandani et al. use the term "multi-dimensional" to describe this type of framework.

Snyder, Daskin and Teo (2003) use an approach similar to Weaver and Church's to solve a scenario-based stochastic version of the joint location-inventory model of Daskin, Coullard and Shen (2002) and Shen, Coullard and Daskin (2003). The original model chooses DC locations to minimize fixed costs, transportation costs, and endogenously computed inventory costs at the DCs in the face of stochastic demands. Snyder et al. allow the demand means and variances themselves to be stochastic, as well as costs, lead times, and other parameters. The non-linear objective function prevents the Lagrangian

subproblem from decoupling by facility as in Weaver and Church’s algorithm. Rather, it decouples by facility–scenario pair, and the scenarios are linked in a subsequent step of the algorithm. (The original location–inventory model is itself an interesting example of the way a stochastic problem can be formulated as a deterministic one whose objective function is a function of the means and variances of the random parameters.)

Louveaux (1986) presents stochastic versions of the capacitated P -median problem (CPMP) and capacitated fixed-charge location problem (CFLP) in which demands, production costs, and selling prices are random. The goal is to choose facility locations, determine their capacities, and decide which customers to serve and from which facilities to maximize the expected utility of profit. Since demands are random and facilities are capacitated, the facilities chosen in the first stage may be insufficient to serve all of the demands in the second stage; hence a penalty for unmet demand is included in the models. To formulate the stochastic CPMP, the constraint requiring P facilities to be opened is replaced by a budget constraint on the total cost; this constraint must be satisfied under any realization of the demand. The budget can be used to determine P . The author shows that under a particular type of budget constraint, the two stochastic models (CFLP and CPMP) are equivalent. Louveaux and Peeters (1992) present a dual-based heuristic for Louveaux’s CFLP model with scenario-based uncertainty, and Laporte, Louveaux and van Hamme (1994) present an optimal algorithm based on the L -shaped method of stochastic programming. Louveaux (1993) reviews applications of these and related models.

Chan, Carter and Burnes (2001) consider “stochastically processed demands”—demands that arise from completed service at a facility. In particular, each facility serves as a processing site, and when it completes processing of its current buffer of items, a vehicle delivers a load of new items for processing. The objective is to minimize the expected transportation cost. The problem is solved heuristically using stochastic decomposition (an extension of Benders decomposition) and space-filling curves.

A somewhat different approach is taken by Ricciardi, Tadei and Grosso (2002), who consider a facility location model with random throughput costs at the DCs. The objective is to minimize the deterministic transportation cost (plant–DC and DC–customer) plus the expected throughput cost at the DCs. The authors first consider the network flow aspect of the problem (assuming the DC locations are given), developing a multi-

nomial logit model for the expected flows. They then embed the expected cost model into a non-linear integer program (NLIP); this model is solved heuristically since for each candidate solution to the location problem, a Lagrangian problem must be solved to compute the expected flows.

3.2.2 Location–Transportation Problems

Balachandran and Jain (1976) present a capacitated facility location model with piecewise linear production costs that need not be either concave or convex. Demands are random and continuous, described by some joint probability distribution. There are penalty and holding costs for producing too little or too much relative to the realized demand. The objective is to minimize the expected cost of location, production, transportation, and underage and overage. The authors present a branch-and-bound algorithm in which the production functions are replaced by linear underestimates; the resulting problem is a capacitated convex transportation problem whose optimal solution can be used to obtain both lower and upper bounds on the optimal objective value for the original problem. Branching corresponds to partitioning the production cost function and estimating each portion separately.

LeBlanc (1977) considers a similar problem, but with linear production (or transportation) costs. For fixed facility locations, the problem reduces to the stochastic transportation problem (Williams 1963). LeBlanc presents a Lagrangian heuristic, while França and Luna (1982) solve the problem optimally using generalized Benders decomposition. França and Luna also permit additional linear configuration constraints, for example to limit the number of facilities to be constructed.

A simpler but more tractable model is presented by Gregg, Mulvey and Wolpert (1988), who minimize a weighted sum of the production cost (a one-time cost for establishing capacity), the transportation cost, and expected overage and underage costs. By varying the weights, the modeler can express different preferences among the objectives. The model is solved using MINOS and is illustrated by a case study involving the Queens borough public library system in New York City.

In all of these problems, only first-stage decisions are available; there are no recourse decisions. Once production and transportation levels are set, they cannot be changed

after the uncertainty is resolved. The objectives therefore include the expected recourse cost—the cost of holding and stockouts.

3.2.3 Dynamic Location Problems

Berman and Odoni (1982) study a single-facility location problem in which travel times are stochastic and the facility (e.g., an ambulance) may be relocated at a cost as conditions change. Travel times are scenario-based, and scenario transitions occur according to a discrete-time Markov process. The objective is to choose a facility location for each scenario to minimize expected transportation and relocation costs. The authors show that the Hakimi property applies to this problem and that the problem on a tree is equivalent to the deterministic problem; any scenario can be used to determine the optimal location since the 1-median on a tree is independent of the edge lengths. They then present a heuristic for the problem on a general network that involves iteratively fixing the location in all but one scenario and solving what amounts to a 1-median problem. They also discuss simple bounds on the optimal objective value of the p -facility problem. Berman and LeBlanc (1984) introduce a heuristic for this problem that loops through the scenarios, performs local exchanges within each, and then performs exchanges to link the scenarios in an effort to reduce relocation costs.

Carson and Batta (1990) present a case study of a similar problem in which a single ambulance is to be relocated on SUNY Buffalo's Amherst campus as the population moves about the campus throughout the day (from classroom buildings to dining halls to dormitories, etc.). Given the difficulties inherent in identifying probability distributions and estimating relocation costs in practice, Carson and Batta simply divide the day into four unequal time periods and solve a 1-median problem in each. Relocation costs are not explicitly considered, but the decision to use four time periods was arrived at in response to the tradeoff between frequent relocation and increased response times.

Another dynamic, stochastic facility location problem was studied by Jornsten and Bjorndal (1994), who choose where and when to locate facilities over time to minimize the expected time-discounted cost; production and distribution costs are random. Their algorithm uses scenario aggregation and an augmented Lagrangian approach. A closely related model is the capacity planning problem studied by Eppen, Martin and

Schrage (1989). Although their model is not a facility location model, it and other capacity planning models have a similar flavor to facility location models since capacity expansion entails both an up-front fixed cost and an on-going processing cost. In Eppen et al.'s multi-period model, demands and selling prices are random; capacity levels (for all time periods) are first-stage decisions and production levels are recourse decisions. Their model chooses capacity configurations at each plant in each time period, subject to a re-tooling cost for changing capacity. The objective is to maximize the time-discounted profit subject to a limit on expected downside risk (EDR). Their algorithm involves successively tightening the EDR constraint and re-solving, resulting in multiple solutions; the decision maker can choose among these solutions based on the tradeoff between expected profit and EDR. The formulation presented in the paper is solved by a general-purpose MIP solver but is very large, with as many as five subscripts on some parameters and variables, making it impractical for large problems.

An earlier stochastic capacity planning problem was considered by Manne (1961), who assumes that future demands follow a random walk with an upward trend. The objective is to minimize the expected discounted cost. Manne shows that it is sufficient to consider a deterministic problem in which the interest rate is replaced with one that depends on the demand variance. Bean, Hagle and Smith (1992) relax some of Manne's assumptions about the demand process and cost structures and show similar results using stochastic programming.

3.2.4 Competitive Location Problems

Ghosh and McLafferty (1982) introduce a model for locating multiple stores to maximize market share in a competitive environment with demand uncertainty (actually, uncertainty as to which stores a competitor plans to close, but in this setting they amount to the same thing). The authors discuss a model from the marketing literature for estimating market share given fixed store locations. The location model itself is formulated as a multi-objective model, with each objective representing the market-share-maximization objective in a given scenario. Ultimately, the objectives are combined into a weighted sum to be minimized. If the weights represent scenario probabilities, the objective is equivalent to minimizing the expected cost; otherwise, the weights can be adjusted sys-

tematically to find non-dominated solutions (solutions for which no objective can be improved without degrading another objective). For a given set of weights, the problem is solved using an exchange heuristic. On a small sample problem, 3 non-inferior solutions were found, and the authors provide some discussion as to how to choose among them.

A similar model introduced by Ghosh and Craig (1983) uses competitive equilibrium theory to estimate the expected profit to be enjoyed by a given facility location in a competitive and uncertain environment. Uncertainty comes both from changing customer demographics and from the actions of the firm's competitors. After the authors estimate the expected profit, they solve a single-facility location problem heuristically.

De Palma, Ginsburgh, Labbé and Thisse (1989) study a multi-firm competitive facility location with random consumer utilities. A consumer's utility for firm i is expressed as a constant a_i (the mean utility for the firm) minus the distance from the consumer to the firm's nearest facility minus a random error term. After choosing its maximum-utility firm, each consumer will choose the nearest facility within that firm. Firm i will open m_i facilities to maximize its expected sales (market share). The authors prove that if the m_i -median solution is unique for all i and if the consumers' tastes are sufficiently diverse, then there exists a unique location equilibrium, and in that equilibrium firm i locates its facilities at the m_i -median solution. The problem therefore reduces to solving a separate PMP for each firm.

