New England University Transportation Center
77 Massachusetts Avenue, E40-279
Cambridge, MA 02139
utc.mit.edu

Year 25 Final Report
Grant Number: DTRT13-G-UTC31

Stochastic Analysis of Logistics Capacity in Disaster Response Networks

Project Number:	Project End Date:	Submission Date:
MITR25-7	8/31/18	11/1/18

Principal Investigator:	Jarrod Goentzel	**Co-Principal Investigator:**	Alexander Rothkopf
Title:		**Title:**	
University:	MIT	**University:**	MIT
Email:		**Email:**	
Phone:		**Phone:**	

1 Table of contents

2 Executive summary

Quickly deploying relief items is key to reducing a population's burden in case of sudden onset disasters. Emergency response organizations, such as FEMA or local and state agencies hold a strategic stockpile of critical relief items and contract for contingency stock in preparation for emergencies. Their response capacity depends on their decision to stock items at different depots, contracts with contingency suppliers, and procurement of transportation capacity to move these items.

Building on prior work of Acimovic & Goentzel (2016) we develop a stochastic linear programming model to capture carrier capacity and contingency suppliers. Inputs to the model are a *risk portfolio* reflecting the particular disasters and the inherent uncertainty with respect to when an organization needs to address a large or a small disaster. Further inputs are the organic stockpile of critical relief items, referred to as the *inventory portfolio*, contracts with contingency suppliers, which we term the *supplier portfolio*, and the *portfolio of carrier*s at any depot location.

The model allows to conduct a *system assessment* and a *system optimization*. System assessment evaluates the current state and answers how well the current inventory, supplier, and carrier portfolio is able to meet a given risk portfolio. We present aggregate metrics to assess a system in three dimensions. We evaluate *service metrics* to answer how well the network meets demand of the affected population and how rapidly we reach the affected population, and *efficiency metrics* to indicate how much resources are necessary to meet demand. Taken together these metrics allow to evaluate the state of an emergency response network.

System optimization identifies the optimal allocation of inventory for a given supplier and carrier portfolio against a given risk portfolio. The models provides the above mentioned metrics for a decision-maker to compare to optimal network to the current on. In addition, we prescribe an *inventory balance,* a *carrier contract,* and a *carrier utilization metric* to capture the value of improvement.

In both – system assessment and system optimization – we evaluate a time-based model and a cost-based model to capture the inherent cost-time trade-off. Typically, more responsive suppliers and carriers are more expensive and less responsive suppliers and carriers are less expensive. When choosing where to allocate inventory, and which suppliers and carriers to contract an organization has to resolve this trade-off between cost and time. Our model provides insight into this trade-off and the impact on different performance metrics.

We use data from the openFEMA API to construct a new risk portfolio and estimate an inventory and a carrier portfolio to show the feasibility and functionality of our approach.

3 Introduction

Quickly deploying relief items is key to reducing a population's burden in case of sudden onset disasters. Emergency response organizations, such as FEMA or local and state agencies hold a strategic stockpile of critical relief items and contract for contingency stock in preparation for emergencies. Their response capacity depends on their decision to stock items at different depots, contracts with contingency suppliers, and procurement of transportation capacity to move these items.

Assessing and optimizing response capacity is difficult for multiple reasons. Disasters occur unforeseen and the location and number of people affected varies significantly. Transportation capacity is typically constrained following a disaster, reducing the response time while increasing costs on the spot market. Procurement lead times vary across items and amongst suppliers for any single item, which affects the contingency plans during response and the stockpile resupply that follows.

Acimovic and Goentzel (2016) prescribe a model and develop a set of performance metrics to assess and improve an organization's response capacity. These metrics enable decision-makers to characterize the sufficiency of an organization's capacity against a portfolio of disaster scenarios and delineate strategies to improve performance. Key inputs to the model are disaster scenarios and inventory levels of relevant emergency response organizations. Their main objective is to characterize and evaluate the mix of air and ground level transportation to meet the needs of an affected population considering a global network of disaster relief organizations. The key trade-off of a decision-maker in this case is using airfreight, which is very responsive but expensive, or truck, which is less expensive but typically incurs longer delivery times.

Their analysis provides valuable insight into the response capability of global networks. However, when evaluating a more localized network such as, for example, the disaster response network in the continental United States, the main mode of transportation is truck. With respect to transportation, the trade-off at hand is using responsive but more expensive carriers and less responsive, less expensive carriers. In addition, it is important to evaluate the availability of carrier capacity at each depot, that is, the mix of different carrier contracts, and how much spot market capacity is available, to ensure that inventory levels in depots correspond to the carrier capacity.

This project funded by New England University Transportation Center (UTC) seeks to advance the model of Acimovic and Goentzel (2106) and its metrics to more accurately consider an emergency response organization's carrier portfolio and supplier capability in the context of continental United States (CONUS) response (though extensions outside CONUS could be developed as well). Our model is applicable to a wide range of disaster response organizations. We seek to collaborate with the Federal Emergency Management Agency (FEMA), a governmental organization who provides first response in large sudden onset disasters in the United States.

We develop a stochastic linear program to capture multiple tiers of carrier capacity and the different terms and conditions from pre-contracted carriers and the spot market to characterize in more detail the impact of a carrier portfolio onto an organization's response capability.

This work strives to provide new metrics to capture transportation capacity and to evaluate the status-quo of the network (system assessment) and provide insight into potential routes to improve performance (system optimization) without forcing the decision-maker to go into too much operational details. We develop metrics that allow a decision-maker to evaluate (i) how efficient the current carrier contracts are, (ii) where to renegotiate carrier contracts, (iii) which items an organization should buy, (iv) where the organization should put the items, and (v) if an organization should transfer items within its network.

The model demands five different input categories. We consider a risk portfolio, i.e. a set of disasters that could happen, the network's inventory portfolio, i.e. the organic inventory owned by the organization, the supply portfolio, i.e. the contracts in place with suppliers of disaster relief items, and the network's carrier portfolio, i.e. the capacity, terms and conditions negotiated with carriers at each depot and each supplier. In addition, we use item specific information to ensure that we capture the reality correctly.

The rest of this report is organized as follows. In the next section we introduce the stochastic linear program (SLP) and explain the aggregate metrics that we suggest to use to evaluate carrier contracts and inventory in a disaster response network. In section 5 we explain the inputs to our model and which data sources we use. Section 6 is dedicated to the results. We conduct a system assessment and a system optimization to highlight the capabilities of the model. In section 7 we discuss the model results and some extensions, how we used the model to engage with FEMA and other stakeholders and also indicate a pathway of going forward.

