Proceedings of the 2021 IISE Annual Conference
A. Ghate, K. Krishnaiyer, K. Paynabar, eds.

A Multistage Stochastic Programming Model for Hurricane Relief Logistics Planning

Murwan Siddig and Yongjia Song
Clemson University
Clemson, SC, USA

Abstract

This paper considers a typical logistics planning problem for transporting relief items over multiple periods to improve the disaster relief efforts for an impending hurricane. Since the demand for relief items can not be known in advance, we assume that the decision-maker (DM) receives forecast advisories about the hurricane's characteristics sequentially over time. This information is used to construct a Markov chain (MC) that predicts the hurricane attributes, such as the landfall location and intensity, which are in turn used to estimate the demand. To solve this problem, we feed this MC into a fully adaptive multistage stochastic programming (MSP) model. This MSP model allows the DM to adapt their operational logistics decisions sequentially over time according to the state of the MC and the state of the logistics system in terms of the relief item inventory allocated. We test this MSP model in a case study and compare its performance to alternative deterministic models. Our preliminary numerical results show the effectiveness of this approach.

Keywords
Multistage stochastic programming. Logistics planning. Hurricane relief.

1. Introduction

Multistage stochastic programming (MSP) is a mathematical programming model for *sequential* decision-making under *uncertainty*. In a typical MSP problem, the goal of the decision-maker (DM) is to construct an optimal policy that prescribes decision-making over multiple stages of time known as the planning horizon T. To do this, in each the decision stage $t = 1,\dots,T$, the DM observes the current state of the system and influence its behavior by choosing the action which optimizes his/her objective (e.g., minimizing the total expected cost or maximizing the total expected profit). MSP models have a colorful range of applications in several areas such as energy [3], finance [5], transportation, and logistics [7], among others. In this paper, we are concerned with applying MSP models for *hurricane relief logistics planning*.

Over the past few decades, many significant hurricanes have hit the United States (US). This resulted in heavy loss of lives, numerous personal injuries, and consequential material damages. According to the National Oceanic and Atmospheric Administration (NOAA), a total of 273 hurricanes struck the US between 1851 and 2004. Since 2005 – the year in which the costliest hurricane in history, *Katrina*, made landfall in Louisiana – there have been at least 43 major hurricanes. Altogether, these major hurricanes have caused more than 600 billion dollars worth of damage in the US and the Caribbean, making hurricanes to be among the costliest and the most frequently observed types of natural disasters. Despite their tragic impacts, unlike other unpredictable natural disasters such as earthquakes, hurricanes can be detected a few days before they make landfall. Hence, the (forecast) information provided by such early detection can be leveraged by humanitarian and governmental agencies to prepare and strategically deploy physical and non-physical resources, such as first-aid commodities, food, water, housing, backup power generators, etc., in what is known as the pre-disaster logistics planning. Given the evolving uncertainty associated with the hurricane *intensity* and the potentially affected *locations* from which DMs expect requests for relief items, it is natural to consider using an MSP model to characterize and improve pre-disaster logistics planning.

In this paper, we consider a typical hurricane relief logistics planning situation for prepositioning relief commodities after a hurricane has been predicted to make landfall within a few days. From the possible locations where the hurricane is predicted to make landfall, we define an affected potential area (APA) that is at risk of being affected by the hurricane. Within the APA, there is a set of demand points (DPs) that represent the locations from which we can expect requests of relief items once the hurricane makes landfall. Such relief items can be shipped from a set of supply points (SPs), including the main distribution center (MDC). To that end, the goal of the DM is to design an *optimal* policy for prepositioning relief items for different SPs and DPs over time. In this context, the optimality of such policy is characterized by minimizing the overall logistics cost plus the cost of failing to meet the demand for relief supplies. To construct such policy, we propose an MSP model which allows the DM to *adapt* the sequential prepositioning decisions over time as new information about the hurricane characteristics arrives. We evaluate this model in a case study against several *deterministic* alternative models that serve as the benchmark policies.