3.2.5 Multi-Echelon Facility Location Problems

Because of the difficulty in solving stochastic facility location problems, research on more complex multi-echelon location models under uncertainty has only begun to appear in the literature in the past five years or so. Many of these models, which are often referred to as supply chain network design models, may be viewed as stochastic extensions of the seminal model by Geoffrion and Graves (1974). The objective in these papers is generally to optimize the expected outcome; to the best of our knowledge, no published research has tackled multi-echelon facility location under the more difficult robust objectives discussed in Section 4.

A qualitative discussion of global supply chain design is given by Vidal and Goetschal-

ckx (1997, 2000); the latter paper also presents a large-scale MIP for choosing plant locations and suppliers that incorporates the suppliers' reliability into the constraints.

MirHassani, Lucas, Mitra, Messina and Poojari (2000) formulate a supply chain network design problem as a stochastic program with fixed recourse; the SP has binary first-stage variables and continuous second-stage variables. The objective function coefficients are deterministic; uncertainty is present only in the right-hand sides of the recourse constraints, which may represent, for example, demands or capacities. The authors focus especially on parallel implementation issues for their proposed Benders decomposition algorithm.

Tsiakis, Shah and Pantelides (2001) consider a multi-product, multi-echelon supply chain under scenario-based demand uncertainty. The goal is to choose middle-echelon facility locations and capacities, transportation links, and flows to minimize expected cost. Transportation costs are piecewise linear concave. The model is formulated as a large-scale MIP and solved using CPLEX.

Alonso-Ayuso, Escudero, Garín, Ortuño and Pérez (2003) introduce a more general model that makes decisions about plant capacities, product mix, vendor selection, and assignments of products to plants. The problem is formulated as a multi-stage SP. First-stage variables are binary and represent capacity, product mix, and sourcing decisions. Decisions in subsequent stages are tactical (production, inventory, transportation, etc.) and are represented by continuous variables. The proposed solution algorithm is based on "branch-and-fix coordination"; it was applied with mixed success to moderate-size test problems.

While the preceding three papers consider scenario-based uncertainty, Santoso, Ahmed, Goetschalckx and Shapiro (2003) study a global supply chain network design problem with continuously random parameters, and thus an infinite number of scenarios (though their technique can also be used in problems with a large but finite number of scenarios). Costs, demands, and capacities are random. The problem is to decide where to build facilities and what machines to build at each facility to minimize the total expected cost, which includes a shortfall penalty in case the constructed capacity falls short of the realized demand. The problem is formulated as a two-stage SP with binary first-stage variables and continuous recourse variables and is solved using accelerated Benders decomposition. For each candidate location vector, the expected second-stage cost is

estimated using “sample average approximation” (SAA), which estimates the expected cost using a random sample of the uncertain parameters.

Finally, Butler, Ammons and Sokol (2003) present a model for designing the supply chain for a new product launch under uncertain costs, demands, capacities, and other parameters. Their objective function incorporates both expected cost and regret, though it can be viewed as an expected-cost-type objective with some scenarios weighted more than others based on their optimal costs. They also explore the relationship between expected cost and some of the robustness measures discussed in Section 4.

3.3 Mean–Variance Models

The mean outcome models discussed above consider only the expected performance of the system, ignoring the variability in performance and the decision maker’s possible risk aversion. However, a small body of literature has incorporated the firm’s level of risk aversion into the decision-making process, typically by using a mean–variance objective function. Jucker and Carlson (1976) use such an objective in a stochastic formulation of the UFLP in which selling price (and hence demand) may be random. They develop solution methods for four types of firms, all risk averse, characterized by which variables (e.g., price) they set and which others (e.g., demand) they accept as a result. Hodder and Jucker (1985) extend Jucker and Carlson’s model to allow for correlation among the random prices. Their model is a quadratic programming model but can be solved easily. Hanink (1984) and Hodder (1984) incorporate the capital asset pricing model (CAPM) into facility location problems and compare it to mean–variance objectives.

Hodder and Dincer (1986) consider the location of capacitated facilities globally under exchange rate uncertainty. The model incorporates the financing aspects of plant construction by endogenously deciding how much of each plant’s total cost to borrow from each country; the per-period cost of this financing is a random variable since the exchange rates are uncertain. In addition, costs and per-unit profits are uncertain. The model minimizes a mean–variance expression concerning the total profit. This objective is quadratic and involves a large variance–covariance matrix, each off-diagonal term of which requires a bilinear term in the objective function. Therefore, the authors propose an approximation scheme that effectively diagonalizes the variance–covariance matrix

so that the objective function contains only squared terms and no bilinear terms. The resulting model is solved using an off-the-shelf quadratic programming solver for small problems and using a gradient search method for larger ones. No discussion is provided concerning the form of uncertainty (discrete or continuous) or the probability distributions governing it, but in theory any approach could be used as long as the random parameters can be expressed adequately in the form needed for the approximation.

Verter and Dincer (1992) review the literature on stochastic facility location and capacity expansion problems, focusing on global manufacturing problems.

3.4 Probabilistic Approaches

In contrast to the models discussed above, which consider the expected value and/or variance of the stochastic objective function, there is a substantial body of literature that considers probabilistic information about the performance of the system—for example, maximizing the probability that the performance is good or constraining the probability that it is bad, under suitable definitions of “good” and “bad.” We discuss three such approaches: max-probability locations, chance-constrained programming, and distribution maps.

3.4.1 Max-Probability Locations

Probably the first attempt to solve stochastic location problems of any type is that of Frank (1966), who considers probabilistic centers and medians on a network with arbitrarily distributed independent random demands. He presents methods for finding “max-probability” centers (points that maximize the probability that the maximum weighted distance from the point is within a given limit) and medians (points that maximize the probability that the total demand-weighted distance from the point is within a given limit). Finding the max-probability center is equivalent to a nonlinear minimization problem, which Frank shows to be easy if the demand distribution is discrete. For medians, he approximates the demand distribution using the normal distribution; the problem is then easy. Frank proves that if the probability distributions are not known but are estimated using sampling, the resulting centers and medians are “good enough” in some sense. He later extends his analysis to jointly distributed normal demands

(Frank 1967).

Berman, Wang, Drezner and Wesolowsky (2003a) consider a problem similar to Frank's using center-type objectives instead of median-type ones, and using uniform instead of normal demands. They propose an enumerative approach. The Hakimi property does not apply, but enumeration is possible because the links can be divided into regions on which the objective function is constant; finding the optimal solution thus reduces to enumerating all $O(n^3)$ of these regions. The same problem is considered on the plane by Berman, Wang, Drezner and Wesolowsky (2003b), who prove that the objective function is convex under a certain assumption about the uniform distribution parameters and present an exact algorithm, based on dual-ascent, for this case. If the assumption does not hold, the authors propose using the dual-ascent algorithm as a heuristic and embedding it into a branch-and-bound scheme based on the triangulation method introduced by Drezner and Suzuki (2004). Problems with 100 nodes or so can be solved in negligible time using Excel's solver.

3.4.2 Chance-Constrained Programming

Chance-constrained programming involves requiring the probability of a certain constraint holding to be sufficiently high. For example, Carbone (1974) chooses P facilities under normally distributed, possibly correlated, demands to minimize an upper bound on the total demand-weighted distance that can be achieved with probability α ; that is, to minimize K such that $P\left(\sum_i \sum_j h_i d_{ij} Y_{ij} \leq K\right) \leq \alpha$, where h_i is the (random) demand of customer i , d_{ij} is the distance from facility j to customer i , Y_{ij} is 1 if customer i is assigned to facility j , and $0 \leq \alpha \leq 1$ is a constant. Carbone transforms this problem into its deterministic equivalent, which is a convex minimization problem.

Shiode and Drezner (2003) use a similar approach for a competitive facility location problem on a tree. There are two players, each of whom locates a single facility in turn. Demands are stochastic to the leader but deterministic to the follower. The objective of the paper is to characterize both players' optimal solutions. In the deterministic version, the follower's problem turns out to be a median-type problem while the leader's problem is a center-type problem. In the stochastic version, the leader's problem is to choose a location to maximize K such that the probability that the follower's market share is no

more than K is at least α . (The follower’s problem is the same as in the deterministic version.) The authors show that the Hakimi property applies to the leader’s problem and propose an efficient solution method that involves solving a 1-dimensional unconstrained convex minimization problem using a line search for each node, then choosing the best node.

3.4.3 Distribution Maps

Whereas max-probability papers attempt to choose a location that maximizes the probability of optimality, the papers discussed in this section evaluate the probability that any given solution will be optimal once the uncertainty is resolved. A complete description of the probability that any point is optimal is called a “distribution map.” Wesolowsky (1977) first used the distribution map concept in facility location, considering the Weber problem on a line with demands whose weights are multivariate normal. He shows that only the demand nodes have non-zero probability of being optimal, a result analogous to the Hakimi property. He also computes the expected value of perfect information (EVPI): the expected difference between the cost of the optimal solution and the cost of the solution found by replacing the random weights with their means.