4 Model

4.1 Model introduction

To characterize the ability of a disaster response organization such as FEMA to respond to sudden onset disasters in the United States we create a stochastic linear program. We consider five types of inputs to our model (Figure 1). The risk portfolio characterizes potential disasters the organization can face. It characterized the source of uncertainty of when a disaster of certain magnitude needs to be served. The inventory portfolio captures the capability to respond to a disaster with specific relief items that the organization owns. We refer to this inventory as organic inventory. The supplier portfolio represents contracts the organization has in place with relief item suppliers that amend the organization's response capabilities. The carrier portfolio characterizes the organization's capability to transport relief items to disaster sites and therefore determines cost and time to respond to disasters. In addition, we use item and user specific information to accurately capture user preferences and need. Before we describe where we obtain data for these inputs (Section 5) we explain the model and the output metrics.

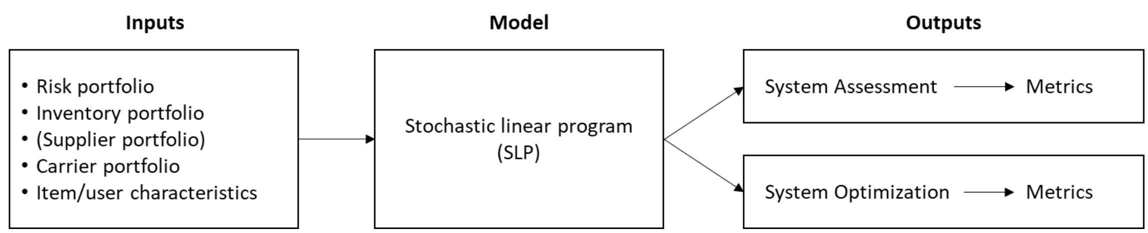

Figure 1: Inputs and outputs to the stochastic linear program (SLP)

4.2 Model formulation

Our model builds on prior research of Acimovic & Goentzel (2016). We extend their model by including carrier capacity that emergency response organizations have in place to respond to a disaster.

Before formulating the model, we define all necessary parameters and variables.

$I^O \ni i^O$ – Set of all depots with organic inventory owned by the emergency response organization.

$I^S \ni i^S$ – Set of all inventory locations with contracted supply.

i^w - The dummy supply node.

$I \equiv I^O \cap I^S \cap i^w$ – Set of all depots available for disaster response including dummy node.

$K \ni k$ – Set of possible disaster locations.

$R^S \ni r^S$ – set of spot market carriers

$R^C \ni r^C$ – set of carriers the organization has pre-negotiated contracts with

$R \equiv R^C \cap R^S$ – the set of all carriers

j^k – Staging area *j* of a disaster in scenario *k*.

$\tau_{i,j}^{k,r}$ – Time/cost to respond from depot *i* to disaster location *j* in disaster scenario *k* with carrier *r*.

$cap_{i,r}$ – Transportation capacity available at depot *i* from carrier *r*.

p^k – Probability of occurrence for disaster scenario *k*. $\sum_k p^k = 1$.

$y_{i,r}^k$ – inventory send from location *i* to supply disaster scenario *k* via carrier node *r*.

y_{r,j^k}^k – inventory send via carrier node r in disaster scenario *k* to supply staging area j^k.

$\boldsymbol{X}$ – The *I*-dimensional vector of starting inventory in each depot X_i.

χ – Starting inventory in the system as a whole, not including the dummy node.

TAP^k – Total affected population in disaster scenario *k*.

β – Units of item demanded per person.

$d^k \equiv \beta\, TAP^k$ – Demand generated by the affected population in disaster scenario *k*.

Next we formulate our basic model:

$$V(\boldsymbol{X}) \equiv \min_y \sum_k p^k \sum_{i \in I,r} \tau_{i,j^k,r} \cdot y_{rj^k}^k \tag{1}$$

s.t.

$$\sum_r y_{rj^k}^k = d_j^k \quad \forall k, j \tag{2}$$

$$\sum_r y_{ir}^k \leq X_i \quad \forall\, i \in I, k \tag{3}$$

$$y_{ir}^k - \sum_j y_{rj^k}^k = 0 \quad \forall\, i \in I, R, k \tag{4}$$

$$y_{ir}^{k} \leq cap_r \quad \forall\, i \in I, R, k \tag{5}$$

$$y_{ir}^{k}, y_{rj^k}^{k} \geq 0 \quad \forall\, i \in I, R, k \tag{6}$$

The objective function (1) minimizes the expected time-to-respond to potential disaster scenarios by calculating the product of probability of occurrence and the time-to-deliver relief goods. Constraints (2) ensure that demand at any disaster location in any scenario is met. The model ensures this by serving demand that cannot be served from organic inventory, i.e. is from depots of the organization, is served from the dummy node. Constraints (3) enforce that the volume shipped from any depot does not exceed the available inventory. Constraint (4) is a flow balance for the carrier nodes at each depot. Also the volume shipped from any location via a carrier cannot exceed the carriers transportation capacity (5) and is nonnegative (6). The model to minimize cost is similar to the one above by substituting time-to-respond for cost.

The measure $V^{w}(\boldsymbol{X})$ represents the expected time (cost) the organization incurs to serve disasters with the available inventory. To optimize the allocation of inventory across the existing network we make $\boldsymbol{X}$ a decision variable and ensure that the entire current inventory χ is used.

$$V^{Opt,w}(\chi) \equiv \left\{ V^{w}(\boldsymbol{X}) : \sum_{i \in I} X_i = \chi, X_i \geq 0 \;\forall\, i \in I \right\}.$$

For later analyses we define the objective values with dummy costs subtracted as

$$V(\,) \equiv V^{w}(\,) - \sum_{k} p^{k}\, \tau_{i^w, j^k}^{k,r}\; y_{i^w}^{k}.$$

4.3 Metrics

Since our model is based on Acimovic & Goentzel (2016) we draw in part on the metrics the authors previously developed to evaluate an organization's response capacity.

4.3.1 Balance metrics

The inventory balance metric

$$\Delta^{I} = \frac{V(\boldsymbol{X})}{V^{Opt}(\chi)}$$

compares the objective function value (corrected for the impact of the dummy variable) of the current inventory allocation with the optimal allocation of inventory. Hence, the balance metric measures the relative benefit of redistributing the inventory across the network similar to the work of Acimovic & Goentzel (2016) and its deterministic version in Acimovic & Graves (2015). That is, a value $\Delta^I > 1$ indicates an out-of-balance state of the current inventory distribution and a positive benefit of redistribution efforts.

The carrier contract metric

$$\Delta^C = \frac{V^S(\boldsymbol{X})}{V^C(\boldsymbol{X})}$$

compares the objective function value (corrected for the impact of the dummy variable) when procuring only the carrier spot market to the (dummy-corrected) objective value when using established carrier contracts and the spot market. Similar to the inventory balance metric, the carrier contract metric, therefore, measures the relative benefit of the established carrier contracts. Note, that we can evaluate the carrier contract metric for the current inventory allocation ($\boldsymbol{X}$) or for the optimal inventory allocation (χ).

