

Planning for Disruptions in Supply Chain Networks

Lawrence V. Snyder

Department of Industrial and Systems Engineering, Lehigh University, Mohler Lab,
200 West Packer Avenue, Bethlehem, Pennsylvania 18013, larry.snyder@lehigh.edu

Maria P. Scaparra

Kent Business School, University of Kent, Canterbury, CT2 7PE, England,
m.p.scaparra@kent.ac.uk

Mark S. Daskin

Department of Industrial Engineering and Management Sciences,
Northwestern University, 2145 Sheridan Road, Evanston, Illinois 60208,
m-daskin@northwestern.edu

Richard L. Church

Department of Geography, University of California, Santa Barbara, California 593106-4060,
church@geog.ucsb.edu

Abstract Recent events have highlighted the need for planners to consider the risk of disruptions when designing supply chain networks. Supply chain disruptions have a number of causes and may take a number of forms. Once a disruption occurs, there is very little recourse regarding supply chain infrastructure because these strategic decisions cannot be changed quickly. Therefore, it is critical to account for disruptions during the design of supply chain networks so that they perform well even after a disruption. Indeed, these systems can often be made substantially more reliable with only small additional investments in infrastructure. Planners have a range of options available to them in designing resilient supply chain networks, and their choice of approaches will depend on the financial resources available, the decision maker's risk preference, the type of network under consideration, and other factors. In this tutorial, we present a broad range of models for designing supply chains resilient to disruptions. We first categorize these models by the status of the existing network: A network may be designed from scratch, or an existing network may be modified to prevent disruptions at some facilities. We next divide each category based on the underlying optimization model (facility location or network design) and the risk measure (expected cost or worst-case cost).

Keywords facility location; network design; disruptions

1. Introduction

1.1. Motivation

Every supply chain faces disruptions of various sorts. Recent examples of major disruptions are easy to bring to mind: Hurricanes Katrina and Rita in 2005 on the U.S. Gulf Coast crippled the nation's oil refining capacity (Mouawad [68]), destroyed large inventories of coffee and lumber (Barrionuevo and Deutsch [3], Reuters [74]), and forced the rerouting of bananas and other fresh produce (Barrionuevo and Deutsch [3]). A strike at two General Motors parts plants in 1998 led to the shutdowns of 26 assembly plants, which ultimately

resulted in a production loss of over 500,000 vehicles and an \$809 million quarterly loss for the company (Brack [13], Simison [88, 89]). An eight-minute fire at a Philips semiconductor plant in 2001 brought one customer, Ericsson, to a virtual standstill while another, Nokia, weathered the disruption (Latour [58]). Moreover, smaller-scale disruptions occur much more frequently. For example, Wal-Mart's Emergency Operations Center receives a call virtually every day from a store or other facility with some sort of crisis (Leonard [60]).

There is evidence that superior contingency planning can significantly mitigate the effect of a disruption. For example, Home Depot's policy of planning for various types of disruptions based on geography helped it get 23 of its 33 stores within Katrina's impact zone open after one day and 29 after one week (Fox [37]), and Wal-Mart's stock prepositioning helped make it a model for post-hurricane recovery (Leonard [60]). Similarly, Nokia weathered the 2001 Phillips fire through superior planning and quick response, ultimately allowing it to capture a substantial portion of Ericsson's market share (Latour [58]).

Recent books and articles in the business and popular press have pointed out the vulnerability of today's supply chains to disruptions and the need for a systematic analysis of supply chain vulnerability, security, and resiliency (Elkins et al. [35], Jüttner et al. [52], Lynn [63], Rice and Caniato [76], Sheffi [84]). One common theme among these references is that the tightly optimized, just-in-time, lean supply chain practices championed by practitioners and OR researchers in recent decades increase the vulnerability of these systems. Many have argued that supply chains should have more redundancy or slack to provide a buffer against various sorts of uncertainty. Nevertheless, companies have historically been reluctant to invest much in additional supply chain infrastructure or inventory, despite the large payoff that such investments can have if a disruption occurs.

We argue that decision makers should take supply uncertainty (of which disruptions are one variety) into account during all phases of supply chain planning, just as they account for demand uncertainty. This is most critical during strategic planning because these decisions cannot easily be modified. When a disruption strikes, there is very little recourse for strategic decisions like facility location and network design. (In contrast, firms can often adjust inventory levels, routing plans, production schedules, and other tactical and operational decisions in real time in response to unexpected events.)

It is easy to view supply uncertainty and demand uncertainty as two sides of the same coin. For example, a toy manufacturer may view stockouts of a hot new toy as a result of demand uncertainty, but to a toy store, the stockouts look like a supply-uncertainty issue. Many techniques that firms use to mitigate demand uncertainty—safety stock, supplier redundancy, forecast refinements—also apply in the case of supply uncertainty. However, it is dangerous to assume that supply uncertainty is a special case of demand uncertainty or that it can be ignored by decision makers, because much of the conventional wisdom gained from studying demand uncertainty does not hold under supply uncertainty. For example, under demand uncertainty, it may be optimal for a firm to operate fewer distribution centers (DCs) because of the risk-pooling effect and economies of scale in ordering (Daskin et al. [27]), while under supply uncertainty, it may be optimal to operate more, smaller DCs so that a disruption to one of them has lesser impact. Snyder and Shen [95] discuss this and other differences between the two forms of uncertainty.

In this tutorial, we discuss models for designing supply chain networks that are resilient to disruptions. The objective is to design the supply chain infrastructure so that it operates efficiently (i.e., at low cost) both normally and when a disruption occurs. We discuss models for facility location and network design. Additionally, we analyze fortification models that can be used to improve the reliability of infrastructure systems already in place and for which a complete reconfiguration would be cost prohibitive. The objective of fortification models is to identify optimal strategies for allocating limited resources among possible mitigation investments.

1.2. Taxonomy and Tutorial Outline

We classify models for reliable supply chain design along three axes.

(1) *Design vs. fortification.* Is the model intended to create a reliable network assuming that no network is currently in place, or to fortify an existing network to make it more reliable?

(2) *Underlying model.* Reliability models generally have some classical model as their foundation. In this tutorial, we consider models based on facility location and network design models.

(3) *Risk measure.* As in the case of demand uncertainty, models with supply uncertainty need some measure for evaluating risk. Examples include expected cost and minimax cost.

This tutorial is structured according to this taxonomy. Section 3 discusses design models, while §4 discusses fortification models, with subsections in each to divide the models according to the remaining two axes. These sections are preceded by a review of the related literature in §2 and followed by conclusions in §5.

2. Literature Review

We discuss the literature that is directly related to reliable supply chain network design throughout this tutorial. In this section, we briefly discuss several streams of research that are indirectly related. For more detailed reviews of facility location models under uncertainty, the reader is referred to Daskin et al. [29], Owen and Daskin [70], and Snyder [90]. See Daskin [26] or Drezner [33] for a textbook treatment of facility location theory. An excellent overview of stochastic programming theory in general is provided in Hige [45].

2.1. Network Reliability Theory

The concept of supply chain reliability is related to network reliability theory (Colbourn [22], Shier [86], Shooman [87]), which is concerned with calculating or maximizing the probability that a graph remains connected after random failures due to congestion, disruptions, or blockages. Typically, this literature considers disruptions to the links of a network, but some papers consider node failures (Eiselt et al. [34]), and in some cases the two are equivalent. Given the difficulty in computing the reliability of a given network, the goal is often to find the minimum-cost network with some desirable property like two-connectivity (Monma [66], Monma and Shalcross [67]), k -connectivity (Bienstock et al. [11], Grötschel et al. [41]), or special ring structures (Fortz and Labbe [36]). The key difference between network reliability models and the models we discuss in this tutorial is that network reliability models are primarily concerned with connectivity; they consider the cost of constructing the network but not the cost that results from a disruption, whereas our models consider both types of costs and generally assume connectivity after a disruption.

2.2. Vector-Assignment Problems

Weaver and Church [104] introduce the *vector-assignment P -median problem* (VAPMP), in which each customer is assigned to several open facilities according to an exogenously determined frequency. For example, a customer might receive 75% of its demand from its nearest facility, 20% from its second nearest, and 5% from its third nearest. This is similar to the assignment strategy used in many of the models below, but in our models the percentages are determined endogenously based on disruptions rather than given as inputs to the model. A vector-assignment model based on the uncapacitated fixed-charge location problem (UFLP) is presented by Pirkul [73].

2.3. Multiple, Excess, and Backup Coverage Models

The maximum covering problem (Church and ReVelle [17]) locates a fixed number of facilities to maximize the demands located within some radius of an open facility. It implicitly

assumes that the facilities (e.g., fire stations, ambulances) are always available. Several subsequent papers have considered the congestion at facilities when multiple calls are received at the same time. The maximum expected covering location model (MEXCLM) (Daskin [24, 25]) maximizes the expected coverage given a constant, systemwide probability that a server is busy at any given time. The constant-busy-probability assumption is relaxed in the maximum availability location problem (MALP) (ReVelle and Hogan [75]). A related stream of research explicitly considers the queueing process at the locations; these “hypercube” models are interesting as descriptive models but are generally too complex to embed into an optimization framework (Berman et al. [10], Larson [56, 57]). See Berman and Krass [6] and Daskin et al. [28] for a review of expected and backup coverage models. The primary differences between these models and the models we discuss in this tutorial are (1) the objective function (coverage versus cost), and (2) the reason for a server’s unavailability (congestion versus disruptions).

2.4. Inventory Models with Supply Disruptions

There is a stream of research in the inventory literature that considers supply disruptions in the context of classical inventory models, such as the EOQ (Parlar and Berkin [72], Berk and Arreola-Risa [5], Snyder [91]), (Q, R) (Gupta [42], Parlar [71], Mohebbi [64, 65]), and (s, S) (Arreola-Risa and DeCroix [1]) models. More recent models examine a range of strategies for mitigating disruptions, including dual sourcing (Tomlin [100]), demand management (Tomlin [99]), supplier reliability forecasting (Tomlin [98], Tomlin and Snyder [101]), and product-mix flexibility (Tomlin and Wang [102]). Few models consider disruptions in multiechelon supply chain or inventory systems; exceptions include Kim et al. [53], Hopp et al. [47], and Snyder and Shen [95].