Drezner and Wesolowsky (1981) extend Wesolowsky’s model to include general l_p distances on the plane. They present a procedure for finding approximate probabilities under the rectilinear metric (l_1 -norm), and another for finding approximate probabilities for the general l_p -norm when locations are restricted to the nodes only. Drezner and Wesolowsky (1980) compute the EVPI in the same setting. They show that the EVPI can be separated into individual one-dimensional EVPIs and derive a general expression for it. They develop exact expressions for the rectilinear case and approximate expressions for the Euclidean case. In addition, they use simulation to show that if the demands are lognormal instead of normal, the expressions are still approximately correct.

3.5 Evaluating Options

Constructing facilities in several countries gives a firm a degree of operational flexibility since it can shift production to countries with favorable exchange rates, local costs, labor

availability, etc. Most papers that discuss this strategy are concerned with evaluating the “option value” of a set of global facility locations—generally quite a complex task—rather than with choosing the optimal set. One can think of these papers as computing the objective value for a given solution to a location problem, as opposed to finding the optimal solution.

Huchzermeier and Cohen (1996) evaluate operational options over multiple time periods under uncertainty in exchange rates. They build a Markov model of exchange rates and solve a supply chain design problem for each scenario, then use stochastic dynamic programming to determine the value of each option. Kogut and Kulatilaka (1994) similarly use dynamic programming to evaluate options when there is a cost for switching production from one site to another. They discuss the threshold at which switching becomes advantageous and make the observation that the model favors countries with volatile exchange rates since they provide greater opportunity to take advantage of fluctuations. This counter-intuitive result illustrates the difference between financial and operational hedging: while financial hedging seeks to eliminate volatility, operational hedging seeks to exploit it. This issue is discussed in a more qualitative setting by Kogut (1985). Lowe, Wendell and Hu (2002) provide a decision-analysis approach for the operational hedging concept, illustrating its use with a popular Harvard Business Review case.

Gutiérrez and Kouvelis (1995) present a model to choose suppliers internationally to hedge against changes in exchange rates and local costs. The model reduces to a robust version of the UFLP with uncertain costs. Unlike the models discussed in the preceding paragraph, which are descriptive, Gutiérrez and Kouvelis’s model is normative, though necessarily less rich than the descriptive models. We discuss their model further in Section 4.2.3.

4 Robust Location Problems

When no probability information is known about the uncertain parameters, the expected cost and other objectives discussed in Section 3 are irrelevant. Many measures of robustness have been proposed for this situation. The two most common are minimax cost and minimax regret, which are closely related to one another and are discussed in

Section 4.1. Other less common measures are discussed in Section 4.2. Several of these robustness measures are discussed in the text on robust optimization by Kouvelis and Yu (1997), though they use somewhat different terminology than we do.

As in the stochastic optimization case, uncertain parameters in robust optimization problems may be modeled as either discrete or continuous. Discrete parameters are modeled using the scenario approach. Continuous parameters are generally assumed to lie in some pre-specified interval, since it is often impossible to consider a “worst case scenario” when parameter values are unbounded. We will describe this type of uncertainty as “interval uncertainty” and describe parameters modeled this way as “interval-uncertain” parameters. Some authors use the term “scenario” in the continuous case to refer to a particular realization of the uncertain parameters.

For a given problem under uncertainty with no probability information, the *minimax cost solution* is the one that minimizes the maximum cost across all scenarios. This measure is, on the one hand, overly conservative, emphasizing the worst possible scenario, and on the other hand, somewhat reckless, since it may produce quite poor solutions for scenarios other than the one with maximum cost, especially if the scenarios have a form like “small demand / moderate demand / large demand.” Minimax cost may be an appropriate measure for situations in which it is critical for the system to function well even in the worst case—for example, the location of hospitals or fire stations—or in which a firm’s competitors are likely to make decisions that make the worst scenario come to pass for the firm.

The other two most common robustness measures consider the *regret* of a solution, which is the difference (absolute or percentage) between the cost of a solution in a given scenario and the cost of the optimal solution for that scenario. Regret is sometimes described as opportunity loss: the difference between the quality of a given strategy and the quality of the strategy that would have been chosen had one known what the future held. Models that seek to minimize the maximum (absolute or relative) regret across all scenarios are called *minimax (absolute or relative) regret* models. Minimax cost problems can often be transformed into equivalent minimax regret problems, and vice-versa, since the cost and regret of a given scenario differ only by a constant. Solution approaches for one criterion are often applicable to the other, as well.

The regret criterion is usually applied in uncertainty situations, in which no proba-

bility information is known. It has been discussed in the context of risk situations as well, but minimizing expected regret is equivalent to minimizing expected cost. To see this, consider a general min-expected-absolute-regret problem with variables x_1, \dots, x_n , feasible set X , scenarios $s \in S$, objective function coefficients c_{is} , scenario probabilities q_s , and optimal scenario objective values z_s^* :

$$\text{minimize} \quad \sum_{s \in S} q_s R_s \quad (1)$$

$$\text{subject to} \quad R_s = \sum_{i=1}^n c_{is} x_i - z_s^* \quad \forall s \in S \quad (2)$$

$$x \in X \quad (3)$$

Substituting the regret variables R_s into the objective function, we get

$$\text{minimize} \quad \sum_{s \in S} q_s \left(\sum_{i=1}^n c_{is} x_i - z_s^* \right) \quad (4)$$

$$\text{subject to} \quad x \in X \quad (5)$$

The objective function of this revised problem is the min-expected-cost objective function minus a constant. The relative regret case is similar. This equivalence is sometimes overlooked in the literature.

4.1 Minimax Models

Minimax regret models are commonly employed in the literature. Generally such problems are solved using problem-specific algorithms, though general-purpose algorithms for minimax regret linear programs with interval-uncertain objective function coefficients were proposed by Mausser and Laguna (1998) for problems with absolute regret and by Mausser and Laguna (1999b) for problems with relative regret. The algorithms rely on the fact that for a given solution, each uncertain parameter is set either to its lower or its upper endpoint in the regret-maximizing scenario. To identify this scenario, the authors solve a MIP that adds one binary variable and a few constraints to the original model for each uncertain parameter. This approach is practical for small- to moderate-size LPs. Mausser and Laguna (1999a) propose a greedy heuristic for the absolute regret problem that contains some methods for diversification to avoid local optima. Its results can be

used on their own or in place of the exact solution to the MIP formulated by Mausser and Laguna (1998).

The general strategy of the algorithm in many minimax regret papers can be described as follows:

1. Choose a candidate solution x .
2. Determine the maximum regret across all scenarios if solution x is chosen. For discrete scenarios, this is easy: just compute the cost of the solution under each scenario and compare it to the optimal cost for the scenario, then choose the scenario with the greatest regret. For interval uncertainty, techniques for finding the regret-maximizing scenario rely on the fact that this scenario typically has each parameter set to an endpoint of its interval. Still, the problem of identifying this scenario can be quite difficult. Solving this problem is the crux of the algorithms by Mausser and Laguna (1998, 1999a, 1999b) discussed above. On the other hand, Averbakh and Berman (2000b) develop an $O(n^2)$ algorithm to determine the regret-maximizing scenario for their problem, the 1-median on a general network.
3. Either repeat steps 1 and 2 for all possible solutions (as in Averbakh and Berman 2000b), or somehow find a new candidate solution whose regret is smaller than the regret determined in step 2 (as in Mausser and Laguna 1998, 1999b).

Regret-based problems tend to be more difficult than stochastic problems because of their minimax structure. In fact, many deterministic problems that are polynomially solvable have robust versions that are not. For example, the economic order quantity (EOQ) model is still easy in its robust form (Yu 1997), but the minimax regret shortest path problem is NP-hard (Yu and Yang 1998). As a result, many papers restrict their attention to problems with some special structure, such as 1-median problems or P -medians on tree networks. The focus in such papers is generally to develop analytical results or polynomial-time algorithms. Kouvelis and Yu (1997) provide detailed analysis of many of these problems. We review papers on specially structured problems in Section 4.1.1. Then, in Section 4.1.2, we examine more general problems, for which heuristics are generally required.

4.1.1 Specially Structured Problems

Chen and Lin (1998) consider the minimax-regret 1-median problem on a tree with interval-uncertain edge lengths and node weights. As with many minimax problems, the Hakimi property does not apply to this problem. Chen and Lin present an $O(n^3)$ algorithm. They require node weights to be positive, but Burkard and Dollani (2001) present a polynomial time algorithm for the case in which node weights may be positive or negative. Vairaktarakis and Kouvelis (1999) similarly consider 1-medians on a tree, but in their problem, edge lengths and node weights may be linear over time (i.e., not stochastic but deterministic and dynamic) or random and scenario-based. They trace the path of the solution over time (in the dynamic case) and present low-order polynomial algorithms for both cases.

The problem on a general network is significantly harder. In fact, Averbakh (2003a) proves that if edge weights are interval-uncertain, then both the 1-median and weighted 1-center problems on a general network are NP-hard. (The (deterministic) weighted P -center problem is to locate P facilities to minimize the maximum weighted distance traveled by any customer to its nearest facility.) His proof involves reducing COVER to the problems under consideration. Both problems are polynomially solvable if only the node weights are uncertain (Averbakh and Berman 1997a, 2000b).