The remainder of this paper is organized as follows. Section 2 reviews the literature on applying mathematical models for hurricane forecast and hurricane relief logistics. In Section 3, we present the problem description. In Section 4, we discuss the proposed MSP formulation. In Section 5, we describe the implementation details and present our proposed approach's numerical results in a case study. Finally, in Section 6, we conclude with some final remarks.

2. Literature Review

In the literature, there is a good deal of work on applying mathematical models to address the different aspects of hurricanes. This includes predicting the hurricane's intensity and track using forecasting models [e.g., 2]; making decisions about the evacuation, rescue missions, and distributions of emergency supplies [e.g., 13]; diesel fuel supply chain [e.g., 6], among others. Amongst the models most closely related to the one presented in this paper are the ones presented in [10, 11]. Here, the authors consider the problem of prepositioning relief items in preparation for the immediate aftermath of hurricane landfall. In this problem, the DM receives *sequential* forecast advisories – issued by the National Hurricane Center (NHC) every six hours – about the predicted wind speed at the time of the hurricane landfall. Then, the DM uses this information to construct a policy for prepositioning the relief items. As it is common to assume that earlier advisories have greater uncertainty, whereas logistic costs are much lower at the earlier stages of planning, the actions ought to be taken by the DM can, therefore, be separated into two layers: (i) when starting the prepositioning process; and (ii) how much relief items should be propositioned at the different predicted demand locations. In [10], the authors approach this problem as an optimal stopping problem, whereas, in [11] the authors use a combination of decision theory and stochastic programming techniques.

One major drawback to the model presented in [10] is that the authors assume that the hurricane's forecast advisories – received by the DM – are statistically independent of one another. This issue, however, is addressed in [11], where the authors model the dependency between sequential forecast advisories using a Markov chain (MC). In this context, it is assumed that the next forecast depends on the history of advisories through only the current state *only*, and not on the previous ones. Another prominent feature of the work presented in [11] is that the model uses a *dynamic* approach for prepositioning the relief items. Here, unlike the *static* model presented in [8], where the DM commits to a single prepositioning decision *only* for the entire planning horizon, this dynamic approach allows the DM to *adapt* those decisions as new information about the hurricane arrives. To construct such dynamic policy, the authors decompose the state of the system throughout the planning horizon into three possible states: (i) an *initial* state where no prepositioning has been made yet; (ii) an *active* state where prepositioning decisions can be made; and (iii) a *final* state where the estimated time until the hurricane makes landfall would not be enough to perform any prepositioning operations. At first (time $t = 1$), the system is assumed to be at the initial state. Then, for each subsequent period, the DM has to decide whether to start the prepositioning operations or wait until the next forecast advisory arrives. If the DM chooses to begin the prepositioning operations, the system evolves into the active state where the DM needs to determine the amount and locations of the relief items to be prepositioned by using the available *less accurate* forecast information but with a *lower* logistics cost. Otherwise, the DM waits until the next time period and faces the same dilemma again with a *more accurate* forecast but at a *higher* logistics cost. Finally, once the estimated remaining time until the hurricane makes landfall becomes smaller than the time required to perform any prepositioning activity, the system evolves into the final state, and no more prepositioning is made. To solve this problem, the authors do the following. First, they consider a set of scenarios S, where each scenario $s \in S$ for a hurricane occurrence is characterized by three attributes: (i) location, (ii) intensity, and (iii) time of landfall. Second, for a given scenario $s \in S$, to determine if a given time period $t_s \leq T$ is the *best* time to start prepositioning, the authors evaluate the cost of acting at time t_s compared to waiting for an additional period (i.e., time $t_s + 1$). If the cost of waiting is lower than the risk of acting, the decision of initiating the prepositioning is postponed for time $t_s + 1$. Otherwise, the prepositioning is initiated at time t_s. To measure the difference between the *cost* of acting now and the *risk* of waiting, the authors use a decision theory approach that minimizes the expected cost and the maximum regret. In this paper, we consider a similar problem of prepositioning relief supplies like the one presented in [11]. We also use the Markovian structure for modeling the underlying stochastic process to predict the hurricane's characteristics over time. However, we impose the adaptability of sequential decision-making to the arrival of new information more explicitly via an MSP model.