4.3.2 Service metrics

We consider two major service metrics. The weighted-fraction-of-disasters-served-completely metric (δ) and the fraction-of-demand-served metric (γ) characterize the level of service the emergency response network provides.

The weighted-fraction-of-disasters-served-completely metric (δ) represents the ratio of disasters that were served completely from inventory

$$\delta = \sum_{k:d^k \leq \chi} p^k.$$

The δ-metric is robust to outliers in demand. Whether a disaster's demand is not met by only 1 unit or 100.000 units does not change the δ-metric. The δ-metric therefore provides a sense of the percentage of very large disasters relative to the network's response capacity. A δ-metric close to 1 indicates very few relatively large disasters that the emergency response network was not able to serve completely.

In addition to the δ-metric the fraction-of-demand-served metric (γ) provides a decision-maker with a sense of how much of the demand was served. We calculate the weighted average demand $\mu = \sum_k p^k d^k$, a measure that is independent of the network's response capability, and the weighted average demand met $\mu' = \sum_k p^k \min(d^k, \chi)$, to find the γ-metric

$$\gamma = \frac{\mu'}{\mu}.$$

The γ-metric is more sensitive to outliers in demand. Missing only 1 unit of demand does impact the γ-metric far less than missing 100,000 units. As such the δ-metric and the γ-metric together capture different facets about the emergency response network's capacity to serve the affected population.

4.3.3 Efficiency metrics

To provide decision-makers with a sense of the network's efficiency, we calculate the average time (or cost) to deliver one unit from a depot as

$$\varphi = \frac{V(\boldsymbol{X})}{\mu'}.$$

The φ-metric represents an aggregate approximation of the time (or cost) to deliver items in a configuration of a response network. Lower values indicate high efficiency.

4.3.4 Carrier Utilization

We characterize the carrier utilization as

$$\rho = \sum_k p^k \frac{\sum_r y_{ir}^k}{\sum_r cap_r} \quad \forall\, r \in R^C$$

to capture how well the carrier contracts are aligned with the inventory at each location. The ρ-metric is provided on an aggregate level for the entire network to indicate how well the match between inventory and contracts is overall. It is also provided on a depot level (ρ_i) to show where lower and higher utilizations are located within the network.

In general a high utilization (close to 1) indicates that a lot of inventory is moved compare to the available contracts. Whereas a lower utilization is indicative of less volume moved compared to the inventory.

4.3.5 Dual variables

The stochastic linear program (SLP) provides dual variables for the constrained resources, that is, for inventory levels and carrier capacity. Dual variables indicate how much increasing a specific constraint increases (or decreases) the objective value of the SLP. Therefore, the dual variables quantify the value of procuring an additional unit of scarce resource to improve the objective function value. Hence, they are sometimes referred to as shadow prices.

We obtain dual variables for the depots i

$$\pi_i' = \sum\nolimits_k \pi_i^k - (1-\delta)\tau^w$$

and for each carrier (tier) r

$$\pi_r' = \sum\nolimits_k \pi_r^k - (1-\delta)\tau^w.$$

We adjust the dual variables $(\pi_i^k, \pi_r^k)$to correct for the impact of the dummy weight τ^w on the objective value, which is rather arbitrary, and may distort the value of the duals and therefore make an interpretation difficult.

A positive carrier dual variable shows the reduction in expected time (cost) if we increase the capacity of a specific carrier (tier) by one unit. Because we assume there is always sufficient overall capacity at each depot from the spot market, carrier duals will always be non-negative and the highest positive carrier dual variable promises the highest benefit from increasing capacity.

A positive depot dual variable shows the reduction in expected time (cost) if we increase the capacity of a depot by one unit. Typically, we will have less inventory available than the largest demand in the risk portfolio. Because adding one unit will result in delivering one additional unit in one of the disaster scenarios, the expected time (cost) will typically increase if we increase the inventory. Therefore, in most cases the depot dual variables will be negative indicating that increasing inventory results in higher expected times and the smallest (negative) depot dual variable promises the highest benefit (lowest negative impact) from increasing inventory.

Note, that this interpretation of the depot duals variables directly ties into a change to the service metrics. Ceteris paribus, increasing inventory at a depot results in satisfying additional demand from organic inventory. Although more units delivered increase the (adjusted) objective function value, the service metrics will improve.

5 Data

5.1 Demand data

We use historical data to create a disaster risk portfolio. We consider disasters in the continental United States from January 1990 until June 2018 and assume that any of these disasters has an equal chance of occurring.

To acquire disaster information we draw on the open FEMA API (FEMA (2018)), an exhaustive collection of US disasters. We collect from open FEMA the information on disaster type, location, and date of occurrence. We focus on sudden onset disasters in the continental United States. Therefore, we include the disaster categories Costal Storm, Earthquake, Fire, Flood, Freezing, Hurricane, Mud/Landslide, Severe Ice Storm(s), Severe Storm, Tornado, and Tsunami in our analysis. We exclude any disaster sites outside the continental United States (e.g. Hawaii, Alaska). Figure 2 shows the number of occurrences of disasters in the past 28 years for each county in the continental United States. The map highlights that there are geographical differences in how strongly certain regions are affected by disasters. More disasters occur in the west, south east, Midwest, and north east of the United States.

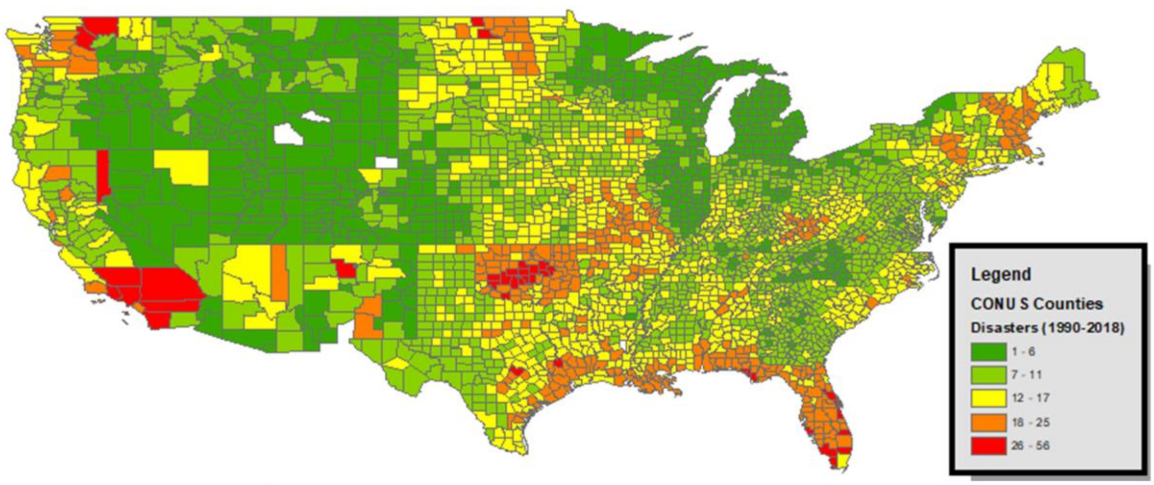

Figure 2: Number of disaster occurrences per United States county 1990-2018

Estimating the number of people affected by a disaster is very difficult because accurate numbers are not public information. We estimate number of people affected by drawing on a method that FEMA uses for internal planning purposes and assume that 26% of the population of an affected area (e.g. county or state) need assistance. We use census data from 2010 to estimate the population numbers (US Census (2010)).