Averbakh and Berman (2000b) consider the minimax-regret 1-median problem on a general network with interval-uncertain node weights. They present the first polynomial-time algorithm for this problem, with complexity $O(mn^2 \log n)$, where m is the number of edges and n is the number of nodes. For the nodal problem, in which facilities may only be located on the nodes, they present an $O(n^3)$ algorithm. They also consider the problem on a tree (with node or edge locations), reducing the known complexity of that problem to $O(n^2)$. The same authors later reduced the complexity of the tree problem further to $O(n \log^2 n)$ (Averbakh and Berman 2003).

Minimax-regret center problems are even more difficult than median problems since center problems already have a minimax structure in their deterministic forms. Averbakh and Berman (2000a) present an $O(n^6)$ algorithm for the minimax regret 1-center problem on a tree with interval-uncertain node weights and edge lengths. Part of the reason for the high order of complexity even in this restrictive setting is that determining

the maximum regret for a given solution is non-trivial. They present an $O(n^2 \log n)$ algorithm for the problem with unweighted nodes. The complexities were reduced to $O(n^3 \log n)$ and $O(n \log n)$ for the weighted and unweighted cases, respectively, by Burkard and Dollani (2002), whose algorithms are based on computational geometry and make use of Averbakh and Berman's (2000a) result that an optimal edge (as opposed to an optimal location on it) can be found in $O(n^2 \log n)$ time.

Labbé, Thisse and Wendell (1991) consider the 1-median problem on a network with uncertain node weights. We are given an estimate for each weight, as well as a range on the possible perturbation for each weight; in other words, node weights are interval-uncertain, but we also have an estimate of them, possibly the midpoint of the interval. Given an optimal solution x^* for the deterministic problem, the authors define the “degree of optimality for tolerance τ ,” denoted $\alpha^*(\tau)$, as the maximum possible difference between the cost of x^* under a perturbed objective and the optimal objective value under the perturbed objective, for all perturbations of size no more than τ for each uncertain weight. In other words, $\alpha^*(\tau)$ is the maximum regret over all “scenarios” in which the weights differ from their estimates by no more than τ . Labbé et al. are not interested in finding an x to minimize this regret measure, but rather in performing sensitivity analysis by characterizing the tradeoff between the uncertainty in the weights (τ) and the worst-case performance of the optimal solution ($\alpha^*(\tau)$). This analysis can be performed for any solution, not just x^* .

Carrizosa and Nickel (2003) consider a measure closely related to that of Labbé et al. Rather than defining a solution's robustness based on how much the cost changes for a given magnitude of change in the parameters, Carrizosa and Nickel define robustness as the minimum magnitude of the parameter perturbation required to violate some desired limit on the cost. If the parameters may change substantially from their estimates without violating the cost cap, the solution is very robust. They consider this measure in the context of the Weber problem with uncertain demands, proposing an iterative solution method that solves, at each iteration, a non-differentiable concave maximization problem.

4.1.2 General Problems

We now turn our attention to multiple-facility problems on general networks under minimax objectives. Naturally, these problems are significantly harder than those discussed in Section 4.1.1, and as a result, they are usually solved heuristically. To our knowledge, the only papers to present analytical results for such problems are those of Averbakh and Berman (1997b) and Averbakh (2003b). Averbakh and Berman (1997b) consider the minimax-regret weighted P -center problem on a general network with interval-uncertain demands. They show that the minimax regret problem can be solved by solving $n + 1$ deterministic weighted P -center problems: n of them on the original network and 1 on an augmented network, where n is the number of nodes in the problem. Since the weighted P -center problem can be solved in polynomial time for the special cases in which $P = 1$ or the network is a tree, this leads to a polynomial-time algorithm for the minimax problem in these special cases.

Averbakh (2003b) extends this concept to more general robust (minimax) combinatorial optimization problems whose deterministic forms already have minimax objectives, including certain facility location, scheduling, and other combinatorial problems. Robustness may be defined as minimax cost or minimax absolute or relative regret, and uncertainty is described by intervals. As in Averbakh and Berman (1997b), the main idea is to reduce the problem to a series of deterministic minimax problems. He proves three main results: (1) the minimax cost problem can be solved by setting all uncertain parameters to their upper bounds and solving the resulting deterministic problem; (2) the minimax absolute and relative regret problems can be solved by solving m deterministic problems, each of which involves setting one parameter to its upper bound and the others to their lower bounds, plus one more deterministic problem (m is the number of uncertain parameters); (3) a polynomial-time algorithm for the deterministic problem implies a polynomial-time algorithm for the minimax cost and minimax relative regret problems, but not necessarily for the minimax absolute regret problem. However, since absolute regret problems can often be transformed into relative regret problems, a polynomial-time algorithm for the relative regret problem will often imply one for the absolute regret problem.

More commonly, heuristic approaches are used to solve these general problems. Serra,

Ratick and ReVelle (1996) solve the maximum capture problem (to locate P facilities in order to capture the maximum market share, given that the firm’s competitors have already located their facilities) under scenario-based demand uncertainty. They consider both maximizing the minimum market share captured (the maximization analog of the “minimax cost” criterion) and minimizing maximum regret. They present a heuristic that involves solving the deterministic problem for each scenario, choosing an initial solution based on those results, and then using an exchange heuristic to improve the solution. A similar approach is used by Serra and Marianov (1998), who solve the minimax cost and minimax regret problems for the PMP, also under scenario-based demand uncertainty. They present a case study involving locating fire stations in Barcelona. Heuristics for minimax-regret versions of the PMP and UFLP are also discussed by Snyder and Daskin (2003); this paper will be discussed further in Section 4.2.3. Current, Ratick and ReVelle (1997) present a model in which facilities are located over time, but the number of facilities that will ultimately be located is uncertain. They call their model NOFUN (“number of facilities uncertain”). The approach to uncertainty is scenario based (scenarios dictate the number of facilities to open), and the authors discuss minimizing either expected or maximum regret. Their proposed formulation is based on the PMP and is solved using a general-purpose MIP solver.

4.2 Other Robustness Measures

We now discuss several other robustness measures in roughly chronological order. Most of these measures have been applied to facility location problems; for those that have not, we provide examples from the broader logistics literature.

4.2.1 Robustness and Stability

One of the earliest robustness measures was proposed by Gupta and Rosenhead (1968) and Rosenhead et al. (1972). In these papers, decisions are made over time, and a solution is considered more robust if it precludes fewer good outcomes for the future. An example in the latter paper concerns a facility location problem in which a firm wants to locate five facilities over time. Suppose all possible five-facility solutions have been enumerated, and N of them have cost less than or equal to some prespecified

value. If facility j is included in p of the N solutions, then its “robustness” is p/N . One should construct the more robust facilities first, then make decisions about future facilities as time elapses and information about uncertain parameters becomes known. Now suppose that the first facility has been constructed and the firm decides (because of budget, politics, shrinking demand, etc.) not to build any of the other facilities. The “stability” of a facility is concerned with how well the facility performs if it is the only one operating. Stability should be used to distinguish among facilities that are nearly equally robust. Note that these definitions of robustness and stability refer to individual facilities, not to solutions as a whole. An obvious disadvantage of this robustness measure is that computing it requires enumerating all possible solutions, which is generally not practical.

Schilling (1982) presents two location models that use this robustness measure (though not explicitly), both of which use stochastic, scenario-based demands. The first model is a set-covering-type model that maximizes the number of facilities in common across scenarios subject to all demands being covered in all scenarios and a fixed number of facilities being located in each scenario. By varying this last parameter, one can obtain a tradeoff curve between the total number of facilities constructed and the number of facilities that are common across scenarios. If the firm is willing to build a few extra facilities, it may be able to substantially delay the time until a single solution must be chosen, since the common facilities can be built first. The second model is a max-covering-type model that maximizes the coverage in each scenario subject to the number of common facilities exceeding some threshold. In this case the tradeoff curve represents the balance between demand coverage and common facilities. Unfortunately, Schilling’s models were shown by Daskin, Hopp and Medina (1992) to produce the worst possible results in some cases. For example, imagine a firm that wants to locate two DCs to serve its three customers in New York, Boston, and Seattle. New York has either 45% or 35% of the demand and Boston has 35% or 45% of the demand, depending on the scenario. The remaining 20% of the demand is in Seattle, in either scenario. If the transportation costs are sufficiently large, the optimal solution in scenario 1 is to locate in New York and Seattle, while the optimal solution in scenario 2 is to locate in Boston and Seattle. Schilling’s method would instruct the firm to build a DC in Seattle first, since that location is common to both solutions, then wait until some of the uncertainty

is resolved before choosing the second site. But then all of the east-coast demand is served from Seattle for a time, a sub-optimal result.

Rosenblatt and Lee (1987) use a similar robustness measure to solve a facility layout problem. Unlike Rosenhead et al.'s measure, which considers the percentage of good solutions that contain a given *element* (e.g., facility), Rosenblatt and Lee consider the percentage of scenarios for which a given *solution* is “good,” i.e., has regret bounded by some pre-specified limit. Like the previous measure, Rosenblatt and Lee’s measure requires enumerating all solutions and evaluating each solution under every scenario, making this approach practical only for very small problems.

