A Stochastic Optimization Model for Staged Hospital Evacuation during Hurricanes

Tarun Rambha*[1,2], Linda K. Nozick[3], Rachel Davidson[4], Wenqi Yi[5], and Kun Yang [4]

[1]Department of Civil Engineering, Indian Institute of Science

[2]Center for infrastructure, Sustainable Transportation and Urban Planning, Indian Institute of Science

[3]School of Civil and Environmental Engineering, Cornell University

[4]Department of Civil and Environmental Engineering, University of Delaware

[5]Systems Engineering, Cornell University

Abstract

Hurricanes result in large scale evacuations almost every year. Of particular concern and difficulty is the decision of whether or not to evacuate hospitals in these emergencies. During an emergency, a hospital is a source of refuge, and evacuating its patients is often viewed as a last resort since it is difficult to provide quality care while transporting them. At the same time, flooding and loss of power and communications put patients and caregivers at very high risk. Most emergency response plans do not have clear guidelines for evacuating or sheltering-in-place. Hurricanes are particularly complicated because there is often considerable uncertainty surrounding their eventual trajectory and intensity. These factors have contributed to, what is in hindsight, poor decisions that have cost lives. The current paper addresses this problem by developing a stochastic optimization formulation, taking into account evolving conditions and, therefore a hopefully robust collection of future flood, wind, and roadway traffic predictions. The model determines the order in which patients should be evacuated over time based on the evolution of the storm by trading off cost and risk. A holistic case study focused on North Carolina and the evolution of Hurricane Isabel is presented by fusing data and model outputs from different sources. The results highlight the advantages of using a recourse formulation that adapts to new information and illustrates the proposed decision-support model's long-term applications.

Keywords: Hurricanes; hospital evacuation; disaster management; scenario trees; multi-stage stochastic programming

1 Introduction

Natural disasters such as hurricanes inflict huge financial and infrastructural losses and often lead to casualties. On average, in the last century, 1.7 major hurricanes strike the US coast each year (Landsea, 2015). Hence, several state and local government agencies have led efforts to develop robust contingency plans that minimize damage due to hurricanes. These plans are typically centered around the issuance of voluntary and/or mandatory evacuation orders, identification of contraflow plans, and rationing of resources such as gasoline. However, evacuation of vulnerable populations, especially those in hospitals, is often neglected and left to hospital administrators (Dosa et al., 2007). While some emergency response programs such as the Federal Emergency Management Agency (FEMA), Medical Reserve Corps (MRC), and Metropolitan Medical Response System (MMRS) do help assemble additional medical personnel and supplies, they have only been mildly successful due to the lack of coordination and consistent guidelines (Franco et al., 2006). Thus, without proper decision support tools and resources, hospital administrators often make poor choices that contribute to risky evacuations, high expenditures, and in some cases, patient fatalities.

*✉ tarunrambha@iisc.ac.in

For instance, 10 out of 11 major hospitals in New Orleans, Louisiana, decided not to evacuate during Hurricane Katrina, resulting in tens of deaths: 19 at Lindy Boggs Medical Center, 45 at Memorial Medical Center, and 22 at Lafon Nursing Home. Nearly 1,700 patients were supposedly evacuated under dangerous conditions during Katrina (Gray and Hebert, 2007). Within a few weeks, during Hurricane Rita, which struck Louisiana and Texas, 23 patients who were being evacuated in a bus died near Houston after it caught fire due to overheated brakes and the oxygen tanks on board (Zachria and Patel, 2006). In another incident, 34 patients at a nursing home in Chalmette, Louisiana reportedly drowned after the management decided to shelter-in-place (Dosa et al., 2007). More recently, during Hurricane Irene, many hospitals in Manhattan preemptively evacuated after the mayor of New York City (NYC) issued mandatory evacuation orders. However, in retrospect, these evacuations were unnecessary and cost millions of dollars in lost revenue since the hurricane was not as threatening as they believed it to be. This led to a certain degree of complacency among hospital administrators when Hurricane Sandy struck the following year. Only 2 out of 5 hospitals in lower Manhattan evacuated in advance (Ricci et al., 2015). In fact, hospitals such as the New York Harbor Healthcare System's Manhattan Campus and the Langone Medical Center, which were blocks away from each other, made drastically different decisions. The former evacuated before the storm, whereas the latter decided to shelter-in-place. Unfortunately, backup generators at the Langone Medical Center failed due to storm surge, and 322 patients (of which 21 were infants in the neonatal intensive care unit (ICU)) had to be evacuated overnight in the middle of the storm. Overall, about 6,300 patients were evacuated from 37 facilities in NYC during and after Hurricane Sandy (McGinty et al., 2017).