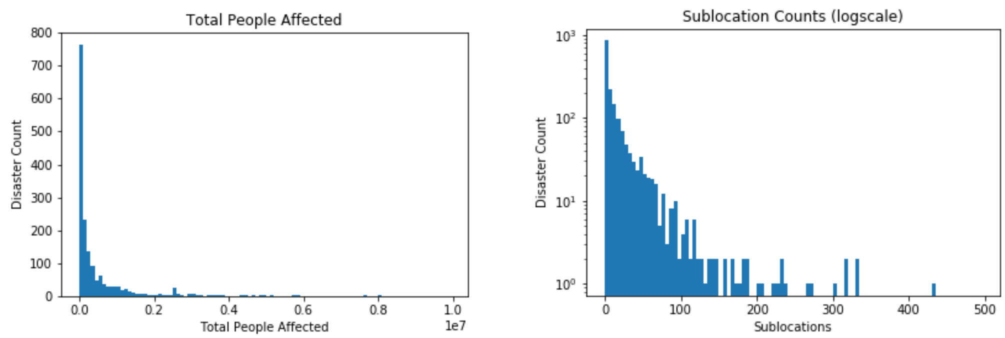

Figure 3: Total people affected per disaster (left) and number of sub-locations per disaster (right)

The data collected from the FEMA API also provides information about the affected areas (such as counties or other political regions) at each disaster. That is, for a disaster the data indicates which areas are affected at the same time. In most cases the areas (sub-locations) refer to US counties. For cases where a reference is given other than county we approximate by finding the corresponding county (less than 1 percent of all entries). We use the sub-locations as a place holder for multiple staging areas. Figure 3 shows that many disasters in our risk portfolio are rather small and that there are fewer disasters that affect millions of people. In Table 1 we report summary statistics for the disasters. We also see that most disasters are geographically limited indicated by only a few counties affected. But the disaster portfolio also includes some instance where hundreds of counties (sub-locations) are involved suggesting large regions were affected. In Table 1 we report also summary statistics for the number of counties per disaster.

	Disaster	County
Count	1,731	34,717
Minimum	168	1
Maximum	66,696,256	2612
Average	634,444	20
Median	138,359	5
1st Quartile	31,058	1
3rd Quartile	530,573	18

Table 1: Summary statistics for the disasters and the counties in the risk portfolio

5.2 Inventory data

Our model evaluates the emergency response capacity to react to sudden onset disasters and serve the affected population with critical relief items. FEMA seeks to supply the affected populations with disaster relief items that allow people to survive in the aftermaths of a disaster. That is, FEMA seeks to supply people in the first 72h after a disaster to bridge the gap to longer-term support from other organizations.

We assume that FEMA operates five warehouse locations strategically positioned across the continental United States (see Figure 4) in Philadelphia (Pennsylvania), Washington D.C., Atlanta (Georgia), Dallas (Texas), and San Francisco (California). In these locations, FEMA holds critical relief items to support people in case of sudden onset disasters. Our analysis concentrates on the following categories: bottled water, shelf-stable meals, cots, tarps, and blankets.

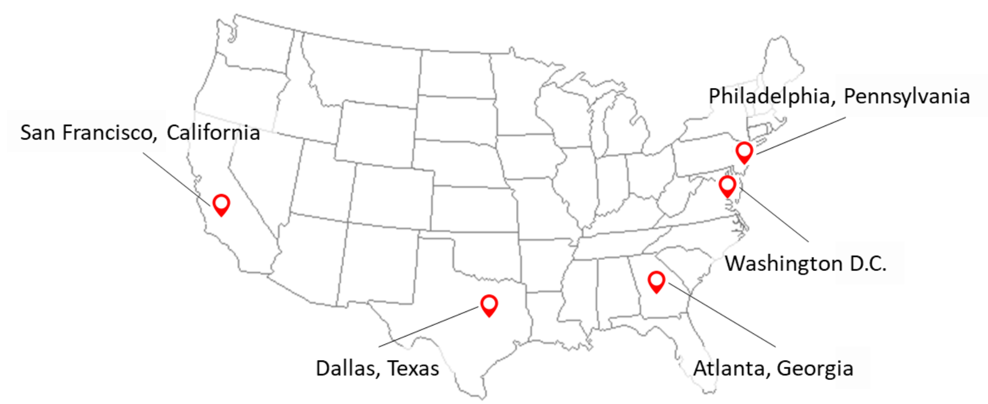

Figure 4: FEMA warehouse locations across the continental United States

Our model evaluates the current state of inventory to conduct a system assessment. For this we are currently in the process of collecting past FEMA inventory levels of the items of interest at important instances in time (for example, the currently held organic inventory, or the inventory held prior to hurricane Sandy in 2012). Different instances of time allow us to conduct scenario analyses and compare the response capability of different inventory portfolios.

To show the capability of the model to analyze the status-quo of a system and provide recommendations for improvement (see Section 6) we use nominal data for inventory levels in Table 4 in the appendix. We elaborate on how we already have collected feedback from FEMA and how we will further engage with FEMA to conduct further validation and implementation of our model (see Section 7).

5.3 Carrier data

We use information on an organization's contracts with carriers at each warehouse location to model transportation capacity. In case of a disaster, an organization such as FEMA may draw on three different tiers of carriers. First, they call on trucks that pick up pre-loaded trailers at the depot locations. After that, an organization uses two options to transport the disaster relief items. On the one hand, they draw on pre-arranged contracts with carriers that report to depots in a pre-defined response time, load the items onto the truck and sent them out to the disaster sites. On the other hand, they procure carrier capacity from the spot market.

Estimating the carrier capacity is not easy because in addition to the time it takes carriers to arrive at the warehouses, the loading dock capacity is limited so that if many trucks arrive at the same time additional waiting time is incurred.