4.2.2 Sensitivity Analysis

Sensitivity analysis does not provide an optimization criterion but is used in an evaluative context. For example, Hodgson (1991) uses simulation to estimate the relative regret of the deterministic P -median solution when the demands and distances are randomly perturbed by scaling each data element by a normally distributed percentage. He finds that the optimal P -median solution is relatively insensitive to errors in the distances and especially to errors in the demands. Labbé et al. (1991) provide more theoretical sensitivity analysis for the 1-median problem with random demands. The goal is to derive a tradeoff between the deterministic solution’s maximum regret and the size of the perturbations. This paper is discussed further in Section 4.1.1.

Rather than evaluating the sensitivity of the deterministic solution to changes in the data, Cooper (1978) considers the sensitivity of the problem itself by evaluating the minimum and maximum possible cost of the optimal solution over all possible scenarios, assuming implicitly that a solution can be chosen after the uncertainty has been resolved. Like Cooper (1974) (discussed in Section 3.2.1), Cooper (1978) considers the Weber problem when the locations of the demand points are not known. The difference is that while Cooper (1974) assumes a probability distribution on the demand points, Cooper (1978) assumes only that the points lie within certain “uncertainty circles.” He argues that the minimum [maximum] cost can be found by solving the deterministic problem and subtracting [adding] the sum of the circle radii from the optimal cost. However, Juel (1980) points out an error in Cooper’s proof and Juel (1981) shows that Cooper’s result

holds for the minimum [maximum] cost problem only if the optimal solution does not lie within [at the center of] any of the uncertainty circles. Drezner (1989) provides an asymptotic analysis for the Weber problem on a sphere with random customer locations; his main result is that as the number of demand points approaches infinity, the relative difference between the minimum and maximum costs approaches zero.

4.2.3 p -Robustness

Rather than minimizing regret, several papers have placed constraints on the maximum regret that may be attained by the solution. This idea was first used by Kouvelis, Kurawarwala and Gutiérrez (1992), who impose a constraint dictating that the relative regret in any scenario must be no greater than p , where $p \geq 0$ is an external parameter. In other words, the cost under each scenario must be within $100(1+p)\%$ of the optimal cost for that scenario. Snyder and Daskin (2003) refer to this measure as “ p -robustness.” For small p , there may be no p -robust solutions for a given problem. Thus, p -robustness adds a feasibility issue not present in most other robustness measures.

The problem considered by Kouvelis et al. is a facility layout problem in which the goal is to construct a list of p -robust solutions, if any exist. The approach used is heuristic in the sense that, although it optimally solves the layout problem for each scenario, there is no guarantee that the resulting list of p -robust solutions is exhaustive. The p -robustness criterion is also used by Gutiérrez and Kouvelis (1995) in the context of an international sourcing problem. They present an algorithm that, for a given p and N , returns either all p -robust solutions (if there are fewer than N of them) or the N solutions with smallest maximum regret. The sourcing problem involves choosing suppliers worldwide so as to hedge against changes in exchange rates and local prices. The problem reduces to the UFLP, so the authors are essentially solving a p -robust version of the UFLP. Their algorithm maintains separate branch-and-bound trees for each scenario, and all trees are explored and fathomed simultaneously. Unfortunately, their algorithm contains an error that makes it return incomplete, and in some cases incorrect, results (Snyder 2003). Gutiérrez, Kouvelis and Kurawarwala (1996) use Benders decomposition to search for a p -robust solution to the uncapacitated network design problem. For each scenario, the authors solve a separate network design problem; these problems are linked

by feasibility cuts that are added simultaneously to all problems.

The three models discussed in the preceding paragraph attempt to find p -robust solutions but do not provide a way to differentiate among those found, which may be numerous. Snyder and Daskin (2003) address this issue by combining the min-expected-cost and p -robustness measures for the PMP and UFLP: the goal is to find the minimum-expected-cost solution that is p -robust. They solve their models using variable splitting (or Lagrangian decomposition). Their method can be used as a heuristic for the minimax-regret PMP or UFLP by systematically varying p until one obtains the smallest p for which the problem is feasible. However, finding a feasible solution, or even determining whether the problem is feasible, can be difficult, especially for small p . Therefore, the authors discuss a mechanism for detecting infeasibility that involves testing the Lagrangian lower bound against an *a priori* upper bound.

4.2.4 Model and Solution Robustness

Mulvey, Vanderbei and Zenios (1995) introduce a framework for robust optimization (the “RO framework”) that involves two types of robustness: “solution robustness” (the solution is nearly optimal in all scenarios) and “model robustness” (the solution is nearly feasible in all scenarios). The definition of “nearly” is left up to the modeler; their objective function has very general penalty functions for both model and solution robustness, weighted by a parameter intended to capture the modeler’s preference between the two. The solution robustness penalty might be the expected cost, maximum regret, or von Neumann–Morgenstern utility function. The model robustness penalty might be the sum of the squared violations of the constraints. Uncertainty may be represented by scenarios or intervals, with or without probability distributions. The authors discuss a number of applications in which the RO framework has been applied. In one example, a power company wants to choose the capacities of its plants to minimize cost while meeting customer demand and satisfying certain physical constraints. In the RO model for this problem, the objective function has the form

$$\text{minimize } E[\text{cost}] + \lambda \text{Var}[\text{cost}] + \omega [\text{sum of squares of infeasibilities}].$$

The first two terms represent solution robustness, capturing the firm’s desire for low costs and its degree of risk aversion, while the third term represents model robustness,

penalizing solutions that fail to meet demand in a scenario or violate other physical constraints like capacity. The RO framework has explicitly been employed in applications as varied as parallel machine scheduling with stochastic interruptions (Laguna, Lino, Pérez, Quintanilla and Valls 2000), relocation of animal species under uncertainty in population growth and future funding (Haight, Ralls and Starfield 2000), production planning (Trafalis, Mishina and Foote 1999), large-scale logistics systems (Yu and Li 2000), and chemical engineering (Darlington, Pantelides, Rustem and Tanyi 2000).

Killmer, Anandalingam and Malcolm (2001) use the RO framework to find solution- and model-robust solutions to a stochastic noxious facility location problem. (Though the authors discuss their model solely in the context of noxious facility location, it is similar to the UFLP and so could be applied to much more general problems.) The RO model for this problem minimizes the expected cost plus penalties for regret, unmet demand, and unused capacity. The expected cost and regret penalty are the solution robustness terms (encouraging solutions to be close to optimal), while the demand and capacity variation penalties are model robustness terms (encouraging solutions to be close to feasible). The non-linear programming model is applied to a small case study in Albany, NY and is solved using MINOS.

4.2.5 Restricting Outcomes

One use of the model robustness term in the RO framework is to penalize solutions for being too different across scenarios (in terms of variables, not costs), thus encouraging the resulting solution to be insensitive to uncertainties in the data. Vladimirou and Zenios (1997) formulate several models for solving this particular realization of the RO framework, which they call “restricted recourse.” Restricted recourse might be appropriate, for example, in a production planning context in which re-tooling is costly. However, there may be a substantial tradeoff between robustness (in this sense) and cost. The authors present three procedures for solving such problems, each of which begins by forcing all second-stage decisions to be equal, and then gradually loosening that requirement until a feasible solution is found. The stochastic programming problems are solved using standard integer SP algorithms. The authors analyze the tradeoff between robustness and cost, and often find large increases in cost as the restricted

recourse constraint is tightened.

In contrast, Paraskevopoulos, Karakitsos and Rustem (1991) present a model for robust capacity planning in which they restrict the sensitivity of the *objective function* (rather than the variables) to changes in the data. Instead of minimizing expected cost, they minimize expected cost plus a penalty on the objective's sensitivity to changes in demand. The penalty is weighted based on the decision maker's level of risk aversion. The advantage of this robustness measure is that the resulting problem looks like the deterministic problem with the uncertain parameters replaced by their means and with an extra penalty term added to the objective. Scenarios and probability distributions do not enter the mix. The down-side is that computing the penalty requires differentiating the cost with respect to the error in the data. For realistic capacity-planning problems, even computing the expected cost (let alone its derivative) is difficult and in some cases must be done using Monte Carlo simulation. For linear models, including most location models, computing the expected cost is easy, but the penalty becomes a constant and the problem reduces to the deterministic problem in which uncertain parameters are replaced by their means; this generally yields poor results. Therefore, Paraskevopoulos et al.'s robustness measure has not been applied to location problems. On the other hand, the requirement that solutions have similar costs across scenarios has some resemblance to the notion of p -robustness, discussed in Section 4.2.3, which has been applied to location problems.

4.2.6 Limiting the Scenario Space

Daskin, Hesse and ReVelle (1997) and Owen (1999) introduce the notion of " α -reliable minimax regret," the idea behind which is that the traditional minimax regret criterion tends to focus on a few scenarios that may be catastrophic but are unlikely to occur. In the α -reliable framework, the maximum regret is computed only over a subset of scenarios, called the "reliability set," whose total probability is at least α . Therefore, the probability that a scenario that was not included in the objective function comes to pass is no more than $1 - \alpha$. The parameter α is specified by the modeler but the reliability set is chosen endogenously. The authors apply the α -reliable concept to the minimax-regret P -median problem, but it could be applied to other problems, and even to other

robustness measures. For instance, one could just as easily formulate, say, the α -reliable minimax cost UFLP. Daskin et al. solve the problem using standard LP/branch-and-bound techniques, while Owen develops a genetic algorithm for it. Daskin et al. also discuss the tradeoff between reliability (large α) and regret (large objective function).