2.5. Process Flexibility

At least five strategies can be employed in the face of uncertain demands: Expanding capacity, holding reserve inventory, improving demand forecasts, introducing product commonality to delay the need for specialization, and adding flexibility to production plants. A complete review of each strategy is beyond the scope of this tutorial. Many of these strategies are fairly straightforward. Process flexibility, on the other hand, warrants a brief discussion. Jordan and Graves [51] compare the expected lost sales that result from using a set of fully flexible plants, in which each plant could produce each product, to a configuration in which each plant produces only two products and the products are chained in such a way that plant A produces products 1 and 2, plant B produces products 2 and 3, and so on, with the last plant producing the final product as well as product 1. They refer to this latter configuration as a 1-chain. They find that a 1-chain provides nearly all of the benefits of total flexibility when measured by the expected number of lost sales. Based on this, they recommend that flexibility be added to create fewer, longer chains of products and plants. Bish et al. [12] study capacity allocation schemes for such chains (e.g., allocate capacity to the nearest demands, to the highest-margin demands, or to a plant’s primary product). They find that if the capacity is either very small or very large relative to the expected demand, the gains from managing flexible capacity are outweighed by the need for additional component inventory at the plants and the costs of order variability at suppliers. They then provide guidelines for the use of one allocation policy relative to others based on the costs of component inventory, component lead times, and profit margins. Graves and Tomlin [40] extend the Jordan and Graves results to multistage systems. They contrast configuration loss with configuration inefficiency. The former measures the difference between the shortfall with total flexibility and the shortfall with a particular configuration of flexible plants. The configuration inefficiency measures the effect of the interaction between stages in causing the shortfall for a particular configuration. They show that this, in turn, is caused by two phenomena: Floating bottlenecks and stage-spanning bottlenecks. Stage-spanning bottlenecks

can arise even if demand is deterministic, as a result of misallocations of capacity across the various stages of the supply chain. Beach et al. [4] and de Toni and Tonchia [30] provide more detailed reviews of the manufacturing flexibility literature.

2.6. Location of Protection Devices

A number of papers in the location literature have addressed the problem of finding the optimal location of protection devices to reduce the impact of possible disruptions to infrastructure systems. For example, Carr et al. [16] present a model for optimizing the placement of sensors in water supply networks to detect maliciously injected contaminants. James and Salhi [49] investigate the problem of placing protection devices in electrical supply networks to reduce the amount of outage time. Flow-interception models (Berman et al. [7]) have also been used to locate protection facilities. For example, Hodgson et al. [46] and Gendreau et al. [39] use flow-interception models to locate inspection stations so as to maximize hazard avoidance and risk reduction in transportation networks. The protection models discussed in this chapter differ from those models in that they do not seek the optimal placement of physical protection devices or facilities. Rather, they aim at identifying the most critical system components to harden or protect with limited protection resources (for example, through structural retrofit, fire safety, increased surveillance, vehicle barriers, and monitoring systems).

3. Design Models

3.1. Introduction

In this section, we discuss *design models* for reliable facility location and network design. These models, like most facility location models, assume that no facilities currently exist; they aim to choose a set of facility locations that perform well even if disruptions occur. It is also straightforward to modify these models to account for facilities that may already exist (e.g., by setting the fixed cost of those facilities to zero or adding a constraint that requires them to be open). In contrast, the *fortification models* discussed in §4 assume that all facility sites have been chosen and attempt to decide which facilities to fortify (protect against disruptions). One could conceivably formulate an integrated design/fortification model whose objective would be to locate facilities and identify a subset of those facilities to fortify against attacks. Formulation of such a model is a relatively straightforward extension of the models we present below, though its solution would be considerably more difficult because it would result in (at least) a tri-level optimization problem.

Most models for both classical and reliable facility location are design models, because as “fortification” is a relatively new concept in the facility location literature. In the subsections that follow, we introduce several design models, classified first according to the underlying model (facility location or network design) and then according to risk measure (expected or worst-case cost).

3.2. Facility Location Models

3.2.1. Expected Cost Models. In this section, we define the reliability fixed-charge location problem (RFLP) (Snyder and Daskin [92]), which is based on the classical uncapacitated fixed-charge location problem (UFLP) (Balinski [2]). There is a fixed set I of customer locations and a set J of potential facility locations. Each customer $i \in I$ has an annual demand of h_i units, and each unit shipped from facility $j \in J$ to customer $i \in I$ incurs a transportation cost of d_{ij} . (We will occasionally refer to d_{ij} as the “distance” between j and i , and use this notion to refer to “closer” or “farther” facilities.) Each facility site has an annual fixed cost f_j that is incurred if the facility is opened. Any open facility may serve any customer (that is, there are no connectivity restrictions), and facilities have unlimited capacity. There is a single product.

Each open facility may fail (be disrupted) with a fixed probability q . (Note that the failure probability q is the same at every facility. This assumption allows a compact description of the expected transportation cost. Below, we relax this assumption and instead formulate a scenario-based model that requires more decision variables but is more flexible.) Failures are independent, and multiple facilities may fail simultaneously. When a facility fails, it cannot provide any product, and the customers assigned to it must be reassigned to a nondisrupted facility.

If customer i is not served by any facility, the firm incurs a penalty cost of θ_i per unit of demand. This penalty may represent a lost-sales cost or the cost of finding an alternate source for the product. It is incurred if all open facilities have failed, or if it is too expensive to serve a customer from its nearest functional facility. To model this, we augment the facility set J to include a dummy “emergency facility,” called u , that has no fixed cost ($f_u = 0$) and never fails. The transportation cost from u to i is $d_{iu} \equiv \theta_i$. Assigning a customer to the emergency facility is equivalent to not assigning it at all.

The RFLP uses two sets of decision variables:

$$X_j = \begin{cases} 1, & \text{if facility } j \text{ is opened,} \\ 0, & \text{otherwise,} \end{cases}$$

$$Y_{ijr} = \begin{cases} 1, & \text{if customer } i \text{ is assigned to facility } j \text{ at level } r, \\ 0, & \text{otherwise.} \end{cases}$$

A “level- r ” assignment is one for which there are r closer open facilities. For example, suppose that the three closest open facilities to customer i are facilities 2, 5, and 8, in that order. Then facility 2 is i ’s level-0 facility, 5 is its level-1 facility, and 8 is its level-2 facility. Level-0 assignments are to “primary” facilities that serve the customer under normal circumstances, while level- r assignments ($r > 0$) are to “backup” facilities that serve it if all closer facilities have failed. A customer must be assigned to *some* facility at each level r unless it is assigned to the emergency facility at some level $s \leq r$. Because we do not know in advance how many facilities will be open, we extend the index r from 0 through $|J| - 1$, but Y_{ijr} will equal 0 for r greater than or equal to the number of open facilities.

The objective of the RFLP is to choose facility locations and customer assignments to minimize the fixed cost plus the expected transportation cost and lost-sales penalty. We formulate it as an integer programming problem as follows.

$$\text{(RFLP) minimize } \sum_{j \in J} f_j X_j + \sum_{i \in I} \sum_{r=0}^{|J|-1} \left[\sum_{j \in J \setminus \{u\}} h_i d_{ij} q^r (1-q) Y_{ijr} + h_i d_{iu} q^r Y_{iur} \right] \quad (1)$$

subject to

$$\sum_{j \in J} Y_{ijr} + \sum_{s=0}^{r-1} Y_{ius} = 1 \quad \forall i \in I, r = 0, \dots, |J| - 1 \quad (2)$$

$$Y_{ijr} \leq X_j \quad \forall i \in I, j \in J, r = 0, \dots, |J| - 1 \quad (3)$$

$$\sum_{r=0}^{|J|-1} Y_{ijr} \leq 1 \quad \forall i \in I, j \in J \quad (4)$$

$$X_j \in \{0, 1\} \quad \forall j \in J \quad (5)$$

$$Y_{ijr} \in \{0, 1\} \quad \forall i \in I, j \in J, r = 0, \dots, |J| - 1 \quad (6)$$

The objective function (1) minimizes the sum of the fixed cost and the expected transportation and lost-sales costs. The second term reflects the fact that if customer i is assigned to facility j at level r , it will actually be served by j if all r closer facilities have failed (which

happens with probability q^r) and if j itself has not failed (which happens with probability $1 - q$). Note that we can compute this expected cost knowing only the *number* of facilities that are closer to i than j , but not which facilities those are. This is a result of our assumption that every facility has the same failure probability. If, instead, customer i is assigned to the emergency facility at level r , then it incurs the lost-sales cost $d_{iu} \equiv \theta_i$ if its r closest facilities have failed (which happens with probability q^r).

Constraints (2) require each customer i to be assigned to some facility at each level r , unless i has been assigned to the emergency facility at level $s < r$. Constraints (3) prevent an assignment to a facility that has not been opened, and constraints (4) prohibit a customer from being assigned to the same facility at more than one level. Constraints (5) and (6) require the decision variables to be binary. However, constraints (6) can be relaxed to nonnegativity constraints because single sourcing is optimal in this problem, as it is in the UFLP.

Note that we do not explicitly enforce the definition of “level- r assignment” in this formulation; that is, we do not require $Y_{ijr} = 1$ only if there are exactly r closer open facilities. Nevertheless, in any optimal solution, this definition will be satisfied because it is optimal to assign customers to facilities by levels in increasing order of distance. This is true because the objective function weights decrease for larger values of r , so it is advantageous to use facilities with smaller d_{ij} at smaller assignment levels. A slight variation of this result is proven rigorously by Snyder and Daskin [92].

