Pre-positioning Disaster Response Facilities At Safe Locations: An Evaluation of Deterministic and Stochastic Modeling Approaches

Anurag Verma*
Industrial & Systems Engineering
Texas A&M University

Gary M. Gaukler†
Drucker School of Management
Claremont Graduate University

Forthcoming in Computers & Operations Research, 2015

April 18, 2015

Abstract: Choosing the locations of disaster response facilities for the storage of emergency supplies is critical to the quality of service provided post-occurrence of a large scale emergency like an earthquake. In this paper, we provide two location models that explicitly take into consideration the impact a disaster can have on the disaster response facilities and the population centers in surrounding areas. The first model is a deterministic model that incorporates distance-dependent damages to disaster response facilities and population centers. The second model is a stochastic programming model that extends the first by directly considering the damage intensity as a random variable. For this second model we also develop a novel solution method based on Benders Decomposition that is generalizable to other 2-stage stochastic programming problems. We provide a detailed case study using large-scale emergencies caused by an earthquake in California to demonstrate the performance of these new models. We find that the locations suggested by the stochastic model in this paper significantly reduce the expected cost of providing supplies when one considers the damage a disaster causes to the disaster response facilities and areas near it. We also demonstrate that the cost advantage of the stochastic model over the deterministic model is especially large when only a few facilities can be placed. Thus, the value of the stochastic model is particularly great in realistic, budget-constrained situations.

*anuragverma@tamu.edu
†gary.gaukler@cgu.edu

1 Introduction

To provide responsive and timely service in the event of natural disasters and terrorist attacks, government agencies are developing large disaster response facilities to pre-position emergency supplies (Balcik and Beamon, 2008). For example, in the United States, significant research interest has been generated in the location planning of these facilities after the Centers for Disease Control and Prevention (CDC) were entrusted with the task of establishing the Strategic National Stockpile (SNS). According to the CDC web site (www.bt.cdc.gov):

> Strategic National Stockpile (SNS) has large quantities of medicine and medical supplies to protect the American public if there is a public health emergency (terrorist attack, flu outbreak, earthquake) severe enough to cause local supplies to run out. Once Federal and local authorities agree that the SNS is needed, medicines will be delivered to any state in the U.S. within 12 hours.

This paper focuses on the optimal placement of disaster response facilities like the SNS that will be used to pre-position emergency supplies. Emergency supplies can include food, medicine, potable water, but also medical equipment, generators, tents etc. In deciding on suitable locations for pre-positioning warehouses, decision makers need to consider disasters that may affect large geographical areas, with the potential to devastate entire cities. Earthquakes are a typical example, but other large-scale disasters where damage to surrounding areas originates from an epicenter are also applicable. This disaster class may include floods, large scale fires, and even non-natural events such as terrorist bomb attacks on certain target structures.

Many models for locating facilities for pre-positioning emergency supplies have assumed that facilities are robust and will be functioning even in the wake of a natural disaster (Balcik and Beamon, 2008; Duran et al., 2011). There exist models that consider facilities that might not be always available at their full capacity, for example (Jia et al., 2007; Paul and Batta, 2008; Beraldi et al., 2004). These models assume that a disaster reduces the capacity of a facility by a certain deterministic fraction. Other models that incorporate damage to facilities, such as Rawls and Turnquist (2010); Noyan (2012); Bozorgi-Amiri et al. (2013), use stochastic formulations, but are scenario-based and model damage exogenously to the model. In contrast to this, both models we develop in this paper explicitly adress the uncertainty in the magnitude of damages caused by a large-scale emergency event through the introduction of a distance-damage function.

We compare the performance of these two models in an exemplary case study, using earthquakes in California as the disaster of interest. Our case study reveals that stochastic treatment of damages can have significant impact on the quality of the pre-positioning decision. Our study also shows that the cost advantage of the stochastic model over the deterministic model is especially large when (1) only a few facilities can be placed, and when (2) the uncertainty over the potential damage inflicted by the disaster is high. Thus, the value of the stochastic model is particularly great in realistic, budget-constrained situations, and when the disaster outcome is hard to predict.

To motivate why making location decisions for large-scale emergencies where facilities may fail is different from making location decisions for general facilities, we consider a simple stylized example: suppose there are two cities A and B where population is concentrated, and four potential facility sites - one at A, one at B, and two between A and B as shown in the Figure 1. We will refer to

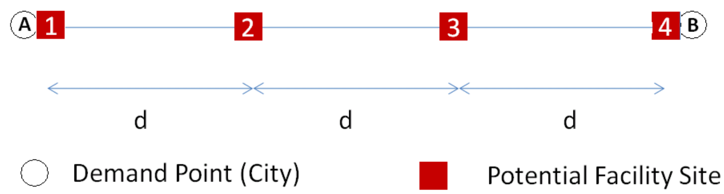

Figure 1: Illustration to show the shortcomings of naive location models for large-scale emergency cases.

the two cities as the demand points. The distances are as marked on the figure, and the chance of a disaster occurring at any of the two cities is the same. For the purpose of exposition, we assume that A and B are high risk areas where disasters might occur, and that the possibility of a disaster occurring at other locations is small enough to be ignored.

Suppose we want to construct two disaster response facilities. Location models such as the traditional k-median model or models developed by Paul and Batta (2008); Beraldi et al. (2004) assume that a reduction in capacity of facilities is unrelated to where the disaster occurs, and would therefore prescribe locating the facilities at sites 1 and 4. However, if a devastating earthquake occurs at city A, most likely facility 1 will be damaged because of its proximity to the disaster and may not be able to satisfy all demands. Aid would have to come from facility 4 which is far away. Similarly, if an earthquake occurs near city B, facility 4 would not be functioning, and aid would have to come all the way from facility 1. This intuitively poor placement decision occurs because these traditional location models assume that facility availability is independent of disaster location, whereas in actuality this is not true.

If one were to condition the functioning of the disaster response facilities on the actual disaster, better solutions might be found. Indeed, the model we present in this paper suggests locating the facilities at sites 2 and 3. When a disaster happens at A, 2 and 3 being relatively far away from the disaster site will still be functioning at a slightly reduced capacity and can combine to send aid. When a disaster happens at B, the same holds true. Locating facilities at sites 2 and 3 saves transportation cost and also reduces response times.

In addition, our paper also addresses other important issues such as the stochastic nature of the damage due to the disaster and the effect of a disaster on multiple cities. We address these issues by explicitly modeling the damage a disaster causes to the cities and facilities in its vicinity as a random variable that is correlated to the location of the disaster via a distance-damage function. Moreover, we provide insight into the impact of the variability of damage intensity on the solution quality through a sensitivity analysis on the coefficient of variation of the random variable describing the damage of the disaster. We also demonstrate the impact that the density of the disaster response facilty network has on the respective solution qualities of the models.

Our modeling approach is based on the intuition that locating a disaster response facility very close to a high risk city or population region may not be optimal as the facility itself might be damaged when needed. This distance-dependence is a reasonable assumption because typically the damage from natural and man-made disasters are highest closest to the primary impact of a disaster such as the epicenter of an earthquake or the track of a hurricane (Schultz et al., 2007).

In this paper, we formulate the distance-dependent large scale emergency pre-positioning model, and we also provide a novel solution algorithm for the stochastic model. This solution algorithm is based on a modification of Benders decomposition, using a greedy heuristic to solve the master problem. To the best of our knowledge, this modification is novel in the literature. Our solution algorithm is formulated for solving the pre-positioning model developed in this paper, but its basic idea should be applicable to a larger class of stochastic location problems. We provide a case study on earthquakes in the state of California to show the performance of our model and to demonstrate the necessity of incorporating the modeling improvements for locating disaster response facilities.

The remainder of this paper is organized as follows. In Section 2, we present an overview of the existing literature. Section 3 provides a new formulation of the pre-positioning problem that considers the effect of a disaster on the facilities and population centers close by, while Section 4 provides effective solution algorithms for solving the model. Section 5 provides a case study of the new models, and Section 6 concludes this paper with a discussion of the contribution of this paper, as well as future research directions.

2 Literature Survey

When examining the literature on facility location under uncertainty, two broad categories of problems stand out: (1) problems where the facilities are more or less constantly in use, like warehouses or fire stations, and (2) problems where facilities come into use after some rare event, such as emergency supply warehouses being used after an earthquake.

In that first category, it is reasonable to have models in which the functioning of a facility is independent of externalities like demand. However, the same does not hold true for the second category. Here the functioning of a facility is coupled with the rare event that causes a demand. In the context of this paper, we define "emergency" as a rare, high-consequence, large-scale event, as opposed to "routine" emergencies such as ambulance or police calls.

The literature on facility location under uncertainty is fairly advanced for the first category of problems that were discussed in the preceding paragraph, as can be seen from (Berman et al., 2003, 2007; Snyder and Daskin, 2005). Berman et al. (2003) discuss the location of facilities whose reliability is dependent on the distance between the facility and the demand point. However, their model was developed for a constant demand class of problems. The chance of providing uninterrupted service goes down as the distance increases. They take into account the uncertainty in roads and transportation links being available and functioning, but they do not take into account the functioning of the facility itself. This is a realistic assumption since these models were designed for a firm providing constant service. However, this assumption makes these models less suitable for emergency facility location. The difference to the emergency facility location problem also becomes evident in their assumption that the probability of a facility not being able to provide service is zero at distance zero, and is a monotonically increasing function of distance. This is in complete disagreement with an emergency scenario, where at distance zero, the facility tends to be severely damaged and will not serve.

Berman et al. (2007) describe a p-median model using a different definition to consider uncertainty in the working of facilities. They define r_j to be the probability that the facility is working at any point of time. Using this setup, they show that co-location of facilities (locating facilities next to each other, or over each other) is a phenomenon observed in such cases. As the failure probability grows, facilities become more and more centralized. Snyder and Daskin (2005) define the reliability k-median problem to incorporate uncertainty into classical facility location models. These models have been discussed keeping a supply chain in mind, and do not apply directly to an emergency facility location problem.