The essence of the α -reliability approach is that the robustness of a solution is evaluated only over a subset of the scenario space, i.e., the reliability set. In that sense, α -reliability is similar to two recently developed approaches to general robust optimization: the “ellipsoidal” model formulated by Ben-Tal and Nemirovski (2000, 2002), in which robustness is evaluated only on an ellipsoid in the scenario space, and the approach advocated by Bertsimas and Sim (2003, 2004), in which robustness is evaluated only over those scenarios in which at most a fixed number of the uncertain parameters are allowed to deviate from their estimates. In the Ben-Tal and Nemirovski approach, the ellipsoid (the analogue to the reliability set) is chosen exogenously, while in both the Bertsimas and Sim approach and the α -reliability approach; the set is chosen endogenously and the solution must be robust over all possible realizations of that set.

The α -reliable approach is extended by Chen, Daskin, Shen and Uryasev (2003), who introduce a measure called “ α -reliable mean-excess regret.” The objective function consists of a weighted sum of the maximum regret over the reliability set (as in α -reliable minimax regret) and the conditional expectation of the regret over the scenarios excluded from the reliability set. Unlike the α -reliable minimax regret model, the mean-excess model ensures that solutions perform reasonably well even in the “ignored” scenarios. The mean-excess model is also significantly more tractable; in fact, Chen et al. suggest a heuristic for solving the α -reliable minimax regret model by solving a series of mean-excess models.

4.2.7 A Game-Theoretic Approach

Blanchini, Rinaldi and Ukovich (1997) consider capacitated network flow and network design problems under uncertainty. Demands are uniformly distributed. They address two questions: (1) does there exist a flow assignment strategy that is guaranteed to be feasible, regardless of the realization of the demands? (2) what is the minimum-cost network design that allows a guaranteed-feasible flow assignment strategy? The problem

is modeled as a two-person game in which each player's moves correspond to setting flows in the network. The first player is the network manager, who must obey the system constraints. The second player represents the demand; this player acts malevolently to push the demands so that the constraints are violated. Finding a winning strategy reduces to solving a flow problem on a line.

5 Discussion

The literature on facility location under uncertainty has been growing steadily. Roughly half of the papers cited in this article were published in the past ten years. The growing interest in these problems is due to the increased recognition of the uncertainties faced by most firms, as well as to improvements in both optimization technology and raw computing power.

When we think of facility location under uncertainty, many of us think only of the two most common objectives: minimizing expected cost (for stochastic problems) and minimizing maximum regret (for robust problems). Yet a wide variety of other approaches has been proposed; this paper discusses at least a dozen such measures. Many of these approaches have modeling, analytical, and computational advantages over the traditional objectives. We have explored these alternative measures with the intention of providing a foundation for researchers doing work in this and related fields. To that end, we have identified four research avenues that we believe are both important and within the grasp of today's OR technology.

1. **Exact algorithms for minimax problems.** Exact algorithms for problems with minimax objectives are still in their infancy. Very few exact algorithms are available for general problems on general networks; most apply only to specially structured problems with limited potential for direct application. Because of the popularity of these models, exact algorithms for solving them would be welcomed by researchers and practitioners.
2. **Meta-heuristics for general problems.** Meta-heuristics have been applied successfully to deterministic location problems, but few, if any, have been developed for their stochastic and robust counterparts. Given the array of approaches dis-

cussed in this paper, meta-heuristics would be particularly valuable since they allow changes to the objective function and constraints with relative ease. Modelers would be able to choose from solutions generated using a number of objectives without requiring special-purpose algorithms for each.

3. **Multi-echelon models.** Only recently have researchers begun to study multi-echelon location or supply chain network design problems under uncertainty. The initial attempts have been promising, but there are many avenues for further research. For example, these problems have been solved using stochastic objectives but not robust ones. Furthermore, there is a need for models that capture the costs of tactical and/or operational functions of the supply chain under uncertainty, including inventory, transportation, and scheduling.
4. **Stochastic programming technology.** The optimization technology developed by stochastic programming researchers has become extremely powerful, but it has only begun to be used in models for facility location and other logistics problems. There is great potential for solving complex, realistic problems by leveraging the available and emerging SP technology.

Undoubtedly, there are many other research avenues that will prove to be productive; we hope that this survey paper helps to facilitate future research in this area.

References

- Alonso-Ayuso, A., Escudero, L. F., Garín, A., Ortuño, M. T. and Pérez, G. (2003), An approach for strategic supply chain planning under uncertainty based on stochastic 0-1 programming, *Journal of Global Optimization* **26**, 97–124.
- Averbakh, I. (2003a), Complexity of robust single facility location problems on networks with uncertain edge lengths, *Discrete Applied Mathematics* **127**(3), 505–522.
- Averbakh, I. (2003b), Minmax regret solutions for minimax optimization problems with uncertainty, *Operations Research Letters* **27**(2), 57–65.
- Averbakh, I. and Berman, O. (1997a), Algorithms for robust center problems, Technical report, presented at INFORMS, October 1997, Dallas, TX.

- Averbakh, I. and Berman, O. (1997b), Minimax regret p -center location on a network with demand uncertainty, *Location Science* **5**(4), 247–254.
- Averbakh, I. and Berman, O. (2000a), Algorithms for the robust 1-center problem on a tree, *European Journal of Operational Research* **123**(2), 292–302.
- Averbakh, I. and Berman, O. (2000b), Minimax regret median location on a network under uncertainty, *INFORMS Journal on Computing* **12**(2), 104–110.
- Averbakh, I. and Berman, O. (2003), An improved algorithm for the minmax regret median problem on a tree, *Networks* **41**(2), 97–103.
- Balachandran, V. and Jain, S. (1976), Optimal facility location under random demand with general cost structure, *Naval Research Logistics Quarterly* **23**, 421–436.
- Balinski, M. L. (1965), Integer programming: methods, uses, computation, *Management Science* **12**(3), 253–313.
- Ball, M. O. and Lin, F. L. (1993), A reliability model applied to emergency service vehicle location, *Operations Research* **41**(1), 18–36.
- Bean, J. C., Higle, J. L. and Smith, R. L. (1992), Capacity expansion under uncertain demands, *Operations Research* **40**(2 supp.), S210–S216.
- Ben-Tal, A. and Nemirovski, A. (2000), Robust solutions of linear programming problems contaminated with uncertain data, *Mathematical Programming Series A* **88**(2), 411–424.
- Ben-Tal, A. and Nemirovski, A. (2002), Robust optimization—methodology and applications, *Mathematical Programming Series B* **92**, 453–480.
- Berman, O. and Krass, D. (2002), Facility location problems with stochastic demands and congestion, in Z. Drezner and H. W. Hamacher, eds, ‘Facility Location: Applications and Theory’, Springer-Verlag, New York, chapter 11.
- Berman, O., Krass, D. and Menezes, M. B. C. (2004), Reliability issues, strategic co-location and centralization in m -median problems, working paper [ok to cite?], Rotman School of Management, University of Toronto.
- Berman, O., Larson, R. C. and Chiu, S. S. (1985), Optimal server location on a network operating as an $M/G/1$ queue, *Operations Research* **33**(4), 746–771.

- Berman, O. and LeBlanc, B. (1984), Location-relocation of mobile facilities on a stochastic network, *Transportation Science* **18**(4), 315–330.
- Berman, O. and Odoni, A. R. (1982), Locating mobile servers on a network with Markovian properties, *Networks* **12**, 73–86.
- Berman, O., Wang, J., Drezner, Z. and Wesolowsky, G. O. (2003*a*), The minimax and maximin location problems on a network with uniform distributed weights, *IIE Transactions* **35**(11), 1017–1025.
- Berman, O., Wang, J., Drezner, Z. and Wesolowsky, G. O. (2003*b*), A probabilistic minimax location problem on the plane, *Annals of Operations Research* **122**, 59–70.
- Bertsimas, D. and Sim, M. (2003), Robust discrete optimization and network flows, *Mathematical Programming Series B* **98**, 49–71.
- Bertsimas, D. and Sim, M. (2004), The price of robustness, *Operations Research* **52**(1), 35–53.
- Birge, J. R. and Louveaux, F. (1997), *Introduction to Stochastic Programming*, Springer-Verlag, New York.
- Blanchini, F., Rinaldi, F. and Ukovich, W. (1997), A network design problem for a distribution system with uncertain demands, *SIAM Journal on Optimization* **7**(2), 560–578.
- Brandeau, M. L. and Chiu, S. S. (1989), An overview of representative problems in location research, *Management Science* **35**(6), 645–674.
- Bundschuh, M., Klabjan, D. and Thurston, D. L. (2003), Modeling robust and reliable supply chains, working paper, University of Illinois at Urbana-Champaign.
- Burkard, R. E. and Dollani, H. (2001), Robust location problems with pos/neg weights on a tree, *Networks* **38**(2), 102–113.
- Burkard, R. E. and Dollani, H. (2002), A note on the robust 1-center problem on trees, *Annals of Operations Research* **110**(1-4), 69–82.
- Butler, R. J., Ammons, J. C. and Sokol, J. (2003), A robust optimization model for strategic production and distribution planning for a new product, working paper, University of Central Florida.