To avoid modeling a detailed arrival and loading process we estimate the carrier capacity available at each warehouse by building tiers of carriers that are characterized by the number of trucks in each tier, a fixed time, a variable time and variable cost. The fixed time is an expert estimate of how much time it takes on average to make a specific number of trucks available at the warehouse and load the items onto the trucks.

The variable time reflects the travel time between a depot and a specific staging area of a disaster. To estimate the distance between warehouses and the staging areas we draw on Google's distance matrix API (Google (2018)). For given latitude and longitude of depot locations and staging areas the API returns the distance between the two locations (in kilometers) and the travel time on the most common route. The travel time directly enters into our time model so that

$$\tau_{i,j}^{k,r} = f^r + v_{i,j}^{k,r}$$

where f^r is the fixed time to make the carrier capacity r available and $v_{i,j}^{k,r}$ is the variable travel time between a depot and the disaster site.

In case of the cost minimization model we use the distance returned from the Google API for each relation (in kilometers) and multiply it by the transportation cost parameter (per kilometer and kilogram) reflecting an organization's contracts with the carrier. Because we are still in the process of validating our model with FEMA representatives, we collect estimates of transportation cost parameters for contracts and for the spot market from DAT (DAT (2018)). Our estimates for the carrier tiers at each depot, the fixed time, and the variable cost can be found in the Appendix in Table 7.

In Figure 5 we present aggregate measures to provide an overview of the carrier portfolio in our analysis. It shows the total carrier capacity at each depot (excluding the spot market), the average fixed time to ready the carrier capacity, and the average variable cost. With 200 trucks available Washington offers the highest capacity in the carrier portfolio and the smallest number of trucks is available in San Francisco and Philadelphia. On average Atlanta's carriers offer the highest fixed time to provide the capacity at the lowest cost. On average carriers are most responsive but also most expensive in San Francisco.

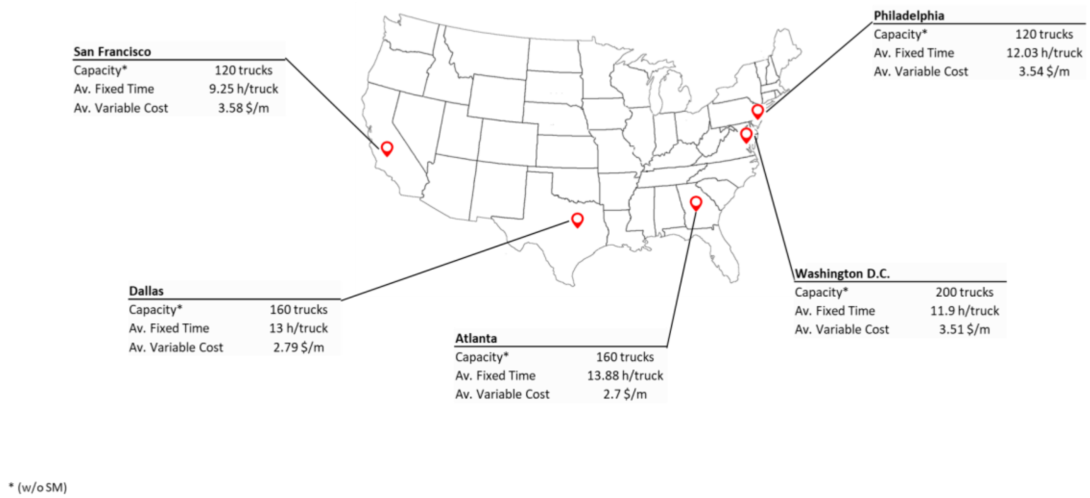

Figure 5: Carrier portfolio

5.4 Item specific information

To capture the reality appropriately we also collected item specific information.

To estimate how many affected people FEMA can serve with one item we estimate conversion rates. In our model we use the item conversion rate β to convert the total affected population (*TAP*) in unit demand. We consulted experts in the field to estimate these conversion rates. We assume that on average within the first 72h of a disaster 9 bottles of water serve a person, 6 units of shelf-stable meals serve one person, every third person needs a cot, a cot needs two blankets, and 0.3 people need a tarp.

Furthermore we collected information on item specific weights and volume (see appendix Table 6). Item weights are important to calculate the variable transportation cost $\tau_{i,j}^{k.r}$. Since carrier rates r are given in price per kilometer per kilogram we multiple the rate by the item specific weight w and the distance between a depot and a disaster site $l_{i,j}$ to obtain transportation cost so that in the cost model

$$\tau_{i,j}^{k.r} = r \, \mathrm{w} \, l_{i,j} \, .$$

Carrier capacity is given as the number of trucks available at each depot. Because different items require different room on a truck we need to convert the number of trucks into an item-specific, i.e. unit-based, transportation capacity. As our items are rather low in weight we do not exceed the weight limits of trucks in the United States. Rather the volume is the limiting factor. We approximate the loading volume of a typical 5-axle tractor semi-trailer with 106 m^3 (3743 ft^3). Dividing this volume by the specific volume of an item and giving 10% for tolerance for odd shaping and other factors we estimate the item-specific capacity.

6 Results

In this section, we present the results of our analyses. Our model allows evaluating the status-quo of an emergency response organization's current network (system assessment) and the optimization of the network to improve performance (system optimization). We first present our result of system assessment in Section 6.1 and then move to optimize the system in Section 6.2.

It is important to note that at the time of this report we have not yet validated the inventory and carrier data. The reported results are based on estimates and will be validated in workshops with FEMA in the coming weeks. In particular, the values for inventory levels and carriers are NOT representative for FEMA and should not be considered as a reflection of FEMA's response capability in the past or today.

6.1 System assessment

First, we evaluate the state of the system for each item that is of interest. If not stated otherwise we report our results in units and not in the number of people affected. Table 2 presents the results and shows the average demand (in units) and the average demand met. We also report the δ-metric, the γ-metric and the average time and cost to serve.

Item	Demand (μ)	Demand met (μ')	Fraction of demand served (γ)	Fraction of disasters served completely (δ)	Average time to serve (Φ)	Average cost to serve (Φ)
Water	5,715,795	3,641,791	63.71%	91.1%	31.25h	$0.42

Table 2: Summary metrics for the time-based model

The results of the system assessment allow a decision-maker to evaluate the network's ability to respond with the current inventory (Figure 6) and carrier portfolio against the risk portfolio on an aggregate level. For example, the service metrics for water bottles suggest that the organization is able to serve 63.71% of the demand and that it serves 91.1% of the disasters completely with the current inventory portfolio and the carrier portfolio. The efficiency measures indicate that on average, it takes the organization 31.25h to ship water bottles to disaster sites at average cost of $0.42 per bottle.