The second category of models (facilities coming into use after a rare event) includes what are typically termed the emergency and disaster logistics applications. The literature on emergency and disaster logistics has experienced extensive growth in recent years. Caunhye et al. (2012) provide a thorough overview of the state of the literature on optimization models in emergency logistics. They identify a total of 64 papers dealing with operations research models for disaster logistics. Caunhye et al. (2012) divide the body of literature into *Facility location based* models, and *Relief distribution and casualty transportation* models. As examples of this category of casualty transportation models, Ben-Tal et al. (2011) use robust optimization to develop emergency response and evacuation transportation plans. Barbarosoglu and Arda (2004) provide a stochastic programming model to plan the transportation for emergency supplies in the event of an earthquake. The model is descriptive with scenarios based on the location and impact of the earthquake. However, this is not a facility location model, but a multi commodity flow model, and the authors are concerned with routing the supplies given the positions of the facilities.

Among Facility location based models, Caunhye et al. (2012) identify papers that focus on location-evacuation problems (where the issue is to determine shelter locations and optimize traffic flow plans), and papers that focus on facility location with relief distribution and stock pre-positioning (where the issue is to determine good locations for storing supplies and other assets in anticipation of a disaster). Our paper falls into that latter category. In the following, we will only discuss the portion of that literature that is closest to our work in the facility location and stock pre-positioning category. We refer the reader to Caunhye et al. (2012) for detailed discussion of extant emergency logistics models in other categories.

Jia et al. (2007) use k-median and k-center models to solve the emergency facility location problem. They account for reduction in service capability through a parameter p_j in their model, which is defined as the fraction of full capacity a facility j is working at in the event of a disaster. Disasters are categorized into disaster scenarios that need to be set up prior to solving the location problem. This external-to-the-model setup of scenarios can quickly become burdensome for large problems. Our model, in contrast, internalizes the scenario creation by directly modeling the damages to a population point or a facility as a function of the proximity to a disaster epicenter. Paul and Batta (2008) present a model for hospital location that is similar to Jia et al. (2007). They also use a factor f_k to denote the fractional capacity of a hospital located at site k in the wake of a disaster. However, f_k is taken to be independent of the disaster here.

Murali et al. (2009) extend the model from Jia et al. (2007) by making the coverage a facility can provide dependent on the distance from the demand point. They argue that the farther a facility is from the demand point, the lesser the chances of providing proper coverage due to damage to the transportation infrastructure.

Duran et al. (2011) present a model to place warehouses to minimize the average of weighted response times for different demand scenarios that might occur in a disaster. Balcik and Beamon (2008) provide a deterministic covering model for determining the locations and capacities of facilities for pre-positioning supplies. Both papers, however, assume that the facilities will keep functioning at their full capacity. Mete and Zabinsky (2010) provide a stochastic model for determining quantities of backup supplies hospitals should store on-site and in warehouses to counter emergencies, but their model does not consider damage to facilities and is not suited for large scale disasters.

Doyen et al. (2012) study a humanitarian relief logistics problem that places both pre- and post-disaster facilities. Pre-disaster facilities are used for pre-positioning relief items, whereas post-disaster facilities are regional and local rescue centers. However, even though Doyen et al. (2012) consider various disaster intensity scenarios, they do not consider the effect of the disaster on the pre-positioned relief items themselves.

Gormez et al. (2011) consider a case study of locating disaster response facilities in the city of Istanbul. The disaster scenario is an earthquake, and facilities are positioned to pre-position relief items as well as carry out post-disaster recovery operations. In contrast to our work, the impact of the earthquake on relief item availability is not considered in their model.

Salmeron and Apte (2010) provide a two-stage stochastic optimization model for pre-disaster budgeting for and positioning of resources. Resources here are warehouses, health care facilities and assets such as transportation vehicles. The first stage determines locations and capacities for these resources, and the second stage covers the operational nature of post-disaster use of assets and means of transportation. Potential disruption to pre-positioned supplies by the disaster is not considered.

Rawls and Turnquist (2010) study a stochastic supply pre-positioning problem where damage to supplies is considered on a scenario basis, similar to Jia et al. (2007). The main difference in model formulation to ours is that supply damage is defined exogenously (through defining what the damage will be at every location, for every disaster scenario), whereas in our model damage to supplies is calculated endogenously through the distance-damage function. Rawls and Turnquist (2011) extend this basic model to include service quality constraints. A constraint on the minimum probability of meeting all relief demand as well as a constraint on the average shipment distance are added. Rawls and Turnquist (2012) extends Rawls and Turnquist (2010) by focusing on short-term time-dependent demands for relief supplies by evacuees at shelter locations. Noyan (2012) extends the Rawls and Turnquist (2010) model by introducing a stochastic formulation based on conditional value-at-risk (in contrast, our stochastic model is expectation-based and risk-neutral). Noyans formulation is also based on Rawls and Turnquists exogenous definition of supply damage. Bozorgi-Amiri et al. (2013) also provide a stochastic programming based pre-positioning model for disaster relief. They provide a case study on planning for earthquake scenarios in Iran. They focus on a multi-objective robust optimization approach, and, like Noyan, treat the damage to supplies as exogenously defined through scenario creation, as opposed to endogenously determined.

Klibi et al. (2013) also present a stochastic programming approach to pre-positioning of relief supplies. Their model assumes multiple natural disasters happening over a discrete time horizon. The model uses a multi-period framework covering deployment, sustainment, recovery, and redeployment. Distribution centers and vendors carry inventory of relief items. Availability of distribution centers and vendors may be affected by the disaster itself: during scenario creation, it is determined which distribution centers and vendors will be able to service demand. Those that are affected by the disaster are taken out of the candidate set. Our model offers a different way of handling supply deterioration and supply availability by allowing for endogenously determined fractional capacities to be available at each center, governed by proximity to the disaster epicenter.

3 Mathematical Formulation

In this section we present two models that incorporate and internalize the assumption of distance-dependent damages to disaster response facilities and population centers. The first is a deterministic model that is used as a baseline benchmark, representing the performance of deterministic large-scale emergency models. The second model is a 2-stage stochastic programming model that directly considers the damage caused by a disaster to facilities and population centers as stochastic. This second model is more realistic, but also presents larger computational challenges while solving for very large instances.

3.1 Deterministic Model

The model presented in this section is considerably different from related models in the literature, because it directly couples damages to population centers and response facilities through a distance-damage function. Besides having the advantage that there is no need for setting up scenarios for damages that are external to the model, the formulation also has several other useful properties: First, the model does not assume that a disaster can only occur on population centers. In our formulation, epicenters can be located anywhere. Secondly, the damage to facilities is explicitly dependent on the location of disasters. Lastly, in case of a disaster, demand on the response facilities is not generated from just one population center, but all population centers depending on their distance from the disaster's epicenter. These modeling choices allow us to more realistically consider disasters that are devastating enough to encompass a large area.

Let I be the set of population centers (demand points) in the region under consideration, F be the set of possible facility locations, and E be the possible epicenters given by a forecasting mechanism (Rundle et al., 2003). In the model provided below, for simplicity and without loss of generality, the demands D_i at demand point $i \in I$ and capacities c_f of a facility $f \in F$ are represented in the same units. These units could represent, for example, the number of people whose needs have to be fulfilled if there is one emergency supply packet stored per person to be served. The probability of a disaster occurring at an epicenter $e \in E$ is denoted by w_e. The quantity $\bar{d}_{if}$ denotes the travel distance from facility f to demand i, whereas d_{ex} denotes the euclidean distance between e and $x \in I \cup F$, and is used to estimate the damage at x due to a disaster at e. The parameter p_{ex} denotes the damage at $x \in I \cup F$ as a fraction of the total demand/capacity. It should be noted that given d_{ex}, we assume that we can obtain p_{ex} by means of a functional relationship $p_{ex}(d_{ex})$. k denotes the upper limit on number of facilities to be built.

The binary variable x_f denotes the decision to open a facility at location $f \in F$. The variable y_{eif} denotes the amount of service provided by facility f to demand point i when a disaster occurs at e. The formulation is given below and is referred to as the *distance-damage* model.

$$\text{Min} \quad \sum_{e\in E}\sum_{i\in I}\sum_{f\in F} w_e \bar{d}_{if} y_{eif} \tag{3.1}$$

$$\text{Subject to} \quad \sum_{f\in F} x_f \leq k \tag{3.2}$$

$$\sum_{i\in I} y_{eif} \leq (1-p_{ef})c_f x_f \qquad \forall e \in E, \forall f \in F \tag{3.3}$$

$$\sum_{f\in F} y_{eif} \geq p_{ei} D_i \qquad \forall e \in E, \forall i \in I \tag{3.4}$$

$$x_f \in \{0,1\} \qquad \forall f \in F \tag{3.5}$$

$$y_{eif} \geq 0 \qquad \forall e \in E, \forall i \in I, \forall f \in F \tag{3.6}$$

The objective function (3.1) is the expected transportation cost over all disaster scenarios assuming the costs are linear in the distance that has to be traveled and the amount of supplies to be shipped. This linearity assumption is common in literature. Constraint (3.2) limits the number of facilities built. Constraints (3.3) limit the total supply a facility can provide to all the demand points that need supplies. Constraints (3.4) make sure all the demands under a scenario are satisfied. Constraint sets (3.5) and (3.6) are binary and non-negativity constraints respectively.

To quantify the extent of damage on facilities and population centers alike, we define the notion of *expected conditional damage*, $p_{ex} \in [0,1]$, at a point $x \in I \cup F$ (population center or facility location) in the event of a disaster at point e (e.g., the epicenter of an earthquake). $p_{ex} = 0$ means that no damage is expected at point x when a disaster occurs at e and $p_{ex} = 1$ means that everything is damaged at x. The exact functional relationship $p_{ex}(d_{ex})$ depends on the circumstances under which the model is being used. For example, in our computational experiments in Section 5, we use a distance-damage function proposed by Howell and Schultz (1975) for earthquakes.