- Carbone, R. (1974), Public facilities location under stochastic demand, *INFOR* **12**(3), 261–270.
- Carrizosa, E. and Nickel, S. (2003), Robust facility location, *Mathematical Methods in Operations Research* **58**, 331–349.
- Carson, Y. M. and Batta, R. (1990), Locating an ambulance on the Amherst campus of the State University of New York at Buffalo, *Interfaces* **20**(5), 43–49.
- Chan, Y., Carter, W. B. and Burnes, M. D. (2001), A multiple-depot, multiple-vehicle, location-routing problem with stochastically processed demands, *Computers & Operations Research* **28**(8), 803–826.
- Chen, B. and Lin, C.-S. (1998), Minmax-regret robust 1-median location on a tree, *Networks* **31**(2), 93–103.
- Chen, G., Daskin, M. S., Shen, Z.-J. and Uryasev, S. (2003), The α -reliable mean-excess regret model for stochastic facility location, working paper, University of Florida, Department of Industrial and Systems Engineering.
- Cooper, L. (1974), A random locational equilibrium problem, *Journal of Regional Science* **14**(1), 47–54.
- Cooper, L. (1978), Bounds on the Weber problem solution under conditions of uncertainty, *Journal of Regional Science* **18**(1), 87–93.
- Cornuejols, G., Fisher, M. L. and Nemhauser, G. L. (1977), Location of bank accounts to optimize float: An analytic study of exact and approximate algorithms, *Management Science* **23**(8), 789–810.
- Current, J., Daskin, M. S. and Schilling, D. (2002), Discrete network location models, in Z. Drezner and H. W. Hamacher, eds, ‘Facility Location: Applications and Theory’, Springer-Verlag, New York, chapter 3.
- Current, J., Ratick, S. and ReVelle, C. (1997), Dynamic facility location when the total number of facilities is uncertain: A decision analysis approach, *European Journal of Operational Research* **110**(3), 597–609.
- Darlington, J., Pantelides, C. C., Rustem, B. and Tanyi, B. A. (2000), Decreasing the sensitivity of open-loop optimal solutions in decision making under uncertainty, *European Journal of Operational Research* **121**, 343–362.

- Daskin, M. S. (1982), Application of an expected covering model to emergency medical service system design, *Decision Sciences* **13**, 416–439.
- Daskin, M. S. (1983), A maximum expected covering location model: Formulation, properties and heuristic solution, *Transportation Science* **17**(1), 48–70.
- Daskin, M. S. (1995), *Network and Discrete Location: Models, Algorithms, and Applications*, Wiley, New York.
- Daskin, M. S., Coullard, C. R. and Shen, Z.-J. M. (2002), An inventory-location model: Formulation, solution algorithm and computational results, *Annals of Operations Research* **110**, 83–106.
- Daskin, M. S., Hesse, S. M. and ReVelle, C. S. (1997), α -reliable p -minimax regret: A new model for strategic facility location modeling, *Location Science* **5**(4), 227–246.
- Daskin, M. S., Hogan, K. and ReVelle, C. (1988), Integration of multiple, excess, backup, and expected covering models, *Environment and Planning B* **15**(1), 15–35.
- Daskin, M. S., Hopp, W. J. and Medina, B. (1992), Forecast horizons and dynamic facility location planning, *Annals of Operations Research* **40**, 125–151.
- Daskin, M. S. and Owen, S. H. (1999), Location models in transportation, in R. W. Hall, ed., ‘Handbook of Transportation Science’, Kluwer Academic, Boston, pp. 311–360.
- De Palma, A., Ginsburgh, V., Labbé, M. and Thisse, J.-F. (1989), Competitive location with random utilities, *Transportation Science* **23**(4), 244–252.
- Drezner, Z. (1989), Stochastic analysis of the Weber problem on the sphere, *Journal of the Operational Research Society* **40**(12), 1137–1144.
- Drezner, Z., ed. (1995), *Facility Location: A Survey of Applications and Methods*, Springer-Verlag, New York.
- Drezner, Z. and Hamacher, H. W., eds (2002), *Facility Location: Applications and Theory*, Springer-Verlag, New York.
- Drezner, Z. and Suzuki, A. (2004), The big triangle small triangle method for the solution of nonconvex facility location problems, *Operations Research* **52**(1), 128–135.
- Drezner, Z. and Wesolowsky, G. O. (1980), The expected value of perfect information in facility location, *Operations Research* **28**(2), 395–402.

- Drezner, Z. and Wesolowsky, G. O. (1981), Optimum location probabilities in the l_p distance Weber problem, *Transportation Science* **15**(2), 85–97.
- Eiselt, H. A., Gendreau, M. and Laporte, G. (1992), Location of facilities on a network subject to a single-edge failure, *Networks* **22**, 231–246.
- Eppen, G. D., Martin, R. K. and Schrage, L. (1989), A scenario approach to capacity planning, *Operations Research* **37**(4), 517–527.
- Erlenkotter, D. (1978), A dual-based procedure for uncapacitated facility location, *Operations Research* **26**(6), 992–1009.
- França, P. M. and Luna, H. P. L. (1982), Solving stochastic transportation-location problems by generalized Benders decomposition, *Transportation Science* **16**(2), 113–126.
- Frank, H. (1966), Optimum locations on a graph with probabilistic demands, *Operations Research* **14**(3), 409–421.
- Frank, H. (1967), Optimum locations on graphs with correlated normal demands, *Operations Research* **15**(3), 552–557.
- Geoffrion, A. M. and Graves, G. W. (1974), Multicommodity distribution system design by Benders decomposition, *Management Science* **20**(5), 822–844.
- Ghosh, A. and Craig, C. S. (1983), Formulating retail location strategy in a changing environment, *Journal of Marketing* **47**(3), 56–68.
- Ghosh, A. and McLafferty, S. L. (1982), Locating stores in uncertain environments: A scenario planning approach, *Journal of Retailing* **58**(4), 5–22.
- Goldman, A. J. (1971), Optimal center location in simple networks, *Transportation Science* **5**(2), 212–221.
- Gregg, S. R., Mulvey, J. M. and Wolpert, J. (1988), A stochastic planning system for siting and closing public service facilities, *Environment and Planning A* **20**, 83–98.
- Gupta, S. K. and Rosenhead, J. (1968), Robustness in sequential investment decisions, *Management Science* **15**(2), B18–B29.
- Gutiérrez, G. J. and Kouvelis, P. (1995), A robustness approach to international sourcing, *Annals of Operations Research* **59**, 165–193.

- Gutiérrez, G. J., Kouvelis, P. and Kurawarwala, A. A. (1996), A robustness approach to uncapacitated network design problems, *European Journal of Operational Research* **94**, 362–376.
- Haight, R. G., Ralls, K. and Starfield, A. M. (2000), Designing species translocation strategies when population growth and future funding are uncertain, *Conservation Biology* **14**(5), 1298–1307.
- Hakimi, S. L. (1964), Optimum locations of switching centers and the absolute centers and medians of a graph, *Operations Research* **12**(3), 450–459.
- Hakimi, S. L. (1965), Optimum distribution of switching centers in a communication network and some related graph theoretic problems, *Operations Research* **13**(3), 462–475.
- Hanink, D. M. (1984), A portfolio theoretic approach to multiplant location analysis, *Geographical Analysis* **16**(2), 149–161.
- Hodder, J. E. (1984), Financial market approaches to facility location under uncertainty, *Operations Research* **32**(6), 1374–1380.
- Hodder, J. E. and Dincer, M. C. (1986), A multifactor model for international plant location and financing under uncertainty, *Computers & Operations Research* **13**(5), 601–609.
- Hodder, J. E. and Jucker, J. V. (1985), A simple plant-location model for quantity-setting firms subject to price uncertainty, *European Journal of Operational Research* **21**, 39–46.
- Hodgson, M. J. (1991), Stability of solutions to the p -median problem under induced data error, *INFOR* **29**(2), 167–183.
- Huchzermeier, A. and Cohen, M. A. (1996), Valuing operational flexibility under exchange rate risk, *Operations Research* **44**(1), 100–113.
- Jornsten, K. and Bjorndal, M. (1994), Dynamic location under uncertainty, *Studies in Regional and Urban Planning* **3**, 163–184.
- Jucker, J. V. and Carlson, R. C. (1976), The simple plant-location problem under uncertainty, *Operations Research* **24**(6), 1045–1055.