Snyder and Daskin [92] present a slightly more general version of this model in which some of the facilities may be designated as “nonfailable.” If a customer is assigned to a nonfailable facility at level r , it does not need to be assigned at any higher level. In addition, Snyder and Daskin [92] consider a multiobjective model that minimizes the weighted sum of two objectives, one of which corresponds to the UFLP cost (fixed cost plus level-0 transportation costs) while the other represents the expected transportation cost (accounting for failures). By varying the weights on the objectives, Snyder and Daskin [92] generate a trade-off curve and use this to demonstrate that the RFLP can produce solutions that are much more reliable than the classical UFLP solution but only slightly more expensive by the UFLP objective. This suggests that reliability can be “bought” relatively cheaply. Finally, Snyder and Daskin [92] also consider a related model that is based on the P -median problem (Hakimi [43, 44]) rather than the UFLP. They solve all models using Lagrangian relaxation.

In general, the optimal solution to the RFLP uses more facilities than that of the UFLP. This tendency toward diversification occurs so that any given disruption affects a smaller portion of the system. It may be viewed as a sort of “risk-diversification effect” in which it is advantageous to spread the risk of supply uncertainty across multiple facilities (encouraging decentralization). This is in contrast to the classical risk-pooling effect, which encourages centralization to pool the risk of demand uncertainty (Snyder and Shen [95]).

Berman et al. [8] consider a model similar to (RFLP), based on the P -median problem rather than the UFLP. They allow different facilities to have different failure probabilities, but the resulting model is highly nonlinear and, in general, must be solved heuristically. They prove that the Hakimi property applies if colocation is allowed. (The Hakimi property says that optimal locations exist at the nodes of a network, even if facilities are allowed on the links.) Berman et al. [9] present a variant of this model in which customers do not know which facilities are disrupted before visiting them and must traverse a path from one facility to the next until an operational facility is found. For example, a customer might walk to the nearest ATM, find it out of order, and then walk to the ATM that is nearest to the current location. They investigate the spatial characteristics of the optimal solution and discuss the value of reliability information.

An earlier attempt at addressing reliability issues in P -median problems is discussed by Drezner [32], who examines the problem of locating P unreliable facilities in the plane so as to minimize expected travel distances between customers and facilities. As in the RFLP,

the unreliable P -median problem in Drezner [32] is defined by introducing a probability that a facility becomes inactive but does not require the failures to be independent events. The problem is solved through a heuristic procedure. A more sophisticated method to solve the unreliable P -median problem was subsequently proposed in Lee [59]. Drezner [32] also presents the unreliable (P, Q) -center problem where P facilities must be located while taking into account that Q of them may become unavailable simultaneously. The objective is to minimize the maximal distance between demand points and their closest facilities.

The formulation given above for (RFLP) captures the expected transportation cost without using explicit scenarios to describe the uncertain events (disruptions). An alternate approach is to model the problem as a two-stage stochastic programming problem in which the location decisions are first-stage decisions and the assignment decisions are made in the second stage, after the random disruptions have occurred. This approach can result in a much larger IP model because $2^{|J|}$ possible failure scenarios exist, and each requires its own assignment variables. That is, in the formulation above we have $|J|$ Y variables for each i, j (indexed $Y_{ijr}, r = 0, \dots, |J| - 1$), while in the scenario-based formulation we have $2^{|J|}$ variables for each i, j . However, formulations built using this approach can be solved using standard stochastic programming methods. They can also be adapted more readily to handle side constraints and other variations.

For example, suppose facility j can serve at most b_j units of demand at any given time. These capacity constraints must be satisfied both by “primary” assignments and by reassignments that occur after disruptions. Let S be the set of failure scenarios such that $a_{js} = 1$ if facility j fails in scenario s , and let q_s be the probability that scenario s occurs. Finally, let Y_{ijs} equal 1 if customer i is assigned to facility j in scenario s and 0 otherwise. The capacitated RFLP can be formulated using the scenario-based approach as follows.

$$\text{(CRFLP)} \quad \text{minimize} \quad \sum_{j \in J} f_j X_j + \sum_{s \in S} q_s \sum_{i \in I} \sum_{j \in J} h_i d_{ij} Y_{ijs} \quad (7)$$

subject to

$$\sum_{j \in J} Y_{ijs} = 1 \quad \forall i \in I, s \in S \quad (8)$$

$$Y_{ijs} \leq X_j \quad \forall i \in I, j \in J, s \in S \quad (9)$$

$$\sum_{i \in I} h_i Y_{ijs} \leq (1 - a_{js}) b_j \quad \forall j \in J, s \in S \quad (10)$$

$$X_j \in \{0, 1\} \quad \forall j \in J \quad (11)$$

$$Y_{ijs} \in \{0, 1\} \quad \forall i \in I, j \in J, s \in S \quad (12)$$

Note that the set J in this formulation still includes the emergency facility u . The objective function (7) computes the sum of the fixed cost plus the expected transportation cost, taken across all scenarios. Constraints (8) require every customer to be assigned to some facility (possibly u) in every scenario, and constraints (9) require this facility to be opened. Constraints (10) prevent the total demand assigned to facility j in scenario s from exceeding j 's capacity and prevent any demand from being assigned if the facility has failed in scenario s . Constraints (11) and (12) are integrality constraints. Integrality can be relaxed to nonnegativity for the Y variables if single-sourcing is not required. (Single-sourcing is no longer optimal because of the capacity constraints.)

(CRFLP) can be modified easily without destroying its structure, in a way that (RFLP) cannot. For example, if the capacity during a disruption is reduced but not eliminated, we can simply redefine a_{js} to be the proportion of the total capacity that is affected by the disruption. We can also easily allow the demands and transportation costs to be scenario dependent.

The disadvantage, of course, is that the number of scenarios grows exponentially with $|J|$. If $|J|$ is reasonably large, enumerating all of the scenarios is impractical. In this case, one generally must use sampling techniques such as sample average approximation (SAA) (Kleywegt et al. [54], Linderoth et al. [62], Shapiro and Homem-de-Mello [83]), in which the optimization problem is solved using a subset of the scenarios sampled using Monte Carlo simulation. By solving a series of such problems, one can develop bounds on the optimal objective value and the objective value of a given solution. Ülker and Snyder [103] present a method for solving (CRFLP) that uses Lagrangian relaxation embedded in an SAA scheme.

An ongoing research project has focused on extending the models discussed in this section to account for inventory costs when making facility location decisions. Jeon et al. [50] consider facility failures in a location-inventory context that is similar to the models proposed recently by Daskin et al. [27] and Shen et al. [85], which account for the cost of cycle and safety stock. The optimal number of facilities in the models by Daskin et al. [27] and Shen et al. [85] is smaller than those in the UFLP due to economies of scale in ordering and the risk-pooling effect. Conversely, the optimal number of facilities is larger in the RFLP than in the UFLP to reduce the impact of any single disruption. The location-inventory model with disruptions proposed by Jeon et al. [50] finds a balance between these two competing tendencies.

3.2.2. Worst-Case Cost Models. Models that minimize the expected cost, as in §3.2.1, take a risk-neutral approach to decision making under uncertainty. Risk-averse decision makers may be more inclined to minimize the worst-case cost, taken across all scenarios. Of course, in this context, it does not make sense to consider all possible scenarios, because otherwise the worst-case scenario is always the one in which *all* facilities fail. Instead, we might consider all scenarios in which, say, at most three facilities fail, or all scenarios with probability at least 0.01, or some other set of scenarios identified by managers as worth planning against. In general, the number of scenarios in such a problem is smaller than in the expected-cost problem because scenarios that are clearly less costly than other scenarios can be omitted from consideration. For example, if we wish to consider scenarios in which at most three facilities fail, we can ignore scenarios in which two or fewer fail.

To formulate the minimax-cost RFLP, we introduce a single additional decision variable U , which equals the maximum cost.

$$\text{(MMRFLP)} \quad \text{minimize } U \quad (13)$$

subject to

$$\sum_{j \in J} f_j X_j + \sum_{i \in I} \sum_{j \in J} h_i d_{ij} Y_{ijs} \leq U \quad \forall s \in S \quad (14)$$

$$\sum_{j \in J} Y_{ijs} = 1 \quad \forall i \in I, s \in S \quad (15)$$

$$Y_{ijs} \leq (1 - a_{js}) X_j \quad \forall i \in I, j \in J, s \in S \quad (16)$$

$$X_j \in \{0, 1\} \quad \forall j \in J \quad (17)$$

$$Y_{ijs} \in \{0, 1\} \quad \forall i \in I, j \in J, s \in S \quad (18)$$

In this formulation, we omit the capacity constraints (10), but they can be included without difficulty. Unfortunately, minimax models tend to be much more difficult to solve exactly, either with general-purpose IP solvers or with customized algorithms. This is true for classical problems as well as for (MMRFLP).

The *regret* of a solution under a given scenario is the relative or absolute difference between the cost of the solution under that scenario and the optimal cost under that scenario. One can modify (MMRFLP) easily to minimize the maximum regret across all scenarios by replacing the right side of (14) with $U + z_s$ (for absolute regret) or $z_s(1 + U)$ (for relative regret).

Here, z_s is the optimal cost in scenario s , which must be determined exogenously for each scenario and provided as an input to the model.

Minimax-regret problems may require more scenarios than their minimax-cost counterparts because it is not obvious a priori which scenarios will produce the maximum regret. On the other hand, they tend to result in a less pessimistic solution than minimax-cost models do. Snyder and Daskin [94] discuss minimax-cost and minimax-regret models in further detail.

One common objection to minimax models is that they are overly conservative because the resulting solution plans against a single scenario, which may be unlikely even if it is disastrous. In contrast, expected-cost models like the CRFLP produce solutions that perform well in the long run but may perform poorly in some scenarios. Snyder and Daskin [94] introduce a model that avoids both problems by minimizing the expected cost (7) subject to a constraint on the maximum cost that can occur in any scenario (in effect, treating U as a constant in (14)). An optimal solution to this model is guaranteed to perform well in the long run (due to the objective function) but is also guaranteed not to be disastrous in any given scenario. This approach is closely related to the concept of p -robustness in robust optimization problems (Kouvelis and Yu [55], Snyder and Daskin [93]). One computational disadvantage is that, unlike the other models we have discussed, it can be difficult (even NP-complete) to find a *feasible* solution or to determine whether a given instance is feasible. See Snyder and Daskin [94] for more details on this model and for a discussion of reliable facility location under a variety of other risk measures.