As in the models present in literature, we consider a demand to be generated only when a disaster occurs at one of the possible locations. Hence, a scenario essentially specifies the point where a disaster has occurred and the effects of the disaster on the facilities and population centers nearby through the distance-damage function. Thus, a scenario e is completely defined by the location of the disaster; the demand at each of the cities under the scenario, $p_{ei}D_i, i \in I$; and the impact of the disaster on the disaster response facilities, $p_{ef}c_f, f \in F$. Note that $p_{ei}D_i, i \in I$ and $p_{ef}c_f, f \in F$ are directly computed within the model using the distance-damage function.

Intuitively, the modifications discussed in the beginning of this section are expected to result in significant changes in the locations of facilities, compared to a standard k-median location approach. Recognizing the impact of the disaster on the facilities themselves, we would like facilities to be at a safe distance away from the disaster sites that affect large populations. Moreover, the model would place more facilities around areas with large accumulation of population centers that can generate large demands collectively.

Our formulation does not directly consider potential damage to the transportation infrastructure, e.g. damage to roadways between emergency response facilities and population centers. We feel that the safety of the facilities and the supplies stored there is of higher importance for a location problem than the distribution aspect. Indeed, if the emergency response facilities are substantially damaged (and thus supplies are destroyed), the distribution problem becomes moot. In addition, typically there are multiple transportation modes available for supplies in emergency situations. For example, if roadways are destroyed, following an earthquake, supplies can be transported via air. During hurricane Katrina and the ensuing flooding, the majority of supplies were transported by boat. Thus, if alternate (albeit high-cost) transportation modes are available, we believe that the functioning of facilities and the safety of supplies, rather than the availability of the primary transportation mode, ought to be the main determinant of suitable locations for emergency response facilities.

3.2 The Stochastic Distance Dependent Model

To incorporate the stochastic nature of damage caused by a disaster, we now present a two-stage stochastic programming model with binary first stage where the random variable $\tilde{p}_{ex}$, which denotes the damage at $x \in I \cup F$ as a fraction of the total demand/capacity, replaces p_{ex}. It is assumed that the distribution of $\tilde{p}_{ex}$ is known or can be obtained. In the case study presented in Section 5, we will represent $\tilde{p}_{ex}$ using (discretized) uniform and normal distributions of earthquake damage.

In the model presented below, all the variable and parameter definitions remain the same as for the deterministic model (3.1)-(3.6). The first stage decision is to choose the locations of a given number (k) of facilities. Once a disaster occurs, the demands and reduced capability of the facilities are known, and the second stage decisions are the amounts to be routed from the opened facilities to the demand points for each disaster scenario.

$$\min \quad E_{\tilde{p}}[f(x, \tilde{p})] \tag{3.7}$$

$$s.t \quad \sum_{f \in F} x_f \leq k, \qquad x \in \{0,1\}^{|F|} \tag{3.8}$$

Where for a particular realization p of damage $\tilde{p}$, we have

$$f(x,p) = \sum_{e \in E} \sum_{f \in F} \sum_{i \in I} w_e \bar{d}_{if} y_{eif} \tag{3.9}$$

$$s.t. \quad \sum_{i \in I} y_{eif} \leq (1 - p_{ef}) c_f x_f \qquad \forall e \in E, \forall f \in F \tag{3.10}$$

$$\sum_{f \in F} y_{eif} \geq p_{ei} D_i \qquad \forall e \in E, \forall i \in I \tag{3.11}$$

$$y_{eif} \geq 0 \qquad \forall e \in E, \forall i \in I, \forall f \in F \tag{3.12}$$

Note that an expectation over all the possible disaster scenarios is already implicit in the objective function (3.9) of the subproblem. The only random variables we take an expectation over in the objective function (3.7) of the master problem are the damages to facilities and cities in each of the disaster scenarios.

Intuitively, we expect that this model will locate facilities in safer locations than the deterministic model to safeguard against the worst case scenarios where the damage to facilities and demand points is very high.

4 Solution Method

The deterministic formulation (3.1)-(3.5) is solved by using the CPLEX 12 mixed integer programming solver. To solve the stochastic formulation (3.7)-(3.12), a modified L-shaped method that optimizes the Sampling Average Approximation (SAA) of the stochastic program is used. In SAA, a stochastic program with a large number of scenarios is approximated by sampling the random variables and generating a small representative set of scenarios. Multiple approximations are generated and solved using Algorithm 1 to obtain a pool of approximate solutions. These solutions are then tested over a larger set of sample scenarios to identify the solution that performs the best on a much closer approximation of the stochastic program. In our computations, three realizations of 500 samples each are first generated for initial approximation, and each of the three is solved using Algorithm 1. The solutions obtained are then tested on an approximation with 10,000 sample scenarios.

Algorithm 1 is a modification of the usual L-shaped method (which is based on Benders algorithm (Benders, 1962)). In the regular Benders method the master problem is optimized to obtain a lower bound on the overall optimal, and the solution supplied to the subproblems. With each iteration, the master program grows in size due to the Benders cut supplied by the sub problem, and solving the master problem to optimality becomes computationally prohibitive. Cote and Laughton (1984) overcome this problem by using any feasible solution to the master problem. However, their algorithm suffers from slow convergence since even the bad solutions are evaluated as long as they are feasible. Poojari and Beasley (2009) use a variation where they solve the master problem in each iteration, but also use a genetic algorithm to create a pool of good solutions and add optimality cuts to the master problem. In this manner, the number of times the master problem has to be solved is reduced by a constant factor. Some researchers have introduced a variant where the master problem is always solved using a heuristic, completely eliminating the need for an exact method, but this approach cannot guarantee optimality of the overall algorithm. Holmberg (1994) provides a good survey of methods that solve the Benders master program approximately.

Our formulation of the master program is given below (refer to (4.13)-(4.16)):

$$\min \quad \eta \tag{4.13}$$

$$s.t \quad \sum_{f \in F} x_f \leq k \tag{4.14}$$

$$\eta + \beta_t^T x \geq \alpha_t \qquad t = 1 \ldots N \tag{4.15}$$

$$x \in \{0,1\}^{|F|} \tag{4.16}$$

In our approach, the master program (4.13)-(4.16) is solved using a heuristic without updating the lower bound until the heuristic provides the same solution in two successive iterations. The

Algorithm 1 `Modified L-shaped`: Two-stage Stochastic Programs with binary first stage

$x^0 \leftarrow$ initial feasible solution
$LB \leftarrow -\infty$
$UB \leftarrow \infty$
while $UB - LB > \epsilon$ **do**
  Solve subproblems using x^t.
  Update UB.
  Add cut $\eta + \beta_t^T x \geq \alpha_t$ to master problem.
  Heuristically solve master problem to obtain x^{t+1}.
  if $x^t = x^{t+1}$ **then**
    Optimally solve master problem and update x^{t+1}.
    Update LB.
  end if
end while

Algorithm 2 `Master Problem Heuristic`

$r_t \leftarrow \alpha_t,\ t = 1 \ldots N$
for $i = 1 \ldots k$ **do**
  Choose f such that $x_f = 0$ and $\max_{t=1\ldots N}(r_t - \beta_{tf})$ is minimized.
  Set $x_f \leftarrow 1$
  Set $r_t \leftarrow r_t - \beta_{tf},\ t = 1 \ldots N$
end for
while g, h, i & j exist such that $x_g = x_h = 0$ and $x_i = x_j = 1$ and $\max_{t=1\ldots N}(r_t - \beta_{tg} - \beta_{th} - \beta_{ti} - \beta_{tj}) < \max_{t=1\ldots N} r_t$ **do**
  $x_g \leftarrow 1,\ x_h \leftarrow 1,\ x_i \leftarrow 0,\ x_j \leftarrow 0$
end while

rationale for this approach is two-fold: 1) if the solution to the heuristic repeats, the same cut will be generated again and the heuristic will keep generating the same solution endlessly; 2) a good heuristic will provide distinct solutions in successive iterations unless the Benders algorithm is close to optimality, thus reducing the number of times an exact algorithm has to be used considerably. Since the heuristic is not guaranteed to find the optimal solution to the master problem, the lower bound value is updated only when the master program is solved to optimality using an exact method.

A greedy solution technique, supplemented by local search described in Algorithm 2 is used as the heuristic solution technique for solving the master problem (obtained after N iterations of Benders algorithm). To the best of our knowledge, this selective use of a heuristic to repeatedly and quickly solve the master program, but only updating the lower bound using an exact master program solution, is new in the literature. The primary advantage of our technique over existing modifications surveyed in Holmberg (1994) is that it allows for significant savings in computation time, while at the same time preserving optimality of the solution.

We find that using Algorithm 1 in conjunction with Algorithm 2 reduces the time spent in solving the master program by an order of magnitude. Computational experiments show a speed up of at least two times using the modified algorithm over regular Benders approach, with solving the 500 scenario subproblems becoming the bottleneck for further speedup.

Although the algorithmic modifications in this section were designed for the models presented here, the essential underlying principles are fairly general. These modifications can be applied to any general problem where Benders decomposition is applicable and solving the reduced master program in each iteration is the bottleneck. In fact, the modifications only require a good heuristic to be designed for the reduced master problem.

5 Case Study: Pre-Positioning for Earthquake Damage

To obtain computational insight into our model formulation, we consider earthquakes as a threat for which we are pre-positioning facilities. Earthquakes are chosen as an example of a large-scale emergency situation because they recur fairly frequently in certain regions, they have the potential to devastate large areas, and, from the perspective of our modeling approach, their distance-damage functions have been measured and are empirically well understood (Howell and Schultz, 1975). Furthermore, fairly advanced literature is available in probabilistic earthquake forecasting, which can provide possible earthquake epicenters required by our model (Rundle et al., 2003).