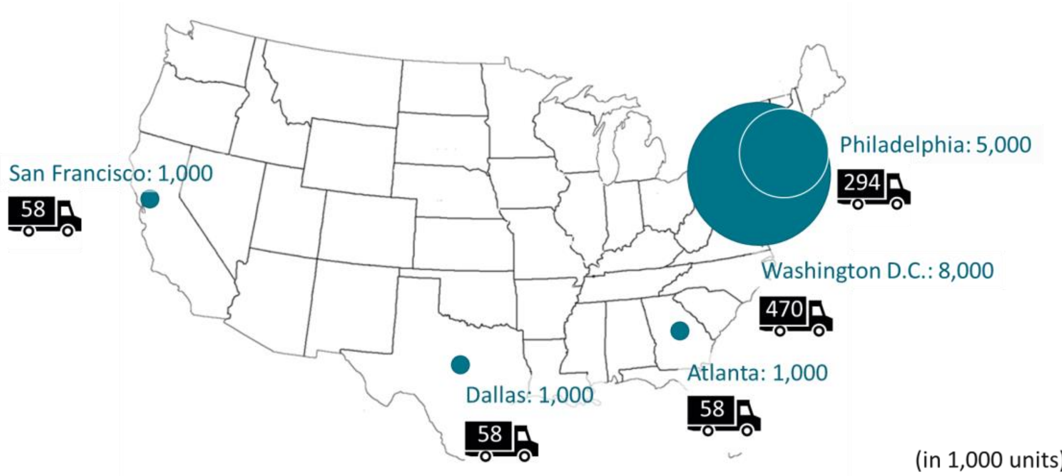

Figure 6: Current inventory portfolio for water and the corresponding number of full truck loads

In Figure 7 we report the fraction of demand served over time. The results show how responsive the network is to serve demand for a given inventory and carrier portfolio. Figure 7 suggests that the network generally needs around 8h to start shipping out the pre-loaded trailers to the destinations and that approximately 55% of the demand is served after the first 48h.

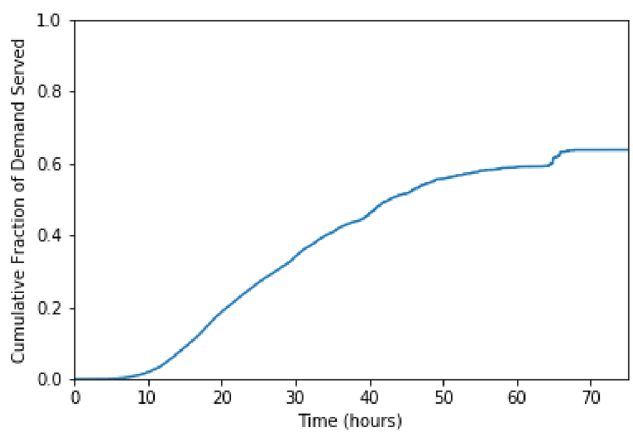

Figure 7: Fraction of demand served (γ) over time in the time-based model for cots

The demand served over time metric in Figure 7 provides the decision maker with a sense of (average) responsiveness to disaster portfolio. The result can be discussed and evaluted against organization's response targets to better understand if the current state of the system is sufficiently responsive or if steps to improve the response capacity should be considered.

Our model also allows to identify the value of the prenegotiated contracts with carriers. The carrier contract metric for the current inventory and carrier portfolio is $\Delta^{C} = 1.4$. This result suggests that relying only on the spot market – instead of using the carrier contracts and the spot market – results in 40% higher response time. That is, having carrier contract in place is extremely valuable for the organization.

The metrics provided in the system assessment provide insight into the state of response capacity of a system and evaluate if the current setting is able to meet the expectations and targets. However, our method also support a decision-maker on how to optimize performance. For this we conduct a system optimization in the next section.

6.2 System optimization

In the previous chapter we performed a system assessment that showed the emergency response network's capacity with the current inventory and carrier portfolio to meet a given disaster risk profile. In this section, we perform a system optimization by re-distributing the entire inventory available to the emergency response organization to minimize delivery time (or cost) given the contracts with the carriers, the spot market availability, and the risk portfolio.

Figure 8 compares the optimal distribution of water bottle inventory across the network. The results suggest that it is optimal to distribute the inventory more evenly across the network and put more inventory in the Dallas depot to account for the large number of disasters in the Midwest and south of the United States (compare Figure 2).

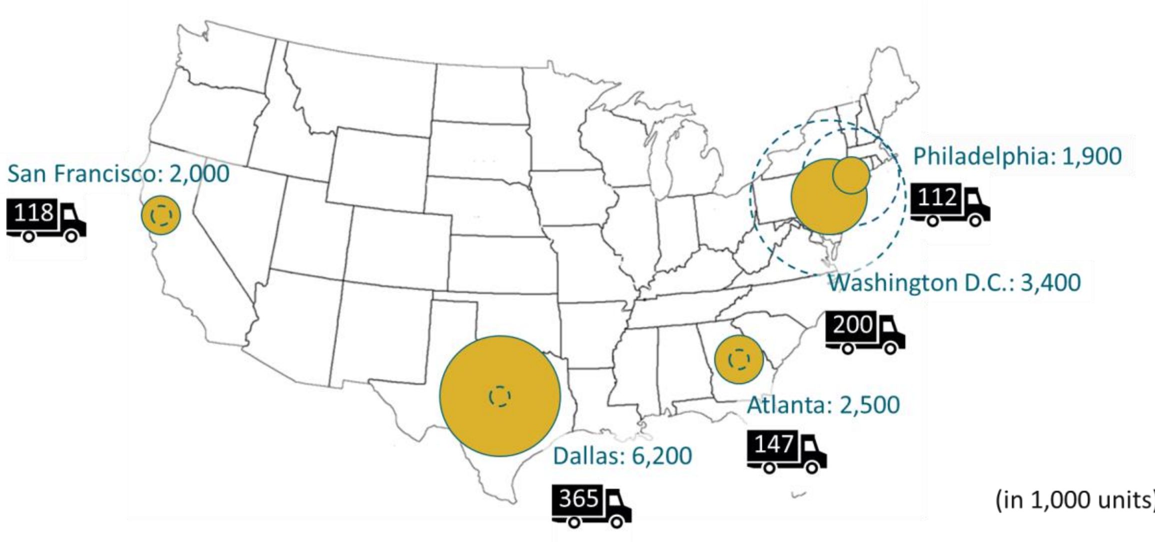

Figure 8: Actual (dashed) and optimal (yellow) allocation of water bottle inventory in the time-based model and full truck load equivalent

Because we do not increase the entire inventory in the network the service metrics remain the same. However, the network becomes more responsive. Table 3 shows the key metrics for the time-based and the cost-based model. The inventory balance metrics indicate that expected time and expected cost to serve decrease relative to the system assessment suggesting that inventory allocation is indeed beneficial. Also the efficiency metrics improve as the same volume can be shipped at less time or cost, respectively.