3. Problem Description

We consider a typical problem setting for relief supply prepositioning in disaster relief logistics planning. The underlying logistics process can be modeled as a *multi-period network flow* problem. The logistics network is modeled as a directed graph $G = (V,A)$, where the set of nodes $V = \{0\} \cup I \cup J$ consists of the MDC (denoted by node 0), a set I of SPs, and a set J of DPs. All relief items can be procured from the MDC. From there, the relief items can be propositioned at the different SPs over a planning horizon of T periods before the hurricane makes landfall at time T. Then, after the hurricane makes landfall, the items can be delivered from those SPs, as well as the MDC, to the DPs. Relief items can also be rerouted between different SPs at any point in time during the planning horizon. The objective is to *minimize* the total *expected* cost for serving the demand for relief items, which is modeled as a random variable denoted by $\tilde{d}_j$, $\forall j \in J$.

The overall cost function consists of two components: (i) *logistics* cost, and (ii) *penalty* cost for failing to serve the demand shortage (if any is present). The logistics cost consists of three different components, (1) transportation cost, (2) inventory holding cost, and (3) procurement cost. We denote the unit cost of transporting relief items from an SP or the MDC $i \in \{0\} \cup I$ to a DP $j \in J$ by c^a_{ij} and the unit cost of transporting/rerouting the relief items from the MDC and an SP $i \in \{0\} \cup I$ to (and between) the different SPs $i' \in I$ by $c^b_{ii',t}$. We also denote the unit inventory cost at an SP $i \in I$ by $c^h_{i,t}$ and the unit procurement cost by h. Moreover, we denote the penalty cost for each unit of demand shortage by p, and in case the total amount of relief items across all the SPs exceeds the total amount of demand across all the DPs, the surplus can be salvaged at a price of q for each unit of overstock. We also make the following assumptions.

1. The MDC has unlimited capacity, whereas the capacity of each SP is fixed and is given by $\bar{x}_i$, $\forall i \in I$.
2. The locations of DPs depend on the hurricane's intensity landfall location, and we do not make decisions regarding the selection of DPs.
3. The forecast accuracy increases as the hurricane approaches landfall, whereas the logistics costs are non-decreasing from one period to the next.
4. All shipments made at the start of period t will arrive at their destinations within one time period, i.e., by the start of period $t+1$.

To define the problem formulation, we introduce the following decision variables:

- $\mathbf{x}_t = \{x_{i,t}\}_{i\in I}$, where $x_{i,t}$ denotes the level of storage at SP $i \in I$ at the end of period t, $\forall t = 1,2,\ldots,T$,
- $\mathbf{f}_t = \{f_{ii',t}\}_{i\in\{0\}\cup I, i'\in I}$, where $f_{ii',t}$ denotes the flow from the MDC/SP $i \in \{0\} \cup I$ to an SP $i' \in I$, $\forall t = 1,2,\ldots,T$,
- $\mathbf{y} = \{y_{ij}\}_{i\in\{0\}\cup I, j\in J}$, where y_{ij} denotes the *final* delivery from the MDC/SP $i \in \{0\} \cup I$ to a DP $j \in J$.

As such, for a given time of landfall T, demand levels $\bar{d}_j$, $\forall j \in J$, and initial inventory $x_{i,0}$ at different SPs, $\forall i \in I$, the corresponding *deterministic* multi-period disaster relief supply prepositioning problem can be formulated as:

$$
\begin{aligned}
\min_{\mathbf{x}_t,\mathbf{f}_t,\mathbf{y}} z_t(\mathbf{x}_t,\mathbf{f}_t,\mathbf{y}) := \quad & \sum_{t=1}^{T}\left(\sum_{i\in\{0\}\cup I}\sum_{i'\in I} c^b_{ii',t} f_{ii',t} + \sum_{i\in I} c^h_i x_{i,t}\right) + h\left(\sum_{t=1}^{T}\sum_{i\in I} f_{0i,t} + \sum_{j\in J} y_{0,j}\right) \\
& + \sum_{i\in\{0\}\cup I}\sum_{j\in J} c^a_{ij} y_{ij} + \sum_{j\in J} p\left(\bar{d}_j - \sum_{i\in\{0\}\cup I} y_{ij}\right) + \sum_{i\in I} q\cdot\left(x_{i,T} - \sum_{j\in J} y_{ij}\right) \\
\text{s.t.} \quad & x_{i,t-1} + \sum_{j\in\{0\}\cup I, j\neq i} f_{ji,t} - \sum_{j\in I, j\neq i} f_{ij,t} = x_{i,t}, && \forall i \in I, \forall t = 1,2,\ldots,T \\
& \sum_{j\in I, j\neq i} f_{ij,t} \le x_{i,t-1}, && \forall i \in I,\ \forall t = 1,2,\ldots,T \\
& 0 \le x_{i,t} \le \bar{x}_i, && \forall i \in I, \forall t = 1,2,\ldots,T \\
& \sum_{j\in J} y_{ij} \le x_{i,T}, && \forall i \in I \\
& \sum_{\{0\}\cup I} y_{ij} \le \bar{d}_j, && \forall j \in J \\
& y_{ij} \ge 0, && \forall i \in \{0\}\cup I, j \in J.
\end{aligned} \tag{1}
$$

4. Markovian MSP Models for Hurricane Relief Logistics Planning

In this section, we present the MSP model where the DM adaptively makes an operational logistics decision at any stage in the planning horizon $t = 1,\dots,T$ given the system state. We consider the *risk-neutral* case, where when making decisions at any point in time $t < T$, the DM tries to minimize the immediate cost plus the *expected* future cost. To that end, let us first discuss how we can model the evolution of the underlying stochastic process.

4.1 modeling the underlying stochastic process

To predict the demand level at different DPs, the DM receives sequential forecast advisories $\{s_t\}$ about the hurricane's attributes at the beginning of every period $t = 1,\dots,T$ of the planning horizon. From these forecast advisories, we consider two attributes of the impending hurricane's characteristics: its predicted intensity and location. These attributes can then be used to estimate the demand $\tilde{d}_j$ for relief items in each DP $j \in J$ once the hurricane makes landfall at time T.

To model the evolution of the underlying stochastic process governing the expected demand $\tilde{d}_j, \forall j \in J$, we assume that the hurricane's intensity and the location at time $t = 1,\dots,T$, are given by two *independent* random variables, α_t and ℓ_t, respectively, that is, $s_t := (\alpha_t, \ell_t)$. We also assume that the evolution of α_t, and ℓ_t can both be modeled as a MC with finite state spaces $\mathcal{A}$ and $\mathcal{L}$, respectively. Moreover, the transition probability matrix for α_t is given by $\mathbb{P}(\alpha_t = n|\alpha_{t-1} = k) = p_\alpha(k,n)$, and the transition probability matrix for ℓ_t is given by $\mathbb{P}(\ell_t = m|\ell_{t-1} = l) = p_\ell(l,m)$. Given the independence of α_t and ℓ_t, we can define the *conditional joint* probability distribution $P_{s_t|s_{t-1}}(n,m) := \mathbb{P}(\alpha_t = n, \ell_t = m|s_{t-1})$ as $\mathbb{P}(\alpha_t = n|\alpha_{t-1} = k) \times \mathbb{P}(\ell_t = m|\ell_{t-1} = l), \forall t = 1,\dots T$. Then, with some abuse of notation, we can rewrite $P_{s_t|s_{t-1}}(n,m) = p_s(n,m) = p_\alpha(k,n) \times p_\ell(l,m)$ and define the state space of the stochastic process formed by s as $\mathcal{S} := \mathcal{A} \times \mathcal{L}$. Finally, once the state space $\mathcal{S}$ is defined and all the transition probabilities are calculated, a demand value $\tilde{d}_j, \forall j \in J$ is defined for each state $s_t \in \mathcal{S}$ (pair of intensity and location realization). These values can then be used in the MSP formulation of the form that we discuss next.