Church et al. [20] use a somewhat different approach to model worst-case cost design problems, the rationale being that the assumption of independent facility failures underlying the previous models does not hold in all application settings. This is particularly true when modeling intentional disruptions. As an example, a union or a terrorist could decide to strike those facilities in which the greatest combined harm (as measured by increased costs, disrupted service, etc.) is achieved. To design supply systems able to withstand intentional harms by intelligent perpetrators, Church et al. [20] propose the resilient P -median problem. This model identifies the best location of P facilities so that the system works as well as possible (in terms of weighted distances) in the event of a maximally disruptive strike. The model is formulated as a bilevel optimization model, in which the upper-level problem of optimally locating P facilities embeds a lower-level optimization problem used to generate the weighted distance after a worst-case loss of R of these located P facilities. This bilevel programming approach has been widely used to assess worst-case scenarios and identify critical components in existing systems and will be discussed in more depth in §4.2.2. Church et al. [20] demonstrate that optimal P -median configurations can be rendered very inefficient in terms of worst-case loss, even for small values of R . They also demonstrate that resilient design configurations can be near optimal in efficiency as compared to the optimal P -median configurations, but at the same time, maintain high levels of efficiency after worst-case loss. A form of the resilient design problem has also been developed for a coverage-type service system (O'Hanley and Church [69]). The resilient coverage model finds the optimal location of a set of facilities to maximize a combination of initial demand coverage and the minimum coverage level following the loss of one or more facilities. There are several approaches that one can employ to solve this problem, including the successive use of super-valid inequalities (O'Hanley and Church [69]), reformulation into a single-level optimization problem when $R = 1$ or $R = 2$ (Church et al. [20]), or by developing a special search tree. Research is underway to model resilient design for capacitated problems.

3.3. Network Design Models

We now turn our attention from reliability models based on facility location problems to those based on network design models. We have a general network $G = (V, A)$. Each node $i \in V$ serves as either a source, sink, or transshipment node. Source nodes are analogous

to facilities in §3.2 while sink nodes are analogous to customers. The primary difference between network design models and facility location ones is the presence of transshipment nodes. Product originates at the source nodes and is sent through the network to the sink nodes via transshipment nodes.

Like the facilities in §3.2, the nonsink nodes in these models can fail randomly. The objective is to make open/close decisions on the nonsink nodes (first-stage variables) and determine the flows on the arcs in each scenario (second-stage variables) to minimize the expected or worst-case cost. (Many classical network design problems involve open/close decisions on arcs, but the two are equivalent through a suitable transformation.)

3.3.1. Expected Cost. Each node $j \in V$ has a supply b_j . For source nodes, b_j represents the available supply and $b_j > 0$; for sink nodes, b_j represents the (negative of the) demand and $b_j < 0$; for transshipment nodes, $b_j = 0$. There is a fixed cost f_j to open each nonsink node. Each arc (i, j) has a cost of d_{ij} for each unit of flow transported on it, and each nonsink node j has a capacity k_j . The node capacities can be seen as production limitations for the supply nodes and processing resource restrictions for the transshipment nodes.

As in §3.2.1, we let S be the set of scenarios, and $a_{js} = 1$ if node j fails in scenario s . Scenario s occurs with probability q_s . To ensure feasibility in each scenario, we augment V by adding a dummy source node u that makes up any supply shortfall caused by disruptions and a dummy sink node v that absorbs any excess supply. There is an arc from u to each (nondummy) sink node; the per-unit cost of this arc is equal to the lost-sales cost for that sink node (analogous to θ_i in §3.2.1). Similarly, there is an arc from each (nondummy) source node to v whose cost equals 0. The dummy source node and the dummy sink node have infinite supply and demand, respectively.

Let $V_0 \subseteq V$ be the set of supply and transshipment nodes, i.e., $V_0 = \{j \in V \mid b_j \geq 0\}$. We define two sets of decision variables. $X_j = 1$ if node i is opened and 0 otherwise, for $j \in V_0$, and Y_{ijs} is the amount of flow sent on arc $(i, j) \in A$ in scenario $s \in S$. Note that the set A represents the augmented set of arcs, including the arcs outbound from the dummy source node and the arcs inbound to the dummy sink node. With this notation, the reliable network design model (RNDP) is formulated as follows.

$$(RNDP) \quad \text{minimize} \quad \sum_{j \in V_0} f_j X_j + \sum_{s \in S} q_s \sum_{(i, j) \in A} d_{ij} Y_{ijs} \tag{19}$$

subject to

$$\sum_{(j, i) \in A} Y_{jis} - \sum_{(i, j) \in A} Y_{ijs} = b_j \quad \forall j \in V \setminus \{u, v\}, s \in S \tag{20}$$

$$\sum_{(j, i) \in A} Y_{jis} \leq (1 - a_{js})k_j X_j \quad \forall j \in V_0, s \in S \tag{21}$$

$$X_j \in \{0, 1\} \quad \forall j \in V_0 \tag{22}$$

$$Y_{ijs} \geq 0 \quad \forall (i, j) \in A, s \in S \tag{23}$$

The objective function computes the fixed cost and expected flow costs. Constraints (20) are the flow-balance constraints for the nondummy nodes; they require the net flow for node j (flow out minus flow in) to equal the node's deficit b_j in each scenario. Constraints (21) enforce the node capacities and prevent any flow emanating from a node j that has not been opened ($X_j = 0$) or has failed ($a_{js} = 1$). Taken together with (20), these constraints are sufficient to ensure that flow is also prevented *into* nodes that are not opened or have failed. Constraints (22) and (23) are integrality and nonnegativity constraints, respectively. Note that in model (19)–(23), no flow restrictions are necessary for the two dummy nodes. The minimization nature of the objective function guarantees that the demand at each sink node is supplied from regular source nodes whenever this is possible. Only if the node disruption is such to prevent some demand node i from being fully supplied will there be a positive

flow on the link (u, i) at the cost $d_{ui} = \theta_i$. Similarly, only excess supply that cannot reach a sink node will be routed to the dummy sink.

This formulation is similar to the model introduced by Santoso et al. [78]. Their model is intended for network design under demand uncertainty, while ours considers supply uncertainty, though the two approaches are quite similar. To avoid enumerating all possible scenarios, Santoso et al. [78] use SAA. A similar approach is called for to solve (RNDP) because, as in the scenario-based models in §3.2.1, if each node can fail independently, we have $2^{|V_0|}$ scenarios.

A scenario-based model for the design of failure-prone multicommodity networks is discussed in Garg and Smith [38]. However, the model in Garg and Smith [38] does not consider the expected costs of routing the commodities through the network. Rather, it determines the minimum-cost set of arcs to be constructed so that the resulting network continues to support a multicommodity flow under any of a given set of failure scenarios. Only a restricted set of failure scenarios is considered, in which each scenario consists of the concurrent failure of multiple arcs. Garg and Smith [38] also discuss several algorithmic implementations of Benders decomposition to solve this problem efficiently.

3.3.2. Worst-Case Cost. One can modify (RNDP) to minimize the worst-case cost rather than the expected cost in a manner analogous to the approach taken in §3.2.2.

$$\text{minimize } U \tag{24}$$

$$\text{subject to } \sum_{i \in V_0} f_i X_i + \sum_{(i,j) \in A} d_{ij} Y_{ijs} \leq U \quad \forall s \in S \tag{25}$$

$$(20)-(23)$$

Similarly, one could minimize the expected cost subject to a constraint on the cost in any scenario, as proposed above. Bundschuh et al. [15] take a similar approach in a supply chain network design model (with open/close decisions on arcs). They assume that suppliers can fail randomly. They consider two performance measures, which they call *reliability* and *robustness*. The reliability of the system is the probability that all suppliers are operable, while robustness refers to the ability of the supply chain to maintain a given level of output after a failure. The latter measure is perhaps a more reasonable goal because adding new suppliers increases the probability that one or more will fail and, hence, *decreases* the system's "reliability." They present models for minimizing the fixed and (nonfailure) transportation costs subject to constraints on reliability, robustness, or both. Their computational results support the claim made by Snyder and Daskin [92, 94] and others that large improvements in reliability can often be attained with small increases in cost.

4. Fortification Models

4.1. Introduction

Computational studies of the models discussed in the previous sections demonstrate that the impact of facility disruptions can be mitigated by the initial design of a system. However, redesigning an entire system is not always reasonable given the potentially large expense involved with relocating facilities, changing suppliers, or reconfiguring networked systems. As an alternative, the reliability of existing infrastructure can be enhanced through efficient investments in protection and security measures. In light of recent world events, the identification of cost-effective protection strategies has been widely perceived as an urgent priority that demands not only greater public policy support (Sternberg and Lee [97]), but also the development of structured and analytical approaches (Jüttner et al. [52]). Planning for facility protection, in fact, is an enormous financial and logistical challenge if one considers the complexity of today's logistics systems, the interdependencies among critical infrastructures, the variety of threats and hazards, and the prohibitive costs involved in securing large

numbers of facilities. Despite the acknowledged need for analytical models able to capture these complexities, the study of mathematical models for allocation of protection resources is still in its infancy. The few *fortification models* that have been proposed in the literature are discussed in this section, together with possible extensions and variations.

4.2. Facility Location Models

Location models that explicitly address the issue of optimizing facility protection assume the existence of a supply system with P operating facilities. Facilities are susceptible to deliberate sabotage or accidental failures, unless protective measures are taken to prevent their disruption. Given limited protection resources, the models aim to identify the subset of facilities to protect to minimize efficiency losses due to intentional or accidental disruptions. Typical measures of efficiency are distance traveled, transportation cost, or captured demand.