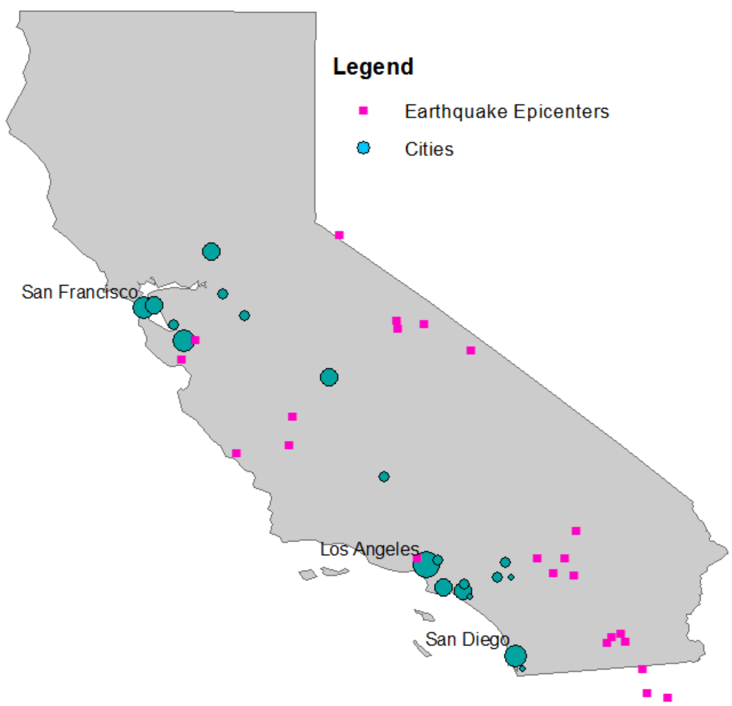

Figure 2: Parameter setting showing the demand points (20 largest cities in California) and the disaster scenarios (23 earthquakes of magnitude larger than 6 that occurred in California since 1973).

We use a case study on California to illustrate the effectiveness and behavior of our models. We use the 20 largest cities (by population) in California as the demand points that might require emergency supplies in case of an earthquake, with demands proportional to their respective populations. The same 20 cities, as well as 38 additional grid points with integer longitude and latitude values spread across California are chosen as potential facility locations.

We place a secondary facility in Utah to serve any remaining demand that is not served by the facilities in California due to unavailability of supplies. The location of that secondary facility was chosen with primary concern for survival of supplies. In reality, either a single central facility or a network of secondary facilities will have to provide supplies in case the primary facilities fail to do so. Using this artificial secondary facility allows us to incorporate some notion of a service level objective in our model: we can now compare the percentage of affected population served by primary facilities in different models, rather than observing feasibility of a solution in one model versus infeasibility in another model. For implementation purposes, the variable corresponding to this secondary facility is fixed to one to indicate that it is open, and the costs associated are multiplied by a factor of 100 to penalize unavailability of supplies in the k primary facilities located in California. The rationale behind assigning a large penalty factor to supplies availed from the secondary facility is that ideally we would prefer all the supplies to come from the k facilities within California. Supplies coming from the secondary facilities are not only delayed, but also undermine the purpose of locating facilities in California. In addition to the extra transportation cost, the factor penalizes the additional response time that will be required to obtain supplies from the secondary facility.

The earthquake scenarios are constructed by using historic data of the 23 earthquakes of magnitude larger than 6 on the Richter scale that have occurred in and around California since 1973 (see USGS)*. An expectation over these 23 scenarios (assumed equiprobable) is taken in the second stage of the stochastic programming formulation. The earthquake epicenters and 20 demand points are shown in Figure 2.

Candidate epicenters are taken directly from historical epicenter locations of high-impact earthquakes, providing a simple model of the future earthquake location distribution. More complex ways of estimating the future epicenter location distribution can be devised and used in our model, but for this case study we will rely on this simple estimation.

In the initial set of experiments, the number of facilities is varied from 1 to 3 and the capacity of each facility is varied between 2, 2.5, 3 & 3.5 million units. The capacity of a facility can be thought of in the same unit dimension as the demand points. For example, the capacity could be measured by the amount of ready-to-deploy disaster response packages containing food, water, first aid supplies etc., that are stored at the facility.

The damage due to an earthquake is modeled using the intensity-distance function provided in Howell and Schultz (1975). For the purpose of our computational study, we assume the damage to be proportional to the intensity and to be given as a function of the distance from the epicenter as $p_{ex}(d_{ex}) = 0.69e^{(0.364-0.130\ln(d_{ex})-0.0019d_{ex})}$ where d_{ex} is measured in kilometers (Howell and Schultz, 1975).

To incorporate stochasticity of earthquake damages, $\tilde{p}_{ex}$ is modeled in the initial set of ex-

*The data used is available at http://people.tamu.edu/~anuragverma/SEFLdata

# Facilities	Cap (mil)	k-Median		Deterministic		Stochastic		
		Z_{KM}	P_{KM}	Z_{DET}	P_{DET}	Z_{STO}	P_{STO}	Runtime
1	2	788.70	50.54	624.93	39.71	624.93	39.71	590.03
	2.5	628.78	40.34	444.83	28.04	444.83	28.04	584.36
	3	495.48	31.78	302.17	18.75	302.17	18.75	643.22
	3.5	389.42	24.94	196.43	11.85	196.43	11.85	655.02
2	2	244.28	15.57	124.72	7.24	121.64	7.00	635.25
	2.5	132.23	8.34	51.38	2.51	45.61	2.01	830.06
	3	74.36	4.61	50.09	2.86	22.21	0.51	2448.47
	3.5	44.40	2.69	34.15	2.00	14.90	0.14	4265.94
3	2	63.29	3.97	30.96	1.64	21.63	0.54	3488.57
	2.5	26.24	1.58	24.53	1.45	9.21	0.25	23517.90
	3	12.54	0.70	11.87	0.66	5.14	0.04	59214.32
	3.5	6.63	0.33	6.34	0.31	3.62	0.04	90145.24

Table 1: Comparison of the expected cost of providing supplies to disaster affected areas in million people-miles using the optimal locations found by the stochastic model (Z_{STO}), the deterministic model (Z_{DET}) and the traditional k-median model without conditional availability (Z_{KM}). The percentage of demand that is met by the secondary facility in Utah is also presented (P_{STO}, P_{DET} and P_{KM}).

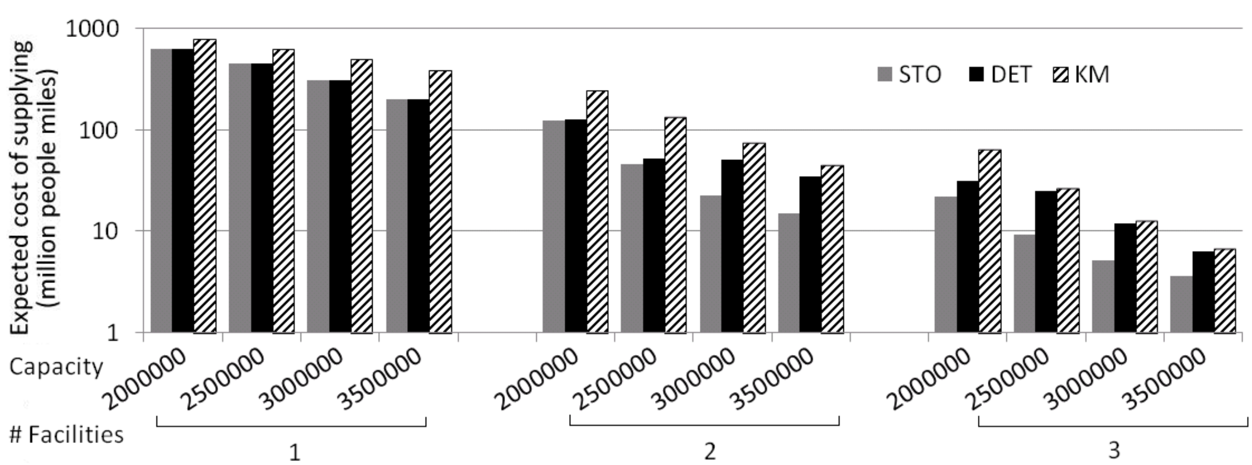

Figure 3: Expected cost of providing supplies to disaster affected areas in million people-miles using the optimal locations found by the stochastic model (Z_{STO}), the deterministic model (Z_{DET}) and the traditional k-median model without conditional availability (Z_{KM}) plotted on the logarithmic scale.