4.2 MSP formulation

A generic form of the MSP model can be written as:

$$\min_{a_1 \in A_1(a_0,\xi_1)} z_1(a_1,\xi_1) + \mathbb{E}_{|\xi_{[1]}}\left[\min_{a_2 \in A_2(a_1,\xi_2)} z_2(a_2,\xi_2) + \mathbb{E}_{|\xi_{[2]}}\left[\cdots + \mathbb{E}_{|\xi_{[T-1]}}\left[\min_{a_T \in A_T(a_{T-1},\xi_T)} z_T(a_T,\xi_T)\right]\right]\right]. \quad (2)$$

Here, $\xi_{[t]} := (\xi_1,\dots,\xi_t)$ denotes the history of the stochastic process up to time t, $z_t(a_t,\xi_t)$ is the *immediate* cost paid by the DM at time t, a_t is the *action* taken at time t, and A_t is the set of feasible actions, $\forall t = 1,\dots,T$. In the context of the relief items prepositioning problem described in the previous section, $\xi_{[t]} := s_t$, $a_t := (\mathbf{x}_t, \mathbf{f}_t)$, $\forall t = 1,\dots,T$ and $a_T := (\mathbf{x}_T, \mathbf{f}_T, \mathbf{y})$. Moreover, $z_t(a_t,\xi_t)$ is given by the first two terms in the objective function of (1) for $t = 1,\dots,T-1$, and $z_T(a_T,\xi_T)$ has the same objective as (1) with the fourth term being replaced by $\sum_{j\in J} p \cdot \max\left\{0, \left(\tilde{d}_j - \sum_{i\in\{0\}\cup I} y_{ij}\right)\right\}$ and constraints $\sum_{\{0\}\cup I} y_{ij} \le \tilde{d}_j, \forall j \in J$ being removed. Finally, $A_t(a_{t-1},s_t)$ is given by the first three sets of constraints in (1) for all $t = 1,\dots,T-1$ and $A_T(a_{T-1},s_{T-1})$ includes all constraints in (1) associated with $t = T$.

Problem (2) is a nested formulation at its present form. A common approach to proceed with the computation is to approximate the underlying process using scenario trees and use a dynamic programming (DP) formulation [1], where for each state $j \in \mathcal{S}$ in stage t, given the system state from previous stage $t-1$, the optimization problem following (2) can be written as:

$$Q_t^j(a_{t-1},s_t) := \min_{a_t}\left\{ z_t(a_t,s_t) + \mathfrak{Q}_{t+1}^j(a_t) \;\middle|\; a_t \in A_t(a_{t-1},s_t)\right\}, \quad (3)$$

where $\mathfrak{Q}_{t+1}^j(a_t) := \sum_{k\in\mathcal{S}} p_s(j,k) \cdot \mathfrak{Q}_{t+1}^k(a_t)$ is the expected *cost-to-go* function for each $t \ne T$, and $\mathfrak{Q}_{T+1}^k(a_T) := 0$, $\forall k \in \mathcal{S}$. This problem can be solved by using a MC variant of the nested Benders decomposition [12].

5. Case Study: Implementation Details and Numerical Results

In this section, we present some preliminary numerical results for the MSP formulation (3) compared to the *clairvoyance* solution (CV) and the *mean value* solution (MV) on a set of N sample paths. The solution of MSP problems is typically non-anticipative, meaning that, when making a decision at time t, the DM does not know the entire sample-path on how the stochastic process will evolve into a leaf node in the scenario tree. If the non-anticipativity constraints are relaxed, the formulation decomposes into independent subproblems for each scenario which can then be optimized individually. This corresponds to the perspective of a clairvoyant, which gives a lower bound to the optimal value obtained by the MSP formulation (3). The MV solution corresponds to the situation where the DM aggregates all of the future exogenous information using a point estimator $\mu_{k,t} = \mathbb{E}[s_t|k]$ for $t = 1,\dots,T$, solve the corresponding *deterministic* problem and evaluates the resulting solution for each scenario $n \in N$.

We consider the same problem setting as described in [8] and [11]. To create a variety of instances, we consider different planning horizons $T \in \{4,8,12,16\}$, a different number of SPs $|I| \in \{3,7,15\}$, and a different number of DPs $|J| \in \{6,14,30\}$. Here, each time period is a six hours time interval, and it is assumed to be enough to preposition/reallocate the relief items between the different SPs. We also consider a sample size of $|N| = 1000$ sample paths. Moreover, we assume that all the components of the logistics (transportation, holding, and procurement) cost scales by a factor of $t \forall t = 1,\dots,T-1$, and that the demand is zero before the hurricane makes landfall at time $t = T$, that is, $\tilde{d}_{j,s_t} = 0, \forall j \in J, s_t \in S, t = 1,\dots,T-1$. Then, once the hurricane makes landfall at time $t = T$, a demand of $\tilde{d}_{j,s_T}$ is incurred at every pair of intensity and location state $s_T \in S$.