4.2.1. Expected Cost Models. In this section, we present the P -median fortification problem (PMFP) (Scaparra [79]). This model builds on the well-known P -median problem (Hakimi [43, 44]). It assumes that the P facilities in the system have unlimited capacity and that the system users receive service from their nearest facility. As in the design model RFLP, each facility may fail or be disrupted with a fixed probability q . A disrupted facility becomes inoperable, so that the customers currently served by it must be reassigned to their closest nondisrupted facility. Limited fortification resources are available to protect Q of the P facilities. A protected facility becomes immune to disruption. The PMFP identifies the fortification strategy that minimizes the expected transportation costs.

The model definition builds on the notation used in the previous sections, with the exception that J now denotes the set of existing, rather than potential, facilities. Additionally, let i_k denote the k th closest facility to customer i , and let d_i^k be the expected transportation cost between customer i and its closest operational facility, given that the $k - 1$ closest facilities to i are not protected, and the k th closest facility to i is protected. These expected costs can be calculated as follows.

$$d_i^k = \sum_{j=1}^{k-1} q^{j-1}(1-q)d_{ii_j} + q^{k-1}d_{ii_k} \tag{26}$$

The PMFP uses two sets of decision variables:

$$Z_j = \begin{cases} 1, & \text{if facility } j \text{ is fortified,} \\ 0, & \text{otherwise} \end{cases}$$

$$W_{ik} = \begin{cases} 1, & \text{if the } k - 1 \text{ closest facilities to customer } i \text{ are not protected but the } k \text{th} \\ & \text{closest facility is,} \\ 0, & \text{otherwise.} \end{cases}$$

Then PMFP can be formulated as the following mixed integer program.

$$\text{(PMFP) minimize } \sum_{i \in I} \sum_{k=1}^{P-Q+1} h_i d_i^k W_{ik} \tag{27}$$

subject to

$$\sum_{k=1}^{P-Q+1} W_{ik} = 1 \quad \forall i \in I, \tag{28}$$

$$W_{ik} \leq Z_{i_k} \quad \forall i \in I, k = 1, \dots, P - Q + 1 \tag{29}$$

$$W_{ik} \leq 1 - Z_{i_{k-1}} \quad \forall i \in I, k = 2, \dots, P - Q + 1 \quad (30)$$

$$\sum_{j \in J} Z_j = Q \quad (31)$$

$$W_{ik} \in \{0, 1\} \quad \forall i \in I, k = 1, \dots, P - Q + 1 \quad (32)$$

$$Z_j \in \{0, 1\} \quad \forall j \in J \quad (33)$$

The objective function (27) minimizes the weighted sum of expected transportation costs. Note that the expected costs d_i^k and the variables W_{ik} need only be defined for values of k between 1 and $P - Q + 1$. In fact, in the worst case, the closest protected facility to customer i is its $(P - Q + 1)$ st-closest facility. This occurs if the Q fortified facilities are the Q furthest facilities from i . If all of the $P - Q$ closest facilities to i fail, customer i is assigned to its $(P - Q + 1)$ st-closest facility. Assignments to facilities that are further than the $(P - Q + 1)$ st-closest facility will never be made in an optimal solution. For each customer i , constraints (28) force exactly one of the $P - Q + 1$ closest facilities to i to be its closest protected facility. The combined use of constraints (29) and (30) ensures that the variable W_{ik} that equals 1 is the one associated with the smallest value of k such that the k th closest facility to i is protected. Constraint (31) specifies that only Q facilities can be protected. Finally, constraints (32) and (33) represent the integrality requirements of the decision variables.

The PMFP is an integer programming model and can be solved with general purpose mixed-integer programming software. Possible extensions of the model include the cases in which facilities have different failure probabilities and fortification only reduces, but does not eliminate, the probability of failure. Unfortunately, (PMFP) cannot be easily adjusted to handle capacity restrictions. As for the design version of the problem, if the system facilities have limited capacities, explicit scenarios must be used to model possible disruption patterns. The capacitated version of (PMFP) can be formulated in an analogous way to the scenario-based model (CRFLP) discussed in §3.2.1.

$$\text{(CPMFP)} \quad \text{minimize} \quad \sum_{s \in S} q_s \sum_{i \in I} \sum_{j \in J} h_i d_{ij} Y_{ijs} \quad (34)$$

subject to

$$\sum_{j \in J} Y_{ijs} = 1 \quad \forall i \in I, s \in S \quad (35)$$

$$\sum_{i \in I} h_i Y_{ijs} \leq (1 - a_{js})b_j + a_{js}b_j Z_j \quad \forall j \in J, s \in S \quad (36)$$

$$\sum_{j \in J} Z_j = Q \quad (37)$$

$$X_j \in \{0, 1\} \quad \forall j \in J \quad (38)$$

$$Y_{ijs} \in \{0, 1\} \quad \forall i \in I, j \in J, s \in S \quad (39)$$

(CPMFP) uses the same parameters a_{js} and set S as (CRFLP) to model different scenarios. It also assumes that the set of existing facilities J is augmented with the unlimited-capacity emergency facility u . CPMFP differs from CRFLP only in a few aspects: No decisions must be made in terms of facility location, so the fixed cost for locating facilities are not included in the objective; the capacity constraints (36) must reflect that if a facility j is protected ($Z_j = 1$), then that facility remains operable (and can supply b_j units of demand) even in those scenarios s that assume its failure ($a_{js} = 1$). Finally, constraint (37) must be added to fix the number of possible fortifications.

Note that in both models (PMFP) and (CPMFP), the cardinality constraints (31) and (37) can be replaced by more general resource constraints, to handle the problem in which

each facility requires a different amount of protection resources and there is a limit on the total resources available for fortification. Alternately, one could incorporate this cost into the objective function and omit the budget constraint. The difference between these two approaches is analogous to that between the P -median problem and the UFLP.

4.2.2. Worst-Case Cost Models. When modeling protection efforts, it is crucial to account for hazards to which a facility may be exposed. It is evident that protecting against intentional attacks is fundamentally different from protecting against acts of nature. Whereas nature hits at random and does not adjust its behavior to circumvent security measures, an intelligent adversary may adjust its offensive strategy depending on which facilities have been protected, for example, by hitting different targets. The expected cost models discussed in §4.2.1 do not take into account the behavior of adversaries and are, therefore, more suitable to model situations in which natural and accidental failures are a major concern. The models in this section have been developed to identify cost-effective protection strategies against malicious attackers.

A natural way of looking at fortification problems involving intelligent adversaries is within the framework of a leader-follower or Stackelberg game [96], in which the entity responsible for coordinating the fortification activity, or *defender*, is the leader and the attacker, or *interdictor*, is the follower. Stackelberg games can be expressed mathematically as bilevel programming problems (Dempe [31]): The upper-level problem involves decisions to determine which facilities to harden, whereas the lower-level problem entails the interdictor's response of which unprotected facilities to attack to inflict maximum harm. Even if in practice we cannot assume that the attacker is always able to identify the best attacking strategy, the assumption that the interdictor attacks in an optimal way is used as a tool to model worst-case scenarios and estimate worst-case losses in response to any given fortification strategy.

The worst-case cost version of PMFP was formulated as a bilevel program by Scaparra and Church [82]. The model, called the R -interdiction median model with fortification (RIMF), assumes that the system defender has resources to protect Q facilities, whereas the interdictor has resources to attack R facilities, with $Q + R < P$. In addition to the fortification variables Z_j defined in §4.2.1, the RIMF uses the following interdiction and assignment variables:

$$S_j = \begin{cases} 1, & \text{if facility } j \text{ is interdicted,} \\ 0, & \text{otherwise} \end{cases}$$

$$Y_{ij} = \begin{cases} 1, & \text{if customer } i \text{ is assigned to facility } j \text{ after interdiction,} \\ 0, & \text{otherwise.} \end{cases}$$

Additionally, the formulation uses the set $T_{ij} = \{k \in J \mid d_{ik} > d_{ij}\}$ defined for each customer i and facility j . T_{ij} represents the set of existing sites (not including j) that are farther than j is from demand i . The RIMF can then be stated mathematically as follows.

$$\text{(RIMF) minimize } H(Z) \tag{40}$$

subject to

$$\sum_{j \in J} Z_j = Q \tag{41}$$

$$Z_j \in \{0, 1\} \quad \forall j \in J, \tag{42}$$

where

$$H(Z) = \text{maximize } \sum_{i \in I} \sum_{j \in J} h_i d_{ij} Y_{ij} \tag{43}$$

$$\sum_{j \in J} Y_{ij} = 1 \quad \forall i \in I \tag{44}$$

$$\sum_{j \in J} S_j = R \tag{45}$$

$$\sum_{h \in T_{ij}} Y_{ih} \leq S_j \quad \forall i \in I, j \in J \tag{46}$$

$$S_j \leq 1 - Z_j \quad \forall j \in J \tag{47}$$

$$S_j \in \{0, 1\} \quad \forall j \in J \tag{48}$$

$$Y_{ij} \in \{0, 1\} \quad \forall i \in I, j \in J \tag{49}$$

In the above bilevel formulation, the leader allocates exactly Q fortification resources (41) to minimize the highest possible level of weighted distances or costs, H , (40) deriving from the loss of R of the P facilities. That H represents worst-case losses after the interdiction of R facilities is enforced by the follower problem, whose objective involves maximizing the weighted distances or service costs (43). In the lower-level interdiction problem (RIM; Church et al. [21]), constraints (44) state that each demand point must be assigned to a facility after interdiction. Constraint (45) specifies that only R facilities can be interdicted. Constraint (46) maintains that each customer must be assigned to its closest open facility after interdiction. More specifically, these constraints state that if a given facility j is not interdicted ($S_j = 0$), a customer i cannot be served by a facility further than j from i . Constraints (47) link the upper- and lower-level problems by preventing the interdiction of any protected facility. Finally, constraints (42), (48), and (49) represent the integrality requirements for the fortification, interdiction, and assignment variables, respectively. Note that the binary restrictions for the Y_{ij} variables can be relaxed, because an optimal solution with fractional Y_{ij} variables only occurs when there is a distance tie between two nondisrupted closest facilities to customer i . Such cases, although interesting, do not affect the optimality of the solution.