Item	Inventory balance metric (Δ^I)	Carrier contract metric (Δ^C)	Average time to serve	Average cost to serve
Water	1.143	1.45	**27.3h**	$0.26
	1.6785	1.84	29.65h	**$0.24**

Table 3: Summary metrics for optimization w.r.t. to time and cost (value being optimized is bold)

Figure 9 shows the γ-metric over time. The results suggest that when optimally allocating the inventory the network is able to respond to about 60% of demand within 48h. That is, we increase the response time to 12% (60% instead of 55%) of the population that we can reach with our inventory. In particular the network becomes more responsive to the larger disasters that take longer to fulfill and particularly the affected population, that had a long waiting time in the current state, is served earlier. This is very important for the disaster response with essential disaster relief items (water-bottles, food, blankets, tents, cots) in the first 72h because reaching the affected population within two instead of three days can make the difference between life and death.

The carrier contract metric for the optimal inventory is $\Delta^C = 1.45$. This indicates that the carrier contracts in place are more valuable in a network with an optimal inventory allocation compared to the carrier contract metric in the system assessment ($\Delta^C = 1.4$). This is because the optimal inventory allocation also takes the carrier contracts at each depot into consideration and seeks to utilize them best. More balanced utilization allows for a better response to disasters.

The service improvement in time is therefore not only because inventory is redistributed and strategically positioned closer to the disaster sites. It is also due to a better utilization of available carrier capacity. Previously, more inventory was located in depots in Washington and Philadelphia without corresponding higher carrier capacity available (see Figure 8). Redistributing the inventory to account for free carrier capacity at other depots allows the organization to improve service.

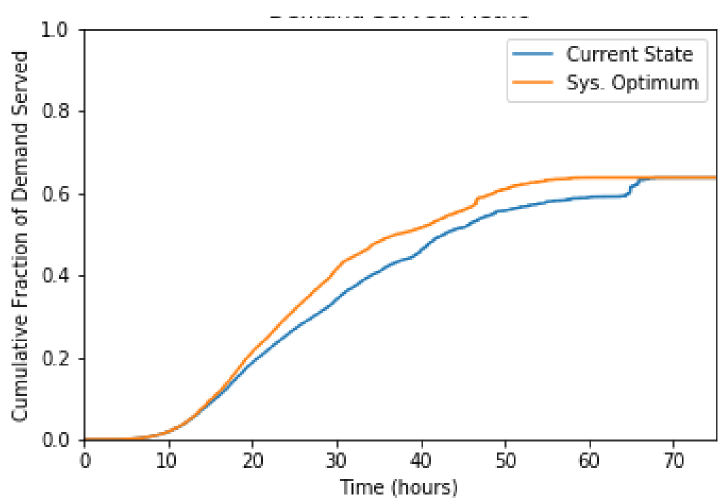

Figure 9: Fraction of demand served (γ) over time in the time-based model for bottled water for the current inventory (blue) and the optimal inventory distribution (yellow)

7 Discussion

7.1 Discussion

The objective of this research was to provide support for decision-makers to evaluate the current state of an emergency response network and suggest areas for improvement. We developed a stochastic linear program that uses a risk portfolio, an inventory portfolio and a carrier portfolio as inputs. We calculate aggregate performance metrics to evaluate the system. Our metrics enable a decision-maker to evaluate the status-quo of a disaster response network and compare these metrics to optimized inventory across the network minimizing either time-to-deliver or cost-to-deliver.

To assess the system, our model reports service metrics. The model calculates how well an emergency response organization can satisfy demand (γ-metric) and how many disaster of a risk portfolio can be met completely (δ-metric). Furthermore, our model allows to calculate a measure of responsiveness – the γ-metric over time – and thereby sheds insight into how timely after a disaster the network responds. Our model also reports the average cost and the average time per unit as a measure of efficiency. These metrics taken together provide insight into how well the network can meet the demand of the affected population in its current setup.

Our model can minimize either expected time-to-deliver or the expected cost-to-deliver. This allows to consider the inherent trade-off of an emergency response network. Delivering in a short period will be very expensive whereas delivering at low cost will result in long delivery-times. The outputs allow a decision-maker to balance cost and time to find a trade-off that aligns with the organization's objectives.

In our work, we are particularly interested in analyzing in more detail the role of contracts with carriers in the network. Clearly, the cost-time trade-off is an important part. Our model also allows us to evaluate the network's responsiveness by plotting the fraction of demand served (γ) over time. This time-based metric sheds light onto how rapidly the network begins to deliver and how long it takes to reach a certain percentage of the affected population. Comparing the γ-metric over time of the system assessment and the system optimization indicates the improvements in network responsiveness through optimal allocation of inventory.

We also include a carrier contract metric that captures the value of pre-negotiated contract to a reference case. For this study we decided to use a less responsive and more expensive spot market as the reference point. We showed that the value of pre-negotiated contracts can be quite high. This supports the idea of this research project of modeling specifically the response networks transportation capacity as the detailed set-up and availability of carrier capacity largely determines an organizations response within the critical initial 72h of a disaster. We can show that if carrier contracts and inventory allocation are not aligned the response of the network is slower and that carrier contracts become more valuable if inventory is optimized. To shed further light into the carrier capacity we also introduced an aggregate and a per-depot carrier utilization rate to flesh out mismatches between pre-negotiated contracts and inventory allocation.

The system optimization not only presents metrics of how service, responsiveness and efficiency change. The results also explain where, i.e. in which depots, a decision-maker should position the inventory to reach the optimum.

Our model also offers insights into the emergency response organization's contracts with carriers. We can extract the adjusted dual variables for the carrier capacity constraints from the SLP. The adjusted duals represent the benefit of increasing the capacity of a carrier by one unit. They provide answers to the question: At which depot and with which carrier should the organization buy more carrier capacity?

Figure 10 shows the adjusted carrier duals for the time-model in the water bottle example. It shows the value of the dual for the eight carrier tiers at each depot in the disaster response network. The first tier at a depot represents pre-loaded carriers in our model, the last tier is the spot market available at a depot, and in between are contracts the organization has established with individual carriers.

Carrier duals are all positive suggesting that expanding carrier capacity at these depots reduces the expected time to deliver. More responsive tiers at these depots have the highest impact on delivery time and should therefore be prioritized in case of carrier capacity adjustments at these depots. That is, contracting more capacity for pre-loaded trailers (C1) results in the highest reduction in expected time whereas the first contract carrier tiers provide lower benefit (C2, C3, etc.). We can also see that higher contract tiers at depots with low inventory (San Francisco, Texas, Atlanta) do not substantially contribute to lowering the expected time. These tiers are never used because of the low inventory and expanding them does not benefit the decision-maker. Rather a decision-maker should consider reducing capacity at these tiers. Whereas lower carrier tiers at depots with a lot of inventory (Philadelphia, Washington D.C.) still contribute to lowering the expected time if capacity is increased.