All of the algorithms were implemented in *Julia* 1.4.0, using *JuMP* 0.18.4 package [4], with commercial solver *Gurobi*, version 9.0.0 [9]. All of the tests were conducted on a high-performance computing cluster, where we used an *R830 Dell Intel Xeon* “big memory” compute node with 2.60GHz, 1.0 TB memory, and 24 cores.

In Table 1, we report the following statistics pertaining to different instances: (i) the lower bound (LB) for the objective value of the MSP formulation denoted by $\underline{z}$; (ii) the sample average $\hat{z}$ plus/minus the sample standard deviation $\hat{\sigma}$ for the upper bound (UB) of the MSP formulation; (iii) the relative gap between the LB $\underline{z}$ and the statistical UB $\bar{z}_+$ with a 95% confidence level, which is given by $\bar{z}_+ = \hat{z} + 1.96\sigma/\sqrt{|N|}$; (iv) the relative optimality gap of the MSP formulation and the MV solutions compared to the CV solution; and (v) the computational time (in seconds) for solving and evaluating the MSP and the MV solutions.

From Table 1, we can make the following observations. First, the relative gap between $\underline{z}$ and $\bar{z}_+$, which is given by $(\bar{z}_+ - \underline{z})/\underline{z}\%$ is small in the test instances (with the maximum being 2.19%). This means that a near-optimal decision policy is obtained from the MSP model using the employed algorithm. Second, we can see that the MSP formulation performs better than the MV solution in terms of the performance of their respective decision policies. Specifically, for the instances with ($|I| = 3, |J| = 6$), ($|I| = 7, |J| = 14$), and ($|I| = 15, |J| = 30$), compared to the CV solution, the MSP formulation has an average of 18.76%, 32.02%, and 30.54% in the optimality gap, whereas, the MV solution has an average of 27.82%, 65.07%, and 53.14% in the optimality gap, respectively. Finally, although the difference in the performance between the MSP and the MV solutions is not as apparent in instances with the smallest number of stage $T = 4$, it is more significant in the instances with a longer planning horizon T. For example, if we consider the average performance across all the instances where $T = 8, T = 12$, and $T = 16$, the relative gap of the MSP formulation compared to the CV solution is 17.80%, 17.71%, and 17.15%, respectively; and the relative gap of the MV solution compared to the CV solution is 39.42%, 42.80%, and 52.58%, respectively. This is to be expected. However, since the number of sample paths increases exponentially as T grows, the MV solution aggregates exogenous information more heavily, deteriorating the quality of the respective solutions.

Instance			Optimal objective value with the MSP formulation			Relative gap to CV		Time (in seconds)	
$\|I\|$	$\|J\|$	T	$\underline{z}$	$\hat{z} \pm 1.96\hat{\sigma}/\sqrt{\|N\|}$	$(\bar{z}_+ - \underline{z})/\underline{z}$	MSP	MV	MSP	MV
3	6	4	21313.00	21465.73 ±314.41	2.19%	51.99%	60.55%	57.16	3.99
		8	59364.25	59516.77 ±406.64	0.94%	11.64%	21.02%	332.53	3.83
		12	138461.72	139310.98 ±626.53	1.07%	6.04%	15.98%	425.91	3.79
		16	206803.88	208315.54 ±1145.40	1.28%	5.37%	13.74%	740.14	3.88
7	14	4	77440.62	77563.14 ±300.40	0.55%	54.97%	57.46%	105.73	6.23
		8	105906.96	106208.31 ±1014.38	1.24%	21.34%	50.20%	423.60	6.47
		12	128864.70	128907.48 ±1534.38	1.22%	24.42%	65.23%	971.30	6.65
		16	144471.32	142657.21 ±1912.50	0.07%	27.32%	87.39%	3280.19	7.37
15	30	4	150119.48	149869.34 ±1107.27	0.57%	60.35%	61.70%	283.91	12.84
		8	253459.61	252879.50 ±1922.69	0.53%	20.43%	47.05%	1910.49	14.00
		12	266961.48	266508.21 ±3194.60	1.03%	22.65%	47.18%	6710.24	13.50
		16	380562.53	382093.32 ±4003.02	1.45%	18.75%	56.61%	22380.76	27.98