Church and Scaparra [18] and Scaparra and Church [81] demonstrate that it is possible to formulate (RIMF) as a single-level program and discuss two different single-level formulations. However, both formulations require the explicit enumeration of all possible interdiction scenarios and, consequently, their applicability is limited to problem instances of modest size. A more efficient way of solving (RIMF) is through the implicit enumeration scheme proposed by Scaparra and Church [82] and tailored to the bilevel structure of the problem.

A stochastic version of (RIMF), in which an attempted attack on a facility is successful only with a given probability, can be obtained by replacing the lower-level interdiction model (43)–(49) with the probabilistic R -interdiction median model introduced by Church and Scaparra [19].

Different variants of the RIMF model, aiming at capturing additional levels of complexity, are currently under investigation. Ongoing studies focus, for example, on the development of models and solution approaches for the capacitated version of the RIMF.

The RIMF assumes that at most R facilities can be attacked. Given the large degree of uncertainty characterizing the extent of man-made and terrorist attacks, this assumption should be relaxed to capture additional realism. An extension of (RIMF) that includes random numbers of possible losses as well as theoretical results to solve this expected loss version to optimality are currently under development.

Finally, bilevel fortification models similar to (RIMF) can be developed for protecting facilities in supply systems with different service protocols and efficiency measures. For example, in emergency service and supply systems, the effects of disruption may be better measured in terms of the reduction in operational response capability. In these problem settings, the most disruptive loss of R facilities would be the one causing the maximal drop in user

demand that can be supplied within a given time or distance threshold. This problem can be modeled by replacing the interdiction model (43)–(49) with the R -interdiction covering problem introduced by Church et al. [21] and by minimizing, instead of maximizing, the upper-level objective function H , which now represents the worst-case demand coverage decrease after interdiction.

4.3. Network Design Models

The literature dealing with the disruption of existing networked systems has primarily focused on the analysis of risk and vulnerabilities through the development of interdiction models. Interdiction models have been used by several authors to identify the most critical components of a system, i.e., those nodes or linkages that, if disabled, cause the greatest disruption to the flow of services and goods through the network. A variety of models, which differ in terms of objectives and underlying network structures, have been proposed in the interdiction literature. For example, the effect of interdiction on the maximum flow through a network is studied by Wollmer [105] and Wood [106]. Israeli and Wood [48] analyze the impact of link removals on the shortest path length between nodes. Lim and Smith [61] treat the multicommodity version of the shortest path problem, with the objective of assessing shipment revenue reductions due to arc interdictions. A review of interdiction models is provided by Church et al. [21].

Whereas interdiction models can help reveal potential weaknesses in a system, they do not explicitly address the issue of optimizing security. Scaparra and Cappanera [80] demonstrate that securing those network components that are identified as critical in an optimal interdiction solution will not necessarily provide the most cost-effective protection against disruptions. Optimal interdiction is a function of what is fortified, so it is important to capture this interdependency within a modeling framework. The models detailed in the next section explicitly addressed the issue of fortification in networked systems.

4.3.1. Expected Cost. In this section, we present the reliable network fortification problem (RNFP), which can be seen as the protection counterpart of the RNDP discussed in §3.3.1. The problem is formulated below by using the same notation as in §3.3.1 and the fortification variables $Z_j = 1$ if node j is fortified, and $Z_j = 0$ otherwise.

$$(RNFP) \quad \text{minimize} \quad \sum_{s \in S} q_s \sum_{(i,j) \in A} d_{ij} Y_{ijs} \tag{50}$$

subject to

$$\sum_{(j,i) \in A} Y_{jis} - \sum_{(i,j) \in A} Y_{ijs} = b_j \quad \forall j \in V \setminus \{u, v\}, s \in S \tag{51}$$

$$\sum_{(j,i) \in A} Y_{jis} \leq (1 - a_{js})k_j + a_{js}k_j Z_j \quad \forall j \in V_0, s \in S \tag{52}$$

$$\sum_{j \in J} Z_j = Q \tag{53}$$

$$Z_j \in \{0, 1\} \quad \forall j \in V_0 \tag{54}$$

$$Y_{ijs} \geq 0 \quad \forall (i, j) \in A, s \in S \tag{55}$$

The general structure of the RNFP and the meaning of most of its components are as in the RNDP. A difference worth noting is that now the capacity constraints (52) maintain that each fortified node preserves its original capacity in every failure scenario.

The RNFP can be easily modified to handle the problem in which fortification does not completely prevent node failures but only reduces the impact of disruptions. As an example, we can assume that a protected node only retains part of its capacity in case of failure and

that the level of capacity that can be secured depends on the amount of protective resources invested on that node. To model this variation, we denote by f_j the fortification cost incurred to preserve one unit of capacity at node j and by B the total protection budget available. Also, we define the continuous decision variables T_j as the level of capacity that is secured at node j (with $0 \leq T_j \leq k_j$). RNFP can be reformulated by replacing the capacity constraints (52) and the cardinality constraints (53) with the following two sets of constraints:

$$\sum_{(j,i) \in A} Y_{jis} \leq (1 - a_{js})k_j + a_{js}T_j \quad \forall j \in V_0, s \in S \tag{56}$$

and

$$\sum_{j \in J} f_j T_j \leq B. \tag{57}$$

4.3.2. Worst-Case Cost. The concept of protection against worst-case losses for network models has been briefly discussed by Brown et al. [14] and Salmeron et al. [77]. The difficulty in addressing this kind of problem is that their mathematical representation requires building tri-level optimization models, to represent fortification, interdiction, and network flow decisions. Multilevel optimization problems are not amenable to solution by standard mixed integer programming methodologies, and no universal algorithm exists for their solutions. To the best of our knowledge, the first attempt at modeling and solving network problems involving protection issues was undertaken by Scaparra and Capanera [80], who discuss two different models: In the first model, optimal fortification strategies are identified to thwart as much as possible the action of an opponent who tries to disrupt the supply task from a supply node to a demand node by disabling or interdicting network linkages. This model is referred to as the shortest path interdiction problem with fortification (SPIF). In the second model, the aim is to fortify network components so as to maximize the flow of goods and services that can be routed through a supply network after a worst-case disruption of some of the network nodes or linkages. This model is referred to as the maximum flow interdiction problem with fortification (MFIF). The two multilevel models incorporate in the lower level the interdiction models described by Israeli and Wood [48] and by Wood [106], respectively.

In both models, there is a supply node o and a demand node d . Additionally, in the SPIF, each arc (i, j) has a penalty of p_{ij} associated with it that represents the cost increase to ship flow through it if the arc is interdicted. (The complete loss of an arc can be captured in the model by choosing p_{ij} sufficiently large.) In the MFIF, each arc has a penalty r_{ij} representing the percentage capacity reduction of the arc deriving from interdiction. (If $r_{ij} = 100\%$, then an interdicted arc (i, j) is completely destroyed.) The remaining notation used by the two models is the same as in §§3.3.1 and 4.3.1.

Note that in both models, it is assumed that the critical components that can be interdicted and protected are the network linkages. However, it is easy to prove that problems in which the critical components are the nodes can be reduced to critical arc models by opportunely augmenting the underlying graph (Corley and Chang [23]). Hence, we describe the more-general case of arc protection and interdiction.

The three-level SPIF can be formulated as follows.

$$\text{(SPIF)} \quad \min_{Z \in F} \max_{S \in D} \min_Y \sum_{(i,j) \in A} (d_{ij} + p_{ij}S_{ij})Y_{ij} \tag{58}$$

subject to

$$\sum_{(j,i) \in A} Y_{ji} - \sum_{(i,j) \in A} Y_{ij} = b_j \quad \forall j \in V \tag{59}$$

$$S_{ij} \leq 1 - Z_{ij} \quad \forall (i, j) \in A \tag{60}$$

$$Y_{ij} \geq 0 \quad \forall (i, j) \in A \tag{61}$$

where $F = \{Z \in \{0, 1\}^n \mid \sum_{(i,j) \in A} Z_{ij} = Q\}$ and $D = \{S \in \{0, 1\}^n \mid \sum_{(i,j) \in A} S_{ij} = R\}$. Also, as in standard shortest path problems, we define $b_o = 1, b_d = -1$, and $b_j = 0$ for all the other nodes j in V . The objective function (58) computes the minimum-cost path after the worst-case interdiction of R unprotected facilities. This cost includes the penalties associated with interdicted arcs. Protected arcs cannot be interdicted (60).

The MFIF model can be formulated in a similar way as follows.

$$\text{(MFIF)} \quad \max_{z \in F} \min_{s \in D} \max_{Y \geq 0} W \quad (62)$$

subject to

$$\sum_{(j,i) \in A} Y_{ji} - \sum_{(i,j) \in A} Y_{ij} = W \quad j = o \quad (63)$$

$$\sum_{(j,i) \in A} Y_{ji} - \sum_{(i,j) \in A} Y_{ij} = 0 \quad \forall j \in V \setminus \{o, d\} \quad (64)$$

$$\sum_{(j,i) \in A} Y_{ji} - \sum_{(i,j) \in A} Y_{ij} = -W \quad j = d \quad (65)$$

$$Y_{ij} \leq k_{ij}(1 - r_{ij}S_{ij}) \quad \forall (i, j) \in A \quad (66)$$

(60)–(61)

In (MFIF), the objective (62) is to maximize the total flow W through the network after the worst-case interdiction of the capacities of R arcs. Capacity reductions due to interdiction are calculated in (66). Constraints (63)–(65) are standard flow conservation constraints for maximum-flow problems.

The two three-level programs (SPIF) and (MFIF) can be reduced to bilevel programs by taking the dual of the inner network flow problems. Scaparra and Cappanera [80] show how the resulting bilevel problem can be solved efficiently through an implicit enumeration scheme that incorporates network optimization techniques. The authors also show that optimal fortification strategies can be identified for relatively large networks (hundreds of nodes and arcs) in reasonable computational time and that significant efficiency gains (in terms of path costs or flow capacities) can be achieved even with modest fortification resources.