- Juel, H. (1980), A note on bounds on the Weber problem solution under conditions of uncertainty, *Journal of Regional Science* **20**(4), 523–524.
- Juel, H. (1981), Bounds in the generalized Weber problem under locational uncertainty, *Operations Research* **29**(6), 1219–1227.
- Katz, I. N. and Cooper, L. (1974), An always-convergent numerical scheme for a random locational equilibrium problem, *SIAM Journal on Numerical Analysis* **11**(4), 683–692.
- Killmer, K. A., Anandalingam, G. and Malcolm, S. A. (2001), Siting noxious facilities under uncertainty, *European Journal of Operational Research* **133**, 596–607.
- Kogut, B. (1985), Designing global strategies: Profiting from operational flexibility, *Sloan Management Review* **27**(1), 27–38.
- Kogut, B. and Kulatilaka, N. (1994), Operating flexibility, global manufacturing, and the option value of a multinational network, *Management Science* **40**(1), 123–139.
- Kouvelis, P., Kurawarwala, A. A. and Gutiérrez, G. J. (1992), Algorithms for robust single and multiple period layout planning for manufacturing systems, *European Journal of Operational Research* **63**, 287–303.
- Kouvelis, P. and Yu, G. (1997), *Robust Discrete Optimization and its Applications*, Kluwer Academic Publishers, Boston.
- Labbé, M., Thisse, J.-F. and Wendell, R. E. (1991), Sensitivity analysis in minimun facility location problems, *Operations Research* **39**(6), 961–969.
- Laguna, M., Lino, P., Pérez, A., Quintanilla, S. and Valls, V. (2000), Minimizing weighted tardiness of jobs with stochastic interruptions in parallel machines, *European Journal of Operational Research* **127**, 444–457.
- Laporte, G., Louveaux, F. V. and van Hamme, L. (1994), Exact solution to a location problem with stochastic demands, *Transportation Science* **28**(2), 95–103.
- Larson, R. C. (1974), A hypercube queuing model for facility location and redistricting in urban emergency services, *Computers and Operations Research* **1**, 67–95.
- Larson, R. C. (1975), Approximating the performance of urban emergency service systems, *Operations Research* **23**(5), 845–868.

- LeBlanc, L. J. (1977), A heuristic approach for large scale discrete stochastic transportation-location problems, *Computers and Mathematics with Applications* **3**, 87–94.
- Louveaux, F. and Thisse, J.-F. (1985), Production and location on a network under demand uncertainty, *Operations Research Letters* **4**(4), 145–149.
- Louveaux, F. V. (1986), Discrete stochastic location models, *Annals of Operations Research* **6**, 23–34.
- Louveaux, F. V. (1993), Stochastic location analysis, *Location Science* **1**(2), 127–154.
- Louveaux, F. V. and Peeters, D. (1992), A dual-based procedure for stochastic facility location, *Operations Research* **40**(3), 564–573.
- Lowe, T. J., Wendell, R. E. and Hu, G. (2002), Screening location strategies to reduce exchange rate risk, *European Journal of Operational Research* **136**(3), 573–590.
- Manne, A. S. (1961), Capacity expansion and probabilistic growth, *Econometrica* **29**(4), 632–649.
- Marianov, V. and ReVelle, C. (1996), The queueing maximal availability location problem: A model for the siting of emergency vehicles, *European Journal of Operational Research* **93**, 110–120.
- Mausser, H. E. and Laguna, M. (1998), A new mixed integer formulation for the maximum regret problem, *International Transactions in Operational Research* **5**(5), 389–403.
- Mausser, H. E. and Laguna, M. (1999a), A heuristic to minimax absolute regret for linear programs with interval objective function coefficients, *European Journal of Operational Research* **117**, 157–174.
- Mausser, H. E. and Laguna, M. (1999b), Minimising the maximum relative regret for linear programmes with interval objective function coefficients, *Journal of the Operational Research Society* **50**(10), 1063–1070.
- Mirchandani, P. B. (1980), Locational decisions on stochastic networks, *Geographical Analysis* **12**(2), 172–183.
- Mirchandani, P. B. and Odoni, A. R. (1979), Locations of medians on stochastic networks, *Transportation Science* **13**(2), 85–97.

- Mirchandani, P. B. and Oudjit, A. (1980), Localizing 2-medians on probabilistic and deterministic tree networks, *Networks* **10**(4), 329–350.
- Mirchandani, P. B., Oudjit, A. and Wong, R. T. (1985), ‘Multidimensional’ extensions and a nested dual approach for the m -median problem, *European Journal of Operational Research* **21**(1), 121–137.
- MirHassani, S. A., Lucas, C., Mitra, G., Messina, E. and Poojari, C. A. (2000), Computational solution of capacity planning models under uncertainty, *Parallel Computing* **26**, 511–538.
- Mulvey, J. M., Vanderbei, R. J. and Zenios, S. A. (1995), Robust optimization of large-scale systems, *Operations Research* **43**(2), 264–281.
- Nel, L. D. and Colbourn, C. J. (1990), Locating a broadcast facility in an unreliable network, *INFOR* **28**(4), 363–379.
- Owen, S. H. (1999), Scenario Planning Approaches to Facility Location: Models and Solution Methods, Ph.D. dissertation, Northwestern University.
- Owen, S. H. and Daskin, M. S. (1998), Strategic facility location: A review, *European Journal of Operational Research* **111**(3), 423–447.
- Paraskevopoulos, D., Karakitsos, E. and Rustem, B. (1991), Robust capacity planning under uncertainty, *Management Science* **37**(7), 787–800.
- ReVelle, C. and Hogan, K. (1989), The maximum availability location problem, *Transportation Science* **23**(3), 192–200.
- Ricciardi, N., Tadei, R. and Grosso, A. (2002), Optimal facility location with random throughput costs, *Computers & Operations Research* **29**, 593–607.
- Rosenblatt, M. J. and Lee, H. L. (1987), A robustness approach to facilities design, *International Journal of Production Research* **25**(4), 479–486.
- Rosenhead, J., Elton, M. and Gupta, S. K. (1972), Robustness and optimality as criteria for strategic decisions, *Operational Research Quarterly* **23**(4), 413–431.
- Santoso, T., Ahmed, S., Goetschalckx, M. and Shapiro, A. (2003), A stochastic programming approach for supply chain network design under uncertainty, working paper, Georgia Institute of Technology, available from Optimization Online.

- Schilling, D. A. (1982), Strategic facility planning: The analysis of options, *Decision Sciences* **13**, 1–14.
- Serra, D. and Marianov, V. (1998), The p -median problem in a changing network: The case of Barcelona, *Location Science* **6**, 383–394.
- Serra, D., Ratick, S. and ReVelle, C. (1996), The maximum capture problem with uncertainty, *Environment and Planning B* **23**, 49–59.
- Shen, Z.-J. M., Coullard, C. R. and Daskin, M. S. (2003), A joint location-inventory model, *Transportation Science* **37**(1), 40–55.
- Sheppard, E. S. (1974), A conceptual framework for dynamic location-allocation analysis, *Environment and Planning A* **6**, 547–564.
- Shiode, S. and Drezner, Z. (2003), A competitive facility location problem on a tree network with stochastic weights, *European Journal of Operational Research* **149**, 47–52.
- Snyder, L. V. (2003), A note on the robust international sourcing algorithm of Gutiérrez and Kouvelis, working paper, Lehigh University.
- Snyder, L. V. and Daskin, M. S. (2003), Stochastic p -robust location problems, working paper, Northwestern University.
- Snyder, L. V. and Daskin, M. S. (2004), Reliability models for facility location: The expected failure cost case, working paper, Lehigh University.
- Snyder, L. V., Daskin, M. S. and Teo, C.-P. (2003), The stochastic location model with risk pooling, working paper, Northwestern University.
- Trafalis, T. B., Mishina, T. and Foote, B. L. (1999), An interior point multiobjective programming approach for production planning with uncertain information, *Computers & Industrial Engineering* **37**, 631–648.
- Tsiakis, P., Shah, N. and Pantelides, C. C. (2001), Design of multi-echelon supply chain networks under demand uncertainty, *Industrial and Engineering Chemistry Research* **40**, 3585–3604.
- Vairaktarakis, G. L. and Kouvelis, P. (1999), Incorporation dynamic aspects and uncertainty in 1-median location problems, *Naval Research Logistics* **46**(2), 147–168.

- Verter, V. and Dincer, M. C. (1992), An integrated evaluation of facility location, capacity acquisition, and technology selection for designing global manufacturing strategies, *European Journal of Operational Research* **60**, 1–18.
- Vidal, C. J. and Goetschalckx, M. (1997), Strategic production-distribution models: A critical review with emphasis on global supply chain models, *European Journal of Operational Research* **98**(1), 1–18.
- Vidal, C. J. and Goetschalckx, M. (2000), Modeling the effect of uncertainties on global logistics systems, *Journal of Business Logistics* **21**(1), 95–120.
- Vladimirou, H. and Zenios, S. A. (1997), Stochastic linear programs with restricted recourse, *European Journal of Operational Research* **101**, 177–192.
- Weaver, J. R. and Church, R. L. (1983), Computational procedures for location problems on stochastic networks, *Transportation Science* **17**(2), 168–180.
- Wesolowsky, G. O. (1977), Probabilistic weights in the one-dimensional facility location problem, *Management Science* **24**(2), 224–229.
- Williams, A. C. (1963), A stochastic transportation problem, *Operations Research* **11**(5), 759–770.
- Yu, C.-S. and Li, H.-L. (2000), A robust optimization model for stochastic logistic problems, *International Journal of Production Economics* **64**, 385–397.
- Yu, G. (1997), Robust economic order quantity models, *European Journal of Operational Research* **100**, 482–493.
- Yu, G. and Yang, J. (1998), On the robust shortest path problem, *Computers & Operations Research* **25**(6), 457–468.