Figure 10: Dual variables for carrier capacity to minimize delivery time of bottled water

Similarly, we can obtain duals for the cost-based model that can be used to better understand at which depots and with which carriers, capacity should be increased to reduce expected delivery costs.

To summarize, the carrier dual variables in the time- and in the cost-model complement and extent the network results of the network assessment and optimization. The system optimization (in Section 6.2) provides intuition into where a decision-maker should move inventory considering – but not changing – the available carrier portfolio, i.e. the entire transportation capacity available at each depot and the responsiveness of each tier. Whereas the analysis of carrier duals provides insights into where a decision-maker should increase carrier capacity considering the terms, i.e. fixed time to respond and the carrier rates, in the carrier contracts. Accordingly, the duals complement the insights from network optimization and provide suggestions which carriers in an organization's carrier portfolio are promising candidates to expand carrier capacity.

7.2 Extensions and future work

Within this project, we have developed a new model to more accurately capture carrier capacity of a (national) disaster response network. We have shown the feasibility of our approach and provided insight into the capabilities of our model. We held a half-day workshop with 20 FEMA decision-makers from different FEMA departments to validate our approach. We learned that our model is very valuable to their problem setting. They intend to conduct a system assessment and optimization workshop in Q1/2018 to evaluate their current network and options to expand their inventory portfolio and the way they work with carriers. For this workshop we will collect inventory- and carrier-specific data from FEMA at different points in time to conduct scenario analyses and show how different settings change service, responsiveness, and efficiency metrics of the network.

We also conducted a live video workshop for the MicroMasters in Supply Chain Management students at the MIT Center of Transportation and Logistics' CAVE lab.[1] We invited two FEMA colleagues to shed light into how the model can help improve FEMA's response capacity and support decision making. The video has been accessed by 800 viewers at the time of this report and contributes to educating supply chain decision makers in more than 190 countries.

The workshop also inspired a discussion that FEMA seeks to conduct a workshop and invite participants of other response organizations and the private sector, who also provides response capacity in case of a disaster, to evaluate the entire response network comprising many organizations in a dynamic workshop in MIT CTL's CAVE lab. Our model can enable this discussion by modeling multiple organizations' assets against a risk portfolio and evaluate where assets should be increased best to improve response time.

Finally, we embark this month on a project with USAID/OFDA to analyze their global response network. We intend to marry the international response network analysis performed by Acimovic & Goentzel (2016) trading of different modes of transportation and the extension for a single-mode, multiple-carrier-tiers setting to analyze a national response network performed in this next research project.

7.3 Limitations

Our work considers each item individually. This approach is reasonable because emergency response organizations typically consider the items independently and ship them out as full-truck loads. We make this assumption because a holistic model, that considers all critical items at once, becomes too complex and its results – if computationally achievable – too intricate to interpret to provide benefit for practitioners.

When interpreting our results, one has to consider that to separate each item and individually assess or optimize the network we have assume that items do not compete for carrier capacity. Essentially, we allocate each carrier (tier) of the entire contract to a specific item and run our model. However, in a disaster the allocation of items to trucks may be dynamic and dependent on local needs that we cannot anticipate with our model.

[1] Video of the live event can be accessed here: https://www.youtube.com/watch?v=rjt9ivoiJQE

8 Literature

Acimovic, J. & Goentzel, J. (2016) Models and metrics to assess humanitarian response capacity, Journal of Operations Management (45), 11- 29

Acimovic, J., Graves, S., 2015. Making better fulfillment decisions on the fly in an online retail environment. Manufacturing & Service Operations Management 17 (1), 34 - 51.

DAT (2018) DAT, https://www.dat.com/company

FEMA (2018) open FEMA API, https://www.fema.gov/openfema

Google (2018) Google Distance Matrix API, https://developers.google.com/maps/documentation/distance-matrix/start

US Census (2010), US Census 2010, https://www.census.gov/programs-surveys/decennial-census/

9 Appendix

Item (units)	Philadelphia, Pennsylvania	Washington D.C.	Atlanta, Georgia	Dallas, Texas	San Francisco, California
Blankets	17,000	270	34,000	37,000	40,000
Cots	26,000	28,000	13,000	24,000	97,000
Meals	4,495,000	3,540,000	1,010,000	5,450,000	5,460,000
Tarps	16,000	47,000	22,000	23,000	27,000
Water	5,000,000	8,000,000	1,000,000	1,000,000	1,000,000

Table 4: Estimated actual inventory distribution for disaster response items

Units/TAP	Blankets	Cots	Meals	Tarps	Water
Conversion Rates (β)	0.666	0.333	0.167	0.3	0.111

Table 5: User item conversion rates

Units/TAP	Blankets	Cots	Meals	Tarps	Water
Weight (kg)	0.00123357	0.0027	0.00022	0.00123357	0.001
Volume (m3)	0.01045	0.5928	0.000286	0.01045	0.001

Table 6: Item specific weights and volumes

Depot	Carrier tier	01	02	03	04	05	06	07	08
Philadelphia	Capacity (# trucks)	16	24	20	20	16	14	10	∞
	Fixed Time (h)	4	9	11	14	15	16	20	26
	Var. Cost ($/FTL)	3.6	3.8	3.5	3.6	3.4	3.5	3.1	9.1
Washington	Capacity (# trucks)	30	32	24	30	30	24	30	∞
	Fixed Time (h)	4	8	10	13	14	16	19	26
	Var. Cost ($/FTL)	3.8	3.6	3.7	3.6	3.2	3.4	3.3	9.1
Atlanta	Capacity (# trucks)	26	32	38	10	32	10	12	∞
	Fixed Time (h)	4	10	15	16	18	19	25	26
	Var. Cost ($/FTL)	3.2	3.0	2.6	2.7	2.2	2.9	2.3	7.9
Dallas	Capacity (# trucks)	20	30	42	28	10	14	16	∞
	Fixed Time (h)	4	9	10	15	20	22	24	26
	Var. Cost ($/FTL)	3.2	2.7	3.0	2.7	2.6	2.5	2.4	6.9
San Francisco	Capacity (# trucks)	20	34	32	6	14	8	6	∞
	Fixed Time (h)	4	8	9	11	13	15	17	26
	Var. Cost ($/FTL)	3.9	3.5	3.7	3.4	3.5	3.0	3.4	9.2

Table 7: Carrier portfolio