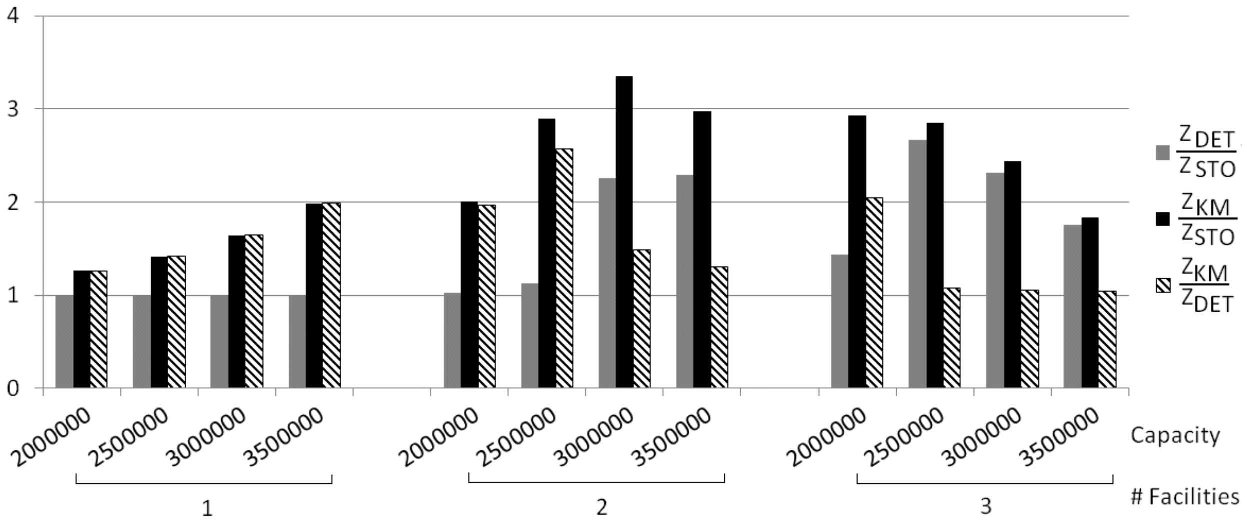

Figure 4: Ratios of the cost of providing supplies to disaster affected areas as obtained by the 2-stage stochastic model, the deterministic model, and the k-median model.

periments as a discretized uniform random variable that can take on the values p_{ex}, 0 and $2p_{ex}$ with equal probability. Note that p_{ex} is the corresponding damage in the deterministic model, and $E[\tilde{p}_{ex}] = p_{ex}$. In subsequent experiments, we explore the sensitivity of model performance to the assumed variability in damage by exploring a (discretized) normal distribution of damage across a wide range of coefficients of variation.

Table 1 provides the expected cost of serving the affected population in case of a disaster using the locations found by the stochastic and deterministic models (Z_{STO} and Z_{DET}) developed in this paper, and also a traditional k-median model. We include the k-median model performance solely to demonstrate the impact of not adequately considering facility unavailability due to damages. We do not advocate using a traditional k-median model for large-scale emergencies. All deterministic models solve within less than 60 seconds; the solution times (in seconds) for the stochastic model are shown in the last row of the table. As one can see, solution times for even modestly-sized instances of the stochastic model require several hours of computing time. As is often the case in stochastic facility location problems, solution times become problematic quickly, even considering our modification to speed up Benders Decomposition.

We observe that as the total built capacity increases (either because the number of facilities or the capacity increases), the expected costs go down. This is apparent in Figure 3, which plots Z_{STO}, Z_{DET} and Z_{KM} against various configurations of capacity and the number of facilities built. The same trend can be observed in the percentage of the population that has to be served by the external facility in Utah on an average across all disaster scenarios (P_{STO}, P_{DET} and P_{KM}). Each of P_{STO}, P_{DET} and P_{KM} decrease as the built capacity of the facilities is increased. Both trends are fairly intuitive, since increasing overall built capacity should reduce the costs and the need to avail the services of the secondary facility.

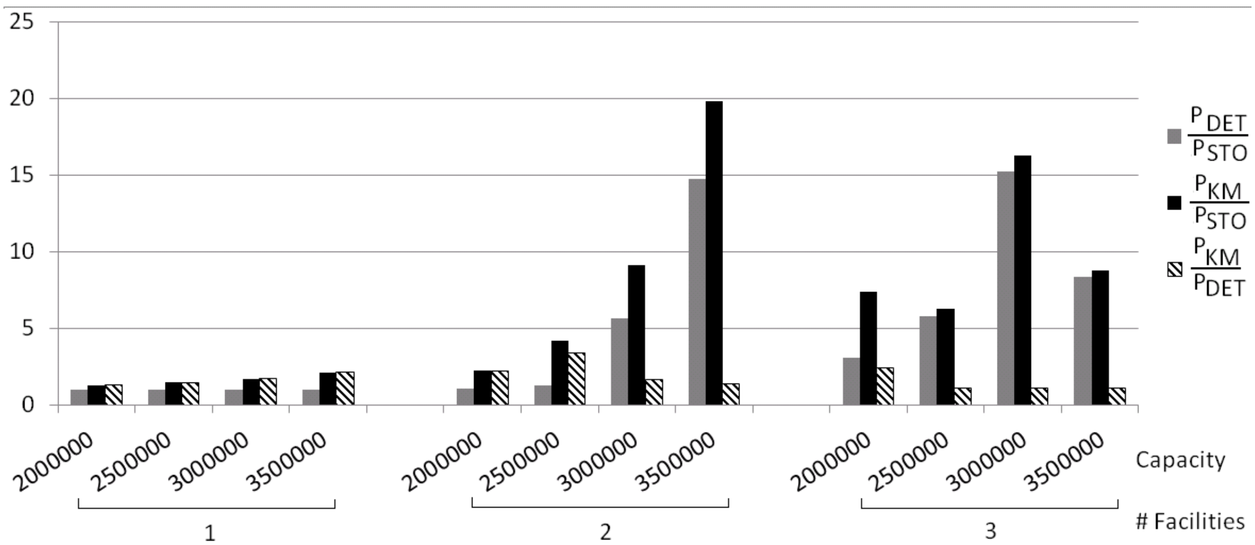

Figure 5: Ratios of the demand that had to be satisfied by the secondary facility when the primary facilities were located using the 2-stage stochastic model, the deterministic model, and the k-median model.

Some more insightful observations can be made when the costs and percentage demand met by the secondary facility are compared between locations found by the stochastic, deterministic and k-median models. To this end, in Figures 4 and 5 we provide some measures that are more demonstrative of the effects of using the different models. For example, Figure 4 provides plots of Z_{DET}/Z_{STO}, which measures the effect of using the stochastic model against the use of the deterministic model, Z_{KM}/Z_{STO}, which draws a similar comparison with the k-median model instead, and Z_{KM}/Z_{DET}, which compares the k-median model to the deterministic model. Although in general the stochastic solution does better overall in these metrics, we observe that there exists a range of built capacity where the stochastic solution offers particularly significant improvement over the deterministic and k-median solutions. In this range of capacity, the stochastic solution is three to four times better than the k-median solution and 2 times better than the deterministic solution when the objective values are compared.

The same trend is also apparent in Figure 5, which plots P_{DET}/P_{STO}, P_{KM}/P_{DET}, and P_{KM}/P_{STO} for different numbers and capacities of the facilities. When the built capacity is on the lower end or the higher end, both Z_{DET}/Z_{STO} and P_{DET}/P_{STO} are close to 1, indicating that both deterministic and stochastic models are providing similar results. This is because when the total built capacity is very low when compared to the demands, there is not much difference between locations found by the deterministic and stochastic models because both models try to place facilities as far away from the disasters to preserve supplies, and obtain more or less the same solutions. On the other hand, when the total built capacity is very high, both the models locate facilities at similar locations since the stochasticity of the damage is absorbed by the high built capacity. Thus, in both cases, the stochasticity is not the primary factor determining the locations. The stochasticity comes into play when the built capacity is such that its utilization by the potential demand is medium to high. In subsequent, expanded experiments, we study this behavior in more detail by keeping total built capacity constant and observing the impact of facility density on model performance.

Figure 4 and 5 also provide Z_{KM}/Z_{DET} and P_{KM}/P_{DET} which show that while in general the deterministic solution is better than the k-median solution, there is a particular range of total built capacity in which the deterministic model is almost two times better than the k-median solution in terms of the objective function. This can again be attributed to the fact that when the built capacity is too low, none of the models is able to provide good locations, while when the capacity is too high, the locations found by either model are similar, because damage to facilities (which results in a reduction of capacity), becomes less relevant.

The quantities Z_{KM}/Z_{STO} and P_{KM}/P_{STO} measure the overall benefits of using the stochastic model over the k-median model. These benefits are in essence the combination of the benefits obtained by using the stochastic model over the deterministic model, plus the benefits of using the deterministic model over the k-median model. As a result, the similar trends in Z_{KM}/Z_{DET} and Z_{DET}/Z_{STO} are accentuated in the metric Z_{KM}/Z_{STO}, resulting in the stochastic solution being three to four times better in terms of the objective function value in the range where the built capacity is neither too high nor too low. The reason for this can be attributed to a combination of the reasons outlined in the previous two paragraphs - in this range, both stochasticity of the damage and the damage to facilities come into play.

It should be noted that a disaster management agency will need to build disaster response facilities with capacities that are both economical and effective. This becomes especially true in the current scenario with federal and state governments restricted in their spending, and occurrences of natural and man-made disasters on a rise (see, for example, the International Disaster Database EM-DAT (EM-DAT)). The results from Table 1, as well as the trends observed in Figure 4 & 5, suggest that our stochastic solution turns out to be significantly better than the deterministic solution in precisely the range of capacities that is of primary interest to decision makers. Thus, the stochastic model appears to provide significant added value over the deterministic model in terms of finding good locations under realistic budget constrained conditions.

5.1 Analysis of the locations suggested by different models

We now analyze the facility locations provided by the three models when the capacities are in the critical range discussed in the previous paragraph. Figure 6 shows the locations of three facilities with capacities set to two million as found by the deterministic model and the k-median model. It can be observed that the deterministic model developed in this paper places two facilities in similar locations of the k-median model, but the third facility is moved from a very earthquake prone location to a much safer location. The cost of supplying facilities from this safer location is higher, but the availability of supplies is higher than the closer but risk-prone facility. With the third facility at the location suggested by the k-median model, a small part of the supplies comes from a close location, but a large penalty is incurred for the remaining demand that is fulfilled from the secondary facility. As a result, the deterministic model finds it beneficial to pay a higher transportation cost to obtain a large part of the required supplies from the safer location.