Model (MFIF) can be easily modified to handle multiple sources and multiple destinations. Also, a three-level model can be built along the same lines as (SPIF) and (MFIF) for multicommodity flow problems. For example, by embedding the interdiction model proposed in Lim and Smith [61] in the three-level framework, it is possible to identify optimal fortification strategies for maximizing the profit that can be obtained by shipping commodities across a network, while taking into account worst-case disruptions.

5. Conclusions

In this tutorial, we have attempted to illustrate the wide range of strategic planning models available for designing supply chain networks under the threat of disruptions. A planner's choice of model will depend on a number of factors, including the type of network under consideration, the status of existing facilities in the network, the firm's risk preference, and the resources available for constructing, fortifying, and operating facilities.

We believe that several promising avenues exist for future research in this field. First, the models we discussed in this tutorial tend to be much more difficult to solve than their reliable-supply counterparts—most have significantly more decision variables, many have additional hard constraints, and some have multiple objectives. For these models to be implemented broadly in practice, better solution methods are required.

The models presented above consider the cost of reassigning customers or rerouting flow after a disruption. However, other potential repercussions should be modeled. For example,

firms may face costs associated with destroyed inventory, reconstruction of disrupted facilities, and customer attrition (if the disruption does not affect the firm's competitors). In addition, the competitive environment in which a firm operates may significantly affect the decisions the firm makes with respect to risk mitigation. For many firms, the key objective may be to ensure that their post-disruption situation is no worse than that of their competitors. Embedding these objectives in a game-theoretic environment is another important extension.

Finally, most of the existing models for reliable supply chain network design use some variation of a minimum-cost objective. Such objectives are most applicable for problems involving the distribution of physical goods, primarily in the private sector. However, reliability is critical in the public sector as well, for the location of emergency services, post-disaster supplies, and so on. In these cases, cost is less important than proximity, suggesting that coverage objectives may be warranted. The application of such objectives to reliable facility location and network design problems will enhance the richness, variety, and applicability of these models.

Acknowledgments

The authors gratefully acknowledge financial support from EPSRC (Ref. 320 21095), the Higher Education Funding Council for England (HEFCE), and the National Science Foundation (Grant DMI-0522725). The authors also thank Michael Johnson for his feedback on earlier drafts of this tutorial.

References

- [1] Antonio Arreola-Risa and Gregory A. DeCroix. Inventory management under random supply disruptions and partial backorders. *Naval Research Logistics* 45:687–703, 1998.
- [2] M. L. Balinski. Integer programming: Methods, uses, computation. *Management Science* 12(3):253–313, 1965.
- [3] Alexei Barrionuevo and Claudia H. Deutsch. A distribution system brought to its knees. *New York Times* (Sept. 1) C1, 2005.
- [4] R. Beach, A. P. Muhlemann, D. H. R. Price, A. Paterson, and J. A. Sharp. A review of manufacturing flexibility. *European Journal of Operational Research* 122:41–57, 2000.
- [5] Emre Berk and Antonio Arreola-Risa. Note on “Future supply uncertainty in EOQ models.” *Naval Research Logistics* 41:129–132, 1994.
- [6] Oded Berman and Dimitri Krass. Facility location problems with stochastic demands and congestion. Zvi Drezner and H. W. Hamacher, eds. *Facility Location: Applications and Theory*. Springer-Verlag, New York, 331–373, 2002.
- [7] O. Berman, M. J. Hodgson, and D. Krass. Flow-interception problems. Zvi Drezner, ed. *Facility Location: A Survey of Applications and Methods*. Springer Series in Operations Research, Springer, New York, 389–426, 1995.
- [8] Oded Berman, Dmitry Krass, and Mozart B. C. Menezes. Facility reliability issues in network p -median problems: Strategic centralization and colocation effects. *Operations Research*. Forthcoming, 2005.
- [9] Oded Berman, Dmitry Krass, and Mozart B. C. Menezes. MiniSum with imperfect information: Trading off quantity for reliability of locations. Working paper, Rotman School of Management, University of Toronto, Toronto, ON, Canada, 2005.
- [10] Oded Berman, Richard C. Larson, and Samuel S. Chiu. Optimal server location on a network operating as an $M/G/1$ queue. *Operations Research* 33(4):746–771, 1985.
- [11] D. E. Bienstock, E. F. Brickell, and C. L. Monma. On the structure of minimum-weight k -connected spanning networks. *SIAM Journal on Discrete Mathematics* 3:320–329, 1990.
- [12] E. K. Bish, A. Muriel, and S. Biller. Managing flexible capacity in a make-to-order environment. *Management Science* 51(2):167–180, 2005.
- [13] Ken Brack. Ripple effect from GM strike build. *Industrial Distribution* 87(8):19, 1998.
- [14] G. G. Brown, W. M. Carlyle, J. Salmerón, and K. Wood. Analyzing the vulnerability of critical infrastructure to attack and planning defenses. H. J. Greenberg, ed., *Tutorials in Operations Research*. INFORMS, Hanover, MD, 102–123, 2005.

- [15] Markus Bundschuh, Diego Klabjan, and Deborah L. Thurston. Modeling robust and reliable supply chains. Working paper, University of Illinois, Urbana-Champaign, IL, 2003.
- [16] R. D. Carr, H. J. Greenberg, W. E. Hart, G. Konjevod, E. Lauer, H. Lin, T. Morrison, and C. A. Phillips. Robust optimization of contaminant sensor placement for community water systems. *Mathematical Programming* 107:337–356, 2005.
- [17] Richard Church and Charles ReVelle. The maximal covering location problem. *Papers of the Regional Science Association* 32:101–118, 1974.
- [18] Richard L. Church and Maria P. Scaparra. Protecting critical assets: The r -interdiction median problem with fortification. *Geographical Analysis*. Forthcoming. 2005.
- [19] R. L. Church and M. P. Scaparra. Analysis of facility systems' reliability when subject to attack or a natural disaster. *Reliability and Vulnerability in Critical Infrastructure: A Quantitative Geographic Perspective*. A. T. Murray and T. H. Grubescic, eds. Springer-Verlag, New York, 2006.
- [20] R. L. Church, M. P. Scaparra, and J. R. O'Hanley. Optimizing passive protection in facility systems. Working paper, ISOLDE X, Spain, 2005.
- [21] Richard L. Church, Maria P. Scaparra, and Richard S. Middleton. Identifying critical infrastructure: The median and covering facility interdiction problems. *Annals of the Association of American Geographers* 94(3):491–502, 2004.
- [22] C. Colbourn. *The Combinatorics of Network Reliability*. Oxford University Press, New York, 1987.
- [23] H. W. Corley and H. Chang. Finding the most vital nodes in a flow network. *Management Science* 21(3):362–364, 1974.
- [24] Mark S. Daskin. Application of an expected covering model to emergency medical service system design. *Decision Sciences* 13:416–439, 1982.
- [25] Mark S. Daskin. A maximum expected covering location model: Formulation, properties and heuristic solution. *Transportation Science* 17(1):48–70, 1983.
- [26] Mark S. Daskin. *Network and Discrete Location: Models, Algorithms, and Applications*. Wiley, New York, 1995.
- [27] Mark S. Daskin, Collette R. Coullard, and Zuo-Jun Max Shen. An inventory-location model: Formulation, solution algorithm and computational results. *Annals of Operations Research* 110:83–106, 2002.
- [28] M. S. Daskin, K. Hogan, and C. ReVelle. Integration of multiple, excess, backup, and expected covering models. *Environment and Planning B* 15(1):15–35, 1988.
- [29] Mark S. Daskin, Lawrence V. Snyder, and Rosemary T. Berger. Facility location in supply chain design. A. Langevin and D. Riopel, eds., *Logistics Systems: Design and Operation*. Springer, New York, 39–66, 2005.
- [30] A. de Toni and S. Tonchia. Manufacturing flexibility: A literature review. *International Journal of Production Research* 36(6):1587–1617, 1998.
- [31] S. Dempe. *Foundations of Bilevel Programming*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2002.
- [32] Z. Drezner. Heuristic solution methods for two location problems with unreliable facilities. *Journal of the Operational Research Society* 38(6):509–514, 1987.
- [33] Zvi Drezner, ed. *Facility Location: A Survey of Applications and Methods*. Springer-Verlag, New York, 1995.
- [34] H. A. Eiselt, Michel Gendreau, and Gilbert Laporte. Location of facilities on a network subject to a single-edge failure. *Networks* 22:231–246, 1992.
- [35] D. Elkins, R. B. Handfield, J. Blackhurst, and C. W. Craighead. 18 ways to guard against disruption. *Supply Chain Management Review* 9(1):46–53, 2005.
- [36] B. Fortz and M. Labbe. Polyhedral results for two-connected networks with bounded rings. *Mathematical Programming Series A* 93:27–54, 2002.
- [37] Justin Fox. A meditation on risk. *Fortune* 152(7):50–62, 2005.
- [38] M. Garg and J. C. Smith. Models and algorithms for the design of survivable multicommodity flow networks with general failure scenarios. *Omega*. Forthcoming. 2006.
- [39] M. Gendreau, G. Laporte, and I. Parent. Heuristics for the location of inspection stations on a network. *Naval Research Logistics* 47:287–303, 2000.
- [40] Stephen C. Graves and Brian T. Tomlin. Process flexibility in supply chains. *Management Science* 49(7):907–919, 2003.