Table 1: The LB and UB statistics of the optimal objective value obtained using the MSP formulation (3); the relative gaps between the LB and the statistical UB; the optimality gaps between the CV and the MSP and MV solutions; and the computational time (in seconds) for solving and evaluating the MSP formulation and obtaining the MV solutions.

6. Conclusion

This work proposes a fully adaptive MSP model for solving a typical hurricane relief logistics planning problem under uncertainty. To solve this problem, we assume that a DM receives sequential forecast advisories about an impending hurricane's attributes. Then he/she uses this information to construct a MC which predicts the demand for relief items. Although we assume that the hurricane time of landfall T is fixed, this model can be generalized to a situation where T is treated as a geometrically distributed random variable in an infinite horizon discounted problem. From our case study we can see that the relative gap between the lower and the statistical upper bound on the objective value of our model is very small, with a less than one percent average across all of the instances. We can also see that compared to the MV solution, the proposed MSP model provides a significantly higher quality decision policy due to its adaptability inherited in the model formulation – especially for instances where the number of stages in the planning horizon is large.

Acknowledgements

The authors acknowledge partial support by the National Science Foundation [Grant CMMI 1854960]. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] Richard Bellman. Dynamic programming. *Princeton University Press*, 1957.

[2] Global Climate and Weather Modeling Branch. The GFS atmospheric model technical report, national weather service, 2003. URL `https://www.emc.ncep.noaa.gov/officenotes/newernotes/on442.pdf`.

[3] Vitor L de Matos, David P Morton, and Erlon C Finardi. Assessing policy quality in a multistage stochastic program for long-term hydrothermal scheduling. *Annals of Operations Research*, 253(2):713–731, 2017.

[4] Iain Dunning, Joey Huchette, and Miles Lubin. Jump: A modeling language for mathematical optimization. *SIAM Review*, 59(2):295–320, 2017.

[5] Jitka Dupačová and Jan Polívka. Asset-liability management for czech pension funds using stochastic programming. *Annals of Operations Research*, 165(1):5–28, 2009.

[6] Daniel Duque, Haoxiang Yang, and David P Morton. Optimizing diesel fuel supply chain operations for hurricane relief. *Optim. Online*, 2020.

[7] Boutheina Fhoula, Adnene Hajji, and Monia Rekik. Stochastic dual dynamic programming for transportation planning under demand uncertainty. In *2013 International Conference on Advanced Logistics and Transport*, pages 550–555. IEEE, 2013.

[8] Gina Galindo and Rajan Batta. Prepositioning of supplies in preparation for a hurricane under potential destruction of prepositioned supplies. *Socio-Economic Planning Sciences*, 47(1):20–37, 2013.

[9] LLC Gurobi Optimization. Gurobi optimizer reference manual, 2019. URL `http://www.gurobi.com`.

[10] Emmett J Lodree Jr and Selda Taskin. Supply chain planning for hurricane response with wind speed information updates. *Computers & Operations Research*, 36(1):2–15, 2009.

[11] Gina Galindo Pacheco and Rajan Batta. Forecast-driven model for prepositioning supplies in preparation for a foreseen hurricane. *Journal of the Operational Research Society*, 67(1):98–113, 2016.

[12] Ragheb Rahmaniani, Teodor Gabriel Crainic, Michel Gendreau, and Walter Rei. The benders decomposition algorithm: A literature review. *European Journal of Operational Research*, 259(3):801–817, 2017.

[13] Michael J Widener and Mark W Horner. A hierarchical approach to modeling hurricane disaster relief goods distribution. *Journal of Transport Geography*, 19(4):821–828, 2011.