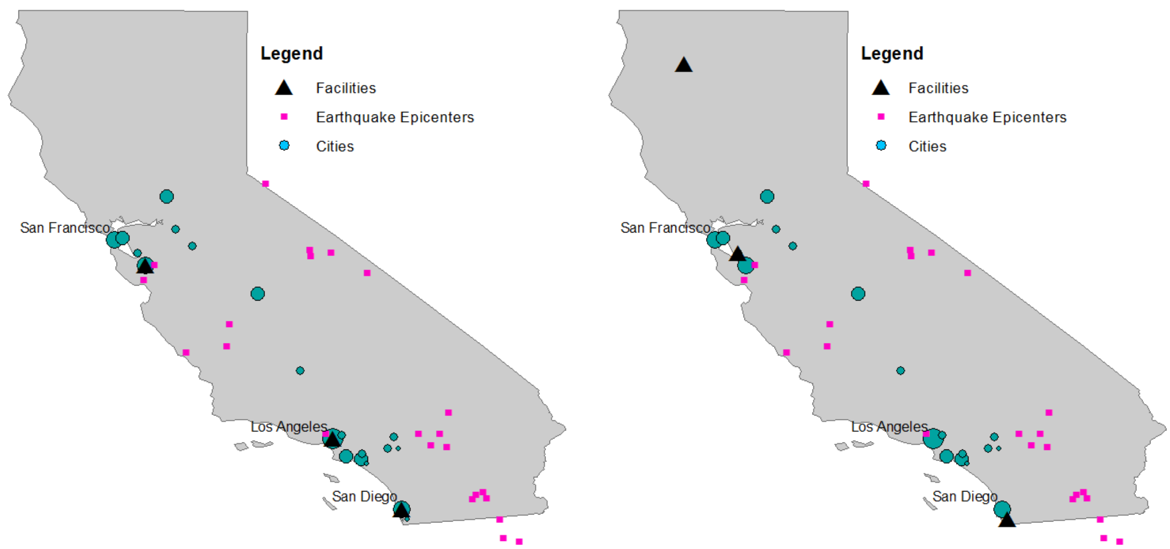

(a) k-Median model (b) Deterministic model

Figure 6: Comparison of locations of three facilities of capacities 2 million as found by the Deterministic and k-Median models.

Figure 7 shows the locations of facilities when the capacities are set to 3.5 million as found by the deterministic model and the stochastic model. A noticeable change in Figure 7(a) from Figure 6(b) is that as the capacities are increased, the deterministic model places facilities in areas of high demand even if they are risk prone. This is because the higher capacities result in little or no penalty even when the facilities are partially damaged. However, in Figure 7(b) the stochastic model still places the one facility far away from the risky Los Angeles area. This can be attributed to the fact that the stochastic program is trying to avoid large penalties resulting from scenarios where the damage from an earthquake is high by placing facilities in relatively safer locations than the deterministic model would.

5.2 Effects of varying the penalty for second-tier sourcing

Table 4 depicts the expected service cost as the penalty of availing supplies from the external facility from Utah is changed between 1 (no penalty), 10 (low penalty) and 100 (high penalty). Penalties act as a parameter in our models by which we can explore the time-sensitive nature of service in the aftermath of a large scale emergency. The choice of the penalty needs to be determined by the decision maker, and will take into consideration various factors that reflect the ”cost” of not providing timely service and having to bring in supplies from an external warehouse. We believe that larger penalties are a more realistic choice in most cases since the service here is critical to a large population.

We can draw the following conclusions from the results: First, as the penalty is increased, for the same number and capacity combination, more demand is satisfied by the model that uses a high penalty. This is intuitive and expected, since a higher penalty forces the model to choose a solution

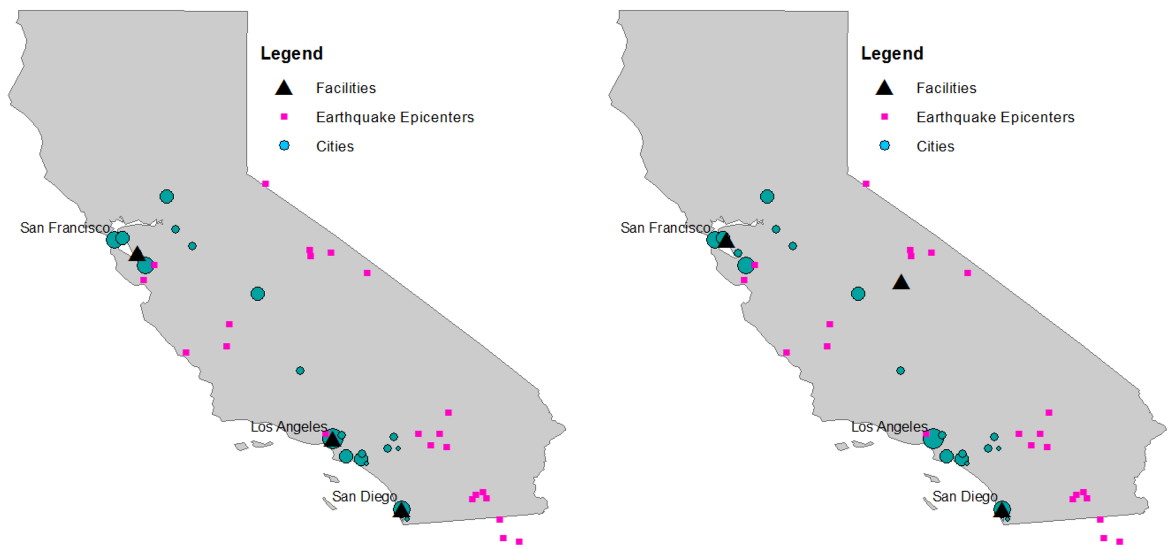

(a) Deterministic model (b) Stochastic model

Figure 7: Comparison of locations of three facilities of capacities 3.5 million as found by the Stochastic and Deterministic models.

that will require less services from the external facilities. Second, when there is a low penalty, or no penalty at all, the stochastic solution does not improve on the deterministic solution by much when compared to the high penalty case. This can be attributed to the fact that when the penalty is low, the cost of availing supplies from the external facility in some extremely bad scenarios is not large enough to change a good location that serves very well during "average" damage scenarios. Thus, the stochastic model also chooses locations similar to the deterministic equivalent model.

5.3 Effects of varying the density of facilities

A key question decision makers are interested in is the following: if the total built capacity is constant, is it better to establish one large facility or many smaller facilities? To obtain insight into this question, we run experiments where we vary the number of facilities while keeping total built capacity constant. We run the following scenarios: 1 facility with capacity 6 million units, 2 with 3 million units, 3 with 2 million units, and 4 with 1.5 million units. Table 2 summarizes the experiment results.

Figures 8 and 9 show that when total capacity is held constant, the cost of serving demand, as well as the fraction of demand that is satisfied by the second-tier facility, decrease for both the deterministic and the stochastic models. It is clear that this is not the result of a capacity saturation effect, because total built capacity is constant among these cases. Instead, a denser network of facilities is beneficial for two main reasons: First, a denser network of facilities helps decrease transportation distances, and, second, it provides a form of diversification of risks: geographically diversified facilities are less likely to be wiped out by a single disaster occurrence.

While both the deterministic and stochastic models are similar in that they show improvement

# Facilities	Cap (mil)	k-Median		Deterministic		Stochastic	
		Z_{KM}	P_{KM}	Z_{DET}	P_{DET}	Z_{STO}	P_{STO}
1	6	145.42	9.12	101.41	5.88	24.26	0.54
2	3	74.36	4.61	50.09	2.86	22.21	0.54
3	2	63.29	3.97	30.96	1.64	21.63	0.51
4	1.5	65.97	4.16	27.97	1.45	21.19	0.54

Table 2: Comparison of the expected cost of providing supplies to disaster affected areas in million people-miles using the optimal locations found by the stochastic model (Z_{STO}), the deterministic model (Z_{DET}) and the traditional k-median model without conditional availability (Z_{KM}). The percentage of demand that is met by the secondary facility in Utah is also presented (P_{STO}, P_{DET} and P_{KM}).

as the facility density increases, a closer look reveals some differences. Figure 10 shows a comparison between the stochastic, deterministic and k-median models in the form of ratios of the cost of transportation. The biggest takeaway from observing Z_{DET}/Z_{STO} is that the benefit of the stochastic model reduces with the number of facilities. In particular, when only a single facility is established, the stochastic solution is approximately 25% of the cost of the deterministic solution. When four facilities are placed, this cost advantage reduces to roughly 75%. This behavior is reasonable since by establishing more (but smaller) facilities, we are hedging against the risk of a disaster wiping out large portions of total supply. In addition, the benefit of the deterministic distance-damage model over the k-median model increases with the number of facilities (Z_{KM}/Z_{DET}). This is because the distance-damage model improves significantly as the number of facilities goes up, while the k-median model also improves, but at a slower rate.

Overall, we observe that even though the deterministic distance-damage model becomes more attractive in relative terms as facility density increases, the stochastic model remains superior in absolute terms, even when the number of facilities is comparatively larger. At the same time, the less densely the facilities are located, the more benefit is derived from the stochastic model solution. From the decision maker's perspective, this means that if budget constraints dictate placing only a few disaster response facilities, those few locations should be determined using the stochastic model.

5.4 Effects of variability in damage

So far we have modeled the variation in damage as a discrete uniform random variable with values p_{ex}, 0, and $2p_{ex}$ with equal probability. To study the effects of different degrees of variability of the damage, we now model the random variable for damage as a normally distributed variable with $\mu = p_{ex}$ and $\sigma = cp_{ex}$, where c is the coefficient of variation (CoV) and takes on the values 0.2, 0.4, 0.8, 1.2, and 1.6. For comparison, the CoV of the discrete uniform random variable base case is 0.81.