- [41] M. Grötschel, C. L. Monma, and M. Stoer. Polyhedral and computational investigations for designing communication networks with high survivability requirements. *Operations Research* 43(6):1012–1024, 1995.
- [42] Diwakar Gupta. The (Q, r) inventory system with an unreliable supplier. *INFOR* 34(2):59–76, 1996.
- [43] S. L. Hakimi. Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations Research* 12(3):450–459, 1964.
- [44] S. L. Hakimi. Optimum distribution of switching centers in a communication network and some related graph theoretic problems. *Operations Research* 13(3):462–475, 1965.
- [45] Julia L. Higle. Stochastic programming: Optimization when uncertainty matters. *Tutorials in Operations Research*. INFORMS, Hanover, MD, 30–53, 2005.
- [46] M. J. Hodgson, K. E. Rosing, and J. Zhang. Locating vehicle inspection stations to protect a transportation network. *Geographical Analysis* 28:299–314, 1996.
- [47] Wallace J. Hopp and Zigeng Yin. Protecting supply chain networks against catastrophic failures. Working paper, Northwestern University, Evanston, IL, 2006.
- [48] E. Israeli and R. K. Wood. Shortest-path network interdiction. *Networks* 40(2):97–111, 2002.
- [49] J. C. James and S. Salhi. A Tabu Search heuristic for the location of multi-type protection devices on electrical supply tree networks. *Journal of Combinatorial Optimization* 6:81–98, 2002.
- [50] Hyong-Mo Jeon, Lawrence V. Snyder, and Z. J. Max Shen. A location-inventory model with supply disruptions. Working paper, Lehigh University, Bethlehem, PA, 2006.
- [51] William C. Jordan and Stephen C. Graves. Principles on the benefits of manufacturing process flexibility. *Management Science* 41(4):577–594, 1995.
- [52] U. Jüttner, H. Peck, and M. Christopher. Supply chain risk management: Outlining an agenda for future research. *International Journal of Logistics: Research and Applications* 6(4):197–210, 2003.
- [53] Hyoungtae Kim, Jye-Chyi Lu, and Paul H. Kvam. Ordering quantity decisions considering uncertainty in supply-chain logistics operations. Working paper, Georgia Institute of Technology, Atlanta, GA, 2005.
- [54] Anton J. Kleywegt, Alexander Shapiro, and Tito Homem-de-Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* 12(2):479–502, 2001.
- [55] Panagiotis Kouvelis and Gang Yu. *Robust Discrete Optimization and Its Applications*. Kluwer Academic Publishers, Boston, MA, 1997.
- [56] Richard C. Larson. A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers and Operations Research* 1:67–95, 1974.
- [57] Richard C. Larson. Approximating the performance of urban emergency service systems. *Operations Research* 23(5):845–868, 1975.
- [58] Almar Latour. Trial by fire: A blaze in Albuquerque sets off major crisis for cell-phone giants—Nokia handles supply chain shock with aplomb as Ericsson of Sweden gets burned—Was Sisu the difference? *Wall Street Journal* (Jan. 29) A1, 2001.
- [59] S. D. Lee. On solving unreliable planar location problems. *Computers and Operations Research* 28:329–344, 2001.
- [60] Devin Leonard. The only lifeline was the Wal-Mart. *Fortune* 152(7):74–80, 2005.
- [61] C. Lim and J. C. Smith. Algorithms for discrete and continuous multicommodity flow network interdiction problems. *IIE Transactions*. Forthcoming, 2006.
- [62] Jeff Linderoth, Alexander Shapiro, and Stephen Wright. The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research* 142:219–245, 2006.
- [63] Barry C. Lynn. *End of the Line: The Rise and Coming Fall of the Global Corporation*. Doubleday, New York, 2005.
- [64] Esmail Mohebbi. Supply interruptions in a lost-sales inventory system with random lead time. *Computers and Operations Research* 30:411–426, 2003.
- [65] Esmail Mohebbi. A replenishment model for the supply-uncertainty problem. *International Journal of Production Economics* 87(1):25–37, 2004.
- [66] C. L. Monma. Minimum-weight two-connected spanning networks. *Mathematical Programming* 46(2):153–171, 1990.

- [67] C. L. Monma and D. F. Shalcross. Methods for designing communications networks with certain 2-connected survivability constraints. *Operations Research* 37:531–541, 4 1989.
- [68] Jad Mouawad. Katrina’s shock to the system. *New York Times* (Sept. 4) 3.1, 2005.
- [69] J. R. O’Hanley and R. L. Church. Planning for facility-loss: A bilevel decomposition algorithm for the maximum covering location-interdiction problem. Working paper, Oxford University, Oxford, England, 2005.
- [70] Susan Hesse Owen and Mark S. Daskin. Strategic facility location: A review. *European Journal of Operational Research* 111(3):423–447, 1998.
- [71] Mahmut Parlar. Continuous-review inventory problem with random supply interruptions. *European Journal of Operational Research* 99:366–385, 1997.
- [72] M. Parlar and D. Berkin. Future supply uncertainty in EOQ models. *Naval Research Logistics* 38:107–121, 1991.
- [73] Hasan Pirkul. The uncapacitated facility location problem with primary and secondary facility requirements. *IIE Transactions* 21(4):337–348, 1989.
- [74] Reuters. Lumber, coffee prices soar in Katrina’s wake. *Reuters* (Sept. 1) 2005.
- [75] Charles ReVelle and Kathleen Hogan. The maximum availability location problem. *Transportation Science* 23(3):192–200, 1989.
- [76] J. B. Rice and F. Caniato. Building a secure and resilient supply network. *Supply Chain Management Review* 7(5):22–30, 2003.
- [77] J. Salmeron, R. K. Wood, and R. Baldick. Analysis of electric grid security under terrorist threat. *IEEE Transactions on Power Systems* 19(2):905–912, 2004.
- [78] Tjendra Santoso, Shabbir Ahmed, Marc Goetschalckx, and Alexander Shapiro. A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research* 167:96–115, 2005.
- [79] M. P. Scaparra. Optimal resource allocation for facility protection in median systems. Working paper, University of Kent, Canterbury, England, 2006.
- [80] M. P. Scaparra and P. Capanera. Optimizing security investments in transportation and telecommunication networks. INFORMS Annual Meeting, San Francisco, CA, 2005.
- [81] Maria P. Scaparra and Richard L. Church. An optimal approach for the interdiction median problem with fortification. Working Paper 78, Kent Business School, Canterbury, England, UK, 2005.
- [82] Maria P. Scaparra and Richard L. Church. A bilevel mixed integer program for critical infrastructure protection planning. *Computers and Operations Research*. Forthcoming. 2006.
- [83] Alexander Shapiro and Tito Homem-de-Mello. A simulation-based approach to two-stage stochastic programming with recourse. *Mathematical Programming* 81:301–325, 1998.
- [84] Yossi Sheffi. *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage*. MIT Press, Cambridge, MA, 2005.
- [85] Zuo-Jun Max Shen, Collette R. Coullard, and Mark S. Daskin. A joint location-inventory model. *Transportation Science* 37(1):40–55, 2003.
- [86] D. R. Shier. *Network Reliability and Algebraic Structures*. Clarendon Press, Oxford, England, 1991.
- [87] Martin L. Shooman. *Reliability of Computer Systems and Networks: Fault Tolerance, Analysis, and Design*. John Wiley & Sons, New York, 2002.
- [88] Robert L. Simison. GM contains its quarterly loss at \$809 million. *Wall Street Journal* (Oct. 14) A2, 1998.
- [89] Robert L. Simison. GM says strike reduced its earnings by \$2.83 billion in 2nd and 3rd periods. *Wall Street Journal* (Aug. 17) 1, 1998.
- [90] Lawrence V. Snyder. Facility location under uncertainty: A review. *IIE Transactions* 38(7):537–554, 2006.
- [91] Lawrence V. Snyder. A tight approximation for a continuous-review inventory model with supplier disruptions. Working paper, Lehigh University, Bethlehem, PA, 2006.
- [92] Lawrence V. Snyder and Mark S. Daskin. Reliability models for facility location: The expected failure cost case. *Transportation Science* 39(3):400–416, 2005.
- [93] Lawrence V. Snyder and Mark S. Daskin. Stochastic p -robust location problems. *IIE Transactions* 38(11):971–985, 2006.

- [94] Lawrence V. Snyder and Mark S. Daskin. Models for reliable supply chain network design. Alan T. Murray and Tony H. Grubescic, eds. *Reliability and Vulnerability in Critical Infrastructure: A Quantitative Geographic Perspective*. Forthcoming. Springer, New York, 2006.
- [95] Lawrence V. Snyder and Z. Max Shen. Disruptions in multi-echelon supply chains: A simulation study. Working paper, Lehigh University, 2005.
- [96] H. Stackelberg. *The Theory of Market Economy*. Oxford University Press, Oxford, England, 1952.
- [97] E. Sternberg and G. Lee. Meeting the challenge of facility protection for homeland security. *Journal of Homeland Security and Emergency Management* 3(1):1–19, 2006.
- [98] Brian T. Tomlin. The impact of supply-learning on a firm’s sourcing strategy and inventory investment when suppliers are unreliable. Working Paper OTIM-2005-05, Kenan-Flagler Business School, University of North Carolina, Chapel Hill, NC, 2005.
- [99] Brian T. Tomlin. Selecting a disruption-management strategy for short life-cycle products: Diversification, contingent sourcing, and demand management. Working Paper OTIM-2005-09, Kenan-Flagler Business School, University of North Carolina, Chapel Hill, NC, 2005.
- [100] Brian T. Tomlin. On the value of mitigation and contingency strategies for managing supply-chain disruption risks. *Management Science* 52(5):639–657, 2006.
- [101] Brian T. Tomlin and Lawrence V. Snyder. Inventory management with advanced warning of disruptions. Working paper, Lehigh University, Bethlehem, PA, 2006.
- [102] Brian Tomlin and Yimin Wang. On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. Working paper, Kenan-Flagler Business School, University of North Carolina, Chapel Hill, NC, 2004.
- [103] Nursen Ş. Ülker and Lawrence V. Snyder. A model for locating capacitated, unreliable facilities. Working paper, Lehigh University, Bethlehem, PA, 2005.
- [104] Jerry R. Weaver and Richard L. Church. A median location model with nonclosest facility service. *Transportation Science* 19(1):58–74, 1985.
- [105] R. Wollmer. Removing arcs from a network. *Operations Research* 12(6):934–940, 1964.
- [106] R. K. Wood. Deterministic network interdiction. *Mathematical and Computer Modelling* 17(2):1–18, 1993.