As in the previous section, we run the following scenarios: 1 facility with capacity 6 million units, 2 with 3 million units, 3 with 2 million units, and 4 with 1.5 million units. Table 3 summarizes

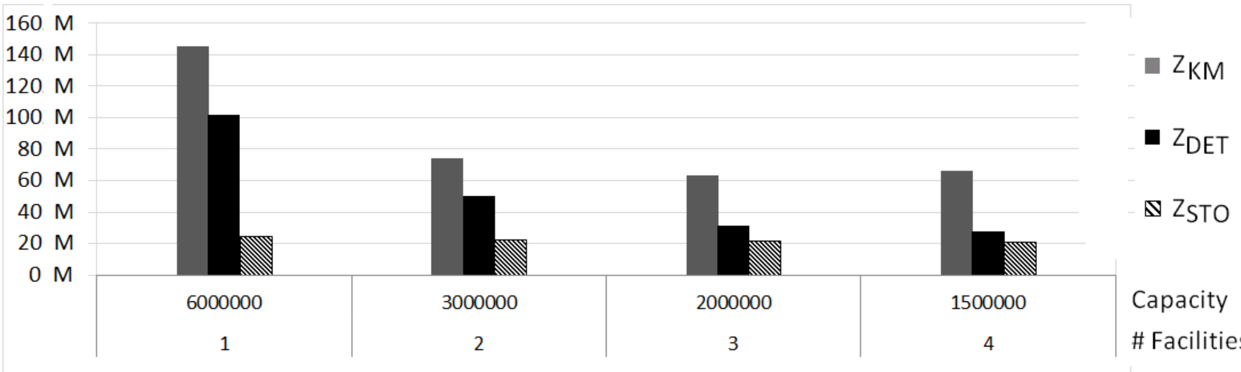

Figure 8: Cost of providing supplies to disaster affected areas as obtained by the 2-stage stochastic model, the deterministic model, and the k-median model. Plotted for different configurations of same total built capacity.

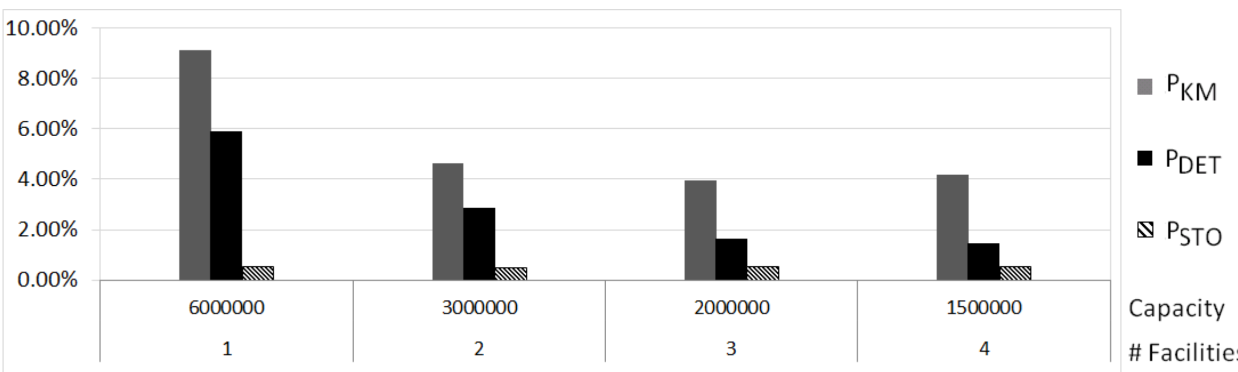

Figure 9: Percentage of supplies to disaster affected areas coming from the facility in Utah, as obtained by the 2-stage stochastic model, the deterministic model, and the k-median model. Plotted for different configurations of same total built capacity.

the experiment results.

Figure 11 shows the measures Z_{DET}/Z_{STO}, Z_{KM}/Z_{DET}, and Z_{STO}/Z_{KM} for the different CoV values and for cases with the same total built capacity, but different number of facilities.

The overall trends in the results are very consistent across the four cases of facility densities. In general, the results suggest that the benefit of using a stochastic model over the deterministic model increases as the CoV of damage increases. This result confirms intuition because with more uncertainty about damage realizations, a deterministic model will perform worse.

Surprisingly though, we also observe that as CoV increases, the advantage of the distance-dependent model over the k-median model (that is, Z_{KM}/Z_{DET}) decreases. This result suggests that as the variability of damage increases, the advantage of better modeling of damage *in a*

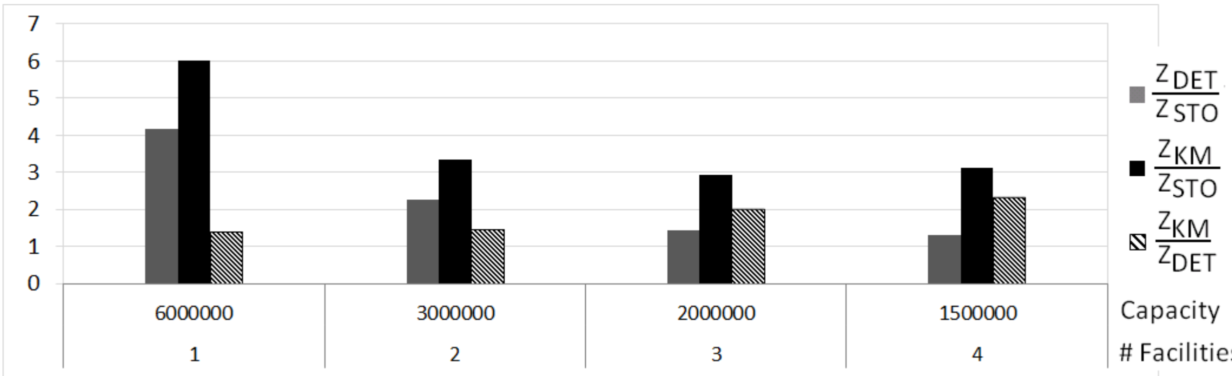

Figure 10: Ratios of the cost of providing supplies to disaster affected areas as obtained by the 2-stage stochastic model, the deterministic model, and the k-median model. Plotted for different configurations of same total built capacity.

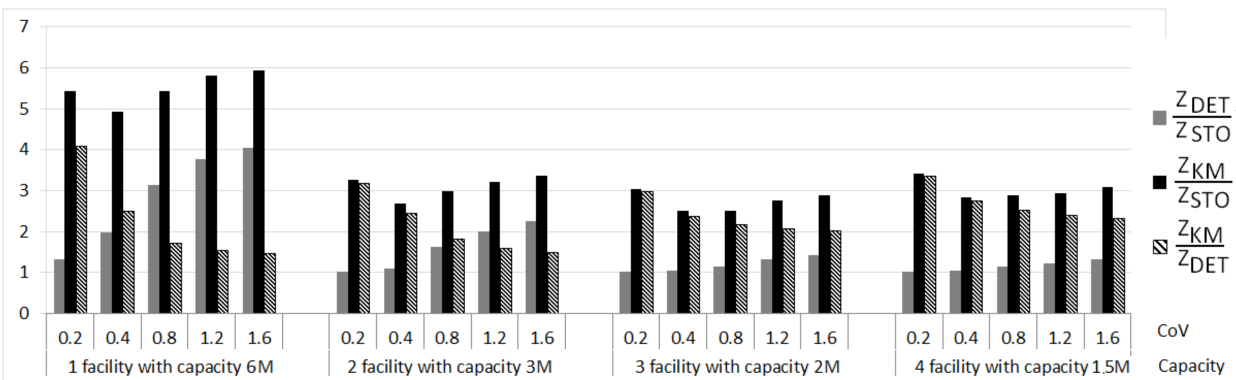

Figure 11: Ratios of the cost of providing supplies to disaster affected areas as obtained by the 2-stage stochastic model, the deterministic model, and the k-median model. Plotted for different configurations of same total built capacity, and by varying CoV.

deterministic setting starts to dissipate. An explanation for this observation is that as realizations of damage become more variable (”random”), what model is used matters less and less, if the model is deterministic and thus does not explicitly consider stochasticity. If the model explicitly considers stochasticity on the other hand, performance will increase over non-stochastic models as variability increases, as previously noted.

# Facilities	Cap (mil)	CoV	k-Median		Deterministic		Stochastic	
			Z_{KM}	P_{KM}	Z_{DET}	P_{DET}	Z_{STO}	P_{STO}
		0.2	74.60 M	4.57%	18.27 M	0.62%	13.73 M	0.21%
		0.4	88.61 M	5.48%	35.59 M	1.73%	18.02 M	0.37%
1	60	0.8	119.00 M	7.43%	68.97 M	3.84%	21.91 M	0.49%
		1.2	138.58 M	8.68%	89.84 M	5.16%	23.90 M	0.52%
		1.6	150.47 M	9.44%	102.55 M	5.96%	25.33 M	0.61%
		0.2	35.51 M	2.14%	11.16 M	0.33%	10.92 M	0.26%
		0.4	41.98 M	2.54%	17.14 M	0.72%	15.70 M	0.39%
2	30	0.8	59.72 M	3.67%	32.70 M	1.72%	20.08 M	0.46%
		1.2	71.47 M	4.42%	44.68 M	2.50%	22.22 M	0.50%
		1.6	78.45 M	4.87%	52.56 M	3.01%	23.40 M	0.58%
		0.2	29.99 M	1.85%	10.07 M	0.33%	9.91 M	0.21%
		0.4	34.54 M	2.12%	14.48 M	0.60%	13.78 M	0.43%
3	20	0.8	49.20 M	3.05%	22.56 M	1.11%	19.66 M	0.73%
		1.2	59.60 M	3.72%	28.66 M	1.49%	21.63 M	0.55%
		1.6	66.02 M	4.13%	32.68 M	1.75%	22.86 M	0.61%
		0.2	33.17 M	2.07%	9.92 M	0.32%	9.72 M	0.27%
		0.4	38.46 M	2.40%	13.97 M	0.57%	13.54 M	0.40%
4	15	0.8	52.75 M	3.31%	20.82 M	1.00%	18.27 M	0.64%
		1.2	62.47 M	3.93%	26.00 M	1.33%	21.31 M	0.76%
		1.6	68.49 M	4.31%	29.49 M	1.55%	22.24 M	0.57%

Table 3: Comparison of the expected cost of providing supplies to disaster affected areas in million people-miles using the optimal locations found by the stochastic model (Z_{STO}), the deterministic model (Z_{DET}) and the traditional k-median model without conditional availability (Z_{KM}) for constant built capacity cases with varying degree of variability in damage function. The percentage of demand that is met by the secondary facility in Utah is also presented (P_{STO}, P_{DET} and P_{KM}).

6 Conclusion

This paper provides two location models for large-scale emergencies that explicitly acknowledge the fact that facility failures will often be caused by the very disasters they are supposed to provide relief from. The paper also presents a variant of Benders Algorithm that has been found to be very effective in solving the models presented and could be extended to other facility location models.

The models proposed in this paper enhance the emergency facility location literature by treating the availability of a facility and the demands placed on a facility directly within the model through the introduction of a general distance-damage function. This model-internal use of disaster scenarios and disaster consequences also allows us to extend the modeling approach into a

Penalty	Capacity	Num. Facilities		
		1	2	3
	2000000	1.00	1.02	1.02
1	2500000	1.00	1.00	1.00
	3000000	1.00	1.00	1.01
	3500000	1.00	1.02	1.00
	2000000	1.00	1.05	1.09
10	2500000	1.00	1.18	1.05
	3000000	1.00	1.06	1.00
	3500000	1.02	1.05	0.99
	2000000	1.00	1.03	1.43
100	2500000	1.00	1.13	2.66
	3000000	1.00	2.26	2.31
	3500000	1.00	2.29	1.75

Table 4: Ratio of the expected cost of serving the disaster affected population by facilities located using the deterministic model and by the stochastic model as the penalty of serving using the external facility is increased.

stochastic-programming treatment of the uncertainty of damages caused at each facility.

Using a case study on pre-positioning supplies for earthquakes at disaster response facilities in California, we demonstrate that naive deterministic models such as the k-median model tend to produce location results that can turn out to be highly undesirable once the potential effects of the disaster on the response facilities themselves are taken into account. Our deterministic distance-dependent model performs reasonably well, especially when considering its very quick solving times of less than one minute for all instances in our case study. The stochastic distance-dependent model provides consistently the best solution quality, at the price of multiple-hour solving times.

Our numerical results indicate that incorporating the stochastic nature of disasters is crucial for a cost effective relief effort. The stochastic model developed here allows us to study a more realistic scenario where predicting the damage caused by a disaster is not straightforward. The superior performance of the stochastic model in our case study highlights the shortcomings of deterministic models for this application. Moreover, we provide insight into the impact of the variability of damage intensity on the solution quality through a sensitivity analysis on the coefficient of variation of the random variable describing the damage of the disaster. We also demonstrate the impact that the density of the disaster response facilty network has on the respective solution qualities of the models.

In particular, our comparison of the performance of the distance-damage deterministic model and the stochastic model shows that the fewer facilities are placed, the more crucial it is to determine placement using the stochastic model rather than the deterministic model. This is an especially important insight in view of shrinking budgets for emergency logistics. Our case study also provides clear evidence that the performance advantage of the stochastic model over the deterministic model increases with the uncertainty about the magnitude of damage. Thus, the more damage can vary (as is the case in particular for earthquakes), the more important it is to determine disaster response facility locations via a model that directly incorporates this stochasticity.

One aspect of facility location in disaster response that still needs further study is managing risk. Many of the models presented in the existing literature, including our present work, aim to minimize the expected damage over the possible disaster scenarios. However, for large scale disasters that are rare events, it may be prudent to locate facilities for worst case scenarios. Building models that do this by incorporating some measure of risk that penalizes the worst case scenarios more may prove to be a valuable future direction. Noyan (2012) has presented work that incorporates a risk-averse decision setting in emergency logistics; it would be interesting to combine a similar conditional-value-at-risk approach with our model. Alternatively, a robust optimization approach similar to the one presented in Ben-Tal et al. (2011) may be workable.

Acknowledgement

We thank Preeti Shankar for preparing the graphics for visual representation of the results using Geographic Information System (GIS) software. We are also grateful to the editor and two anonymous referees whose comments have greatly improved this paper.

References

B. Balcik and B. M. Beamon. Facility location in humanitarian relief. *International Journal of Logistics: Research and Applications*, 11(2):101–121, 2008.

G. Barbarosoglu and Y. Arda. A two-stage stochastic programming framework for transportation planning in disaster response. *The Journal of the Operational Research Society*, 55(1):43–53, 2004.

Aharon Ben-Tal, Byung Do Chung, Supreet Reddy Mandala, and Tao Yao. Robust optimization for emergency logistics planning: Risk mitigation in humanitarian relief supply chains. *Transportation Research Part B*, 45:1177 – 1189, 2011.

J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4:238–252, 1962.

P. Beraldi, M.E. Bruni, and D. Conforti. Designing robust emergency medical service via stochastic programming. *European Journal of Operational Research*, 158:183–193, 2004.

Oded Berman, Zvi Drezner, and George O. Wesolowsky. Locating service facilities whose reliability is distance dependent. *Computers & Operations Research*, 30(11):1683 – 1695, 2003.

Oded Berman, Dmitry Krass, and Mozart B. C. Menezes. Facility reliability issues in network p-median problems: Strategic centralization and co-location effects. *Oper. Res.*, 55(2):332–350, 2007.

Ali Bozorgi-Amiri, M.S. Jabalameli, and S.M.J. Mirzapour Al e Hashem. A multi-objective robust stochastic programming model for disaster relief logistics under uncertainty. *OR Spectrum*, 35: 905 – 933, 2013.

Aakil M. Caunhye, Xiaofeng Nie, and Shaligram Pokharel. Optimization models in emergency logistics: A literature review. *Socio-Economic Planning Sciences*, 46:4 – 13, 2012.

Gilles Cote and Michael A. Laughton. Large-scale mixed integer programming: Benders-type heuristics. *European Journal of Operational Research*, 16(3):327 – 333, 1984.

Alper Doyen, Necati Aras, and Gulay Barbarosoglu. A two-echelon stochastic facility location model for humanitarian relief logistics. *Optimization Letters*, 6:1123 – 1145, 2012.

Serhan Duran, Marco A. Gutierrez, and Pinar Keskinocak. Pre-positioning of emergency items for care international. *Interfaces*, 41:223–237, May 2011. ISSN 0092-2102.

EM-DAT. The international disaster database. URL `http://www.emdat.be/natural-disasters-trends`.

N. Gormez, M. Koksalan, and F.S. Salman. Locating disaster response facilities in istanbul. *Journal of the Operational Research Society*, 62:1239 – 1252, 2011.

Kaj Holmberg. On using approximations of the benders master problem. *European Journal of Operational Research*, 77(1):111 – 125, 1994.

B.F. Howell and T. R. Schultz. Attenuation of modified mercalli intensity with distance from the epicenter. *Bulletin of the Seismological Society of America*, 65:651 – 665, 1975.

H. Jia, F. Ordonez, and M. Dessouky. A modeling framework for facility location of medical services for large-scale emergencies. *IIE Transactions*, 39(1):41–55, 2007.

Walid Klibi, Soumia Ichoua, and Alain Martel. Prepositioning emergency supplies to support disaster relief: A stochastic programming approach. *CIRRELT Technical Publication 19-2013*, 2013. URL `https://www.cirrelt.ca/DocumentsTravail/CIRRELT-FSA-2013-19.pdf`. CIRRELT.

Huseyin Onur Mete and Zelda B. Zabinsky. Stochastic optimization of medical supply location and distribution in disaster management. *International Journal of Production Economics*, 126(1):76 – 84, 2010.

Pavankumar Murali, Fernando Ordonez, and Maged M. Dessouky. Capacitated facility location with distance-dependent coverage under demand uncertainty. *Networks & Spatial Economics*, Submitted to the Special Issue on Facilities Location, 2009. URL `http://www-rcf.usc.edu/m̃aged/publications/facloc_new.pdf`.

Nilay Noyan. Risk-averse two-stage stochastic programming with an application to disaster management. *Computers and Operations Research*, 39:541 – 559, 2012.

Jomon Aliyas Paul and Rajan Batta. Models for hospital location and capacity allocation for an area prone to natural disasters. *International Journal of Operational Research*, 3(5):473–496, 2008.

C.A. Poojari and J.E. Beasley. Improving benders decomposition using a genetic algorithm. *European Journal of Operational Research*, 199(1):89 – 97, 2009.

Carmen G. Rawls and Mark A. Turnquist. Pre-positioning of emergency supplies for disaster response. *Transportation Research Part B: Methodological*, 44(4):521 – 534, 2010.

Carmen G. Rawls and Mark A. Turnquist. Pre-positioning planning for emergency response with service quality constraints. *OR Spectrum*, 33:481 – 498, 2011.

Carmen G. Rawls and Mark A. Turnquist. Pre-positioning and dynamic delivery planning for short-term response following a natural disaster. *Socio-Economic Planning Sciences*, 46:46 – 54, 2012.

J. B. Rundle, D. L. Turcotte, R. Shcherbakov, W. Klein, and C. Sammis. Statistical physics approach to understanding the multiscale dynamics of earthquake fault systems. *Rev. Geophys.*, 41, Feb 2003.

Javier Salmeron and Aruna Apte. Stochastic optimization for natural disaster asset prepositioning. *Production and Operations Management*, 19:561 – 574, 2010.

Carl H. Schultz, Kristi L. Koenig, and Roger J. Lewis. Decisionmaking in hospital earthquake evacuation: Does distance from the epicenter matter? *Annals of Emergency Medicine*, 50(3): 320–326, 2007.

Lawrence V. Snyder and Mark S. Daskin. Reliability models for facility location: The expected failure cost case. *Transportation Science*, 39(3):400–416, 2005.

USGS. URL `http://earthquake.usgs.gov/earthquakes/`.

www.bt.cdc.gov. URL `http://www.bt.cdc.gov/stockpile/`.