Hospital evacuations during hurricanes are a complex, large-scale problem. During an emergency, hospital administrators have to make a myriad of decisions based on partial or uncertain information. Unlike other natural disasters, a hurricane's trajectory and intensity change over time, and no evacuation drill can closely mimic its consequences. Hence, emergency protocols used by a hospital during hurricanes must be different from those used for other evacuation situations such as earthquakes and HazMat. Also, evacuation during hurricanes is a collaborative process, and a hospital must make choices in the context of other decisions taken by state and local public agencies. To get a glimpse of the magnitude of the problem, see Sexton et al. (2007) which details the efforts behind the evacuation of 427 patients in 10 hours from The University of Texas Medical Branch at Galveston during Hurricane Rita.

Hospitals have traditionally been one of the last facilities to evacuate during hurricanes since transporting the sick is risky, and because the chances of an individual's need for medical care during a disaster is very high. However, sheltering-in-place can have adverse consequences because flooding often leads to loss of power and communications, and in some cases patient deaths as cited earlier. During Hurricane Katrina, most of New Orleans was under 20 feet of water and without power for nearly five days (Franco et al., 2006). Hospitals in the area had to thus deal with dysfunctional elevators and a short supply of oxygen, food, medicines, and clothing (Downey et al., 2013). Lack of power led to severe dehydration among patients as the temperatures soared over 100 degrees, and some surgeries were reportedly performed using a flashlight with little anesthesia (Gray and Hebert, 2007). In such cases, hospitals that decide to shelter-in-place may be forced to evacuate after a hurricane's landfall. However, without proper channels for communication, emergency response teams cannot determine patient conditions and needs, making it difficult to prioritize their operations (Franco et al., 2006). Post-flooding evacuations would also likely require boats and helicopters, which are not only risky but are expensive and difficult to gain access to (Gray and Hebert, 2007).

If a hospital decides to evacuate before landfall, it is crucial to consider current and predictions of future roadway traffic conditions. Hurricanes induce region-wide evacuation of households, which clogs transportation networks since they are not designed for such a large influx of vehicles. For instance, during Hurricane Rita, nearly 3 million individuals left the Houston area, which resulted in massive gridlocks on freeways. Soika (2006) reports that it took 12–20 hours to get from Houston to neighboring cities such as San Antonio and Austin. During Hurricane Sandy, nearly 3,500 traffic lights stopped functioning within NYC (Gibbs and Holloway, 2013). In many cases, dedicated emergency lanes are unavailable because they are open to regular traffic (Sexton et al., 2007). A hospital administrator must also decide the order in which patients are evacuated, their destination hospitals, and their travel mode. Vehicles may not be equipped to handle wheelchairs, and shortages caused when ambulance providers do not adhere to contracts are all common issues during an emergency (Dosa et al., 2007). Thus, delaying evacuation decisions can be detrimental, but at the same time, transporting patients early can unnecessarily put them at risk. Optimizing these decisions is, therefore, a near-impossible task without proper decision support tools.

Hospital evacuations also involve several other logistical difficulties such as staffing issues, identifying receiving hospitals, and transferring medical records. French et al. (2002) points out that hospital staff are primarily concerned with family safety, pet care, and personal safety while deciding to continue to serve during an emergency. Lack of an adequate number of physicians, nurses, radiologists, and lab technicians often prohibits a hospital from functioning efficiently. In this context, arrangements to shelter family and the provision of additional compensation have had some success in retaining staff (Downey et al., 2013). Administrators must also identify partner hospitals in the area in advance and sign mutual aid agreements for use during a disaster. Hospitals within a network (such as the US Department of Veterans Affairs) can easily find places with the capacity and expertise to treat their patients. However, for hospitals that are not part of a larger network, finding receiving hospitals for patients may not be easy. For example, many hospitals in Louisiana did not find receiving hospitals during Hurricane Katrina (Gray and Hebert, 2007). In Baton Rouge, the Louisiana State University Health Sciences Center had to set up a temporary field hospital with 800 beds in which over 6,000 patients were treated (Franco et al., 2006). On the other hand, during Hurricane Sandy, the state department of health activated an inter-agency Healthcare Facility Evacuation Center, whose objective was to match evacuated patients with destination facilities. They also played a vital role in communicating with the destination hospitals, which allowed them to decompress and accept transferred patients (Adalja et al., 2014). In some cases, patients were admitted without knowing who would pay for their treatment, and reimbursements were an issue. Finally, hospitals also face difficulties in transferring medical records during an emergency. The lack of a national-level electronic system for health records has, in the past, forced hospitals to send limited or sometimes no information along with transferred patients, which leads to delays in treatment at receiving hospitals. Tracking patients who are evacuated is another matter of concern. According to Gray and Hebert (2007), a few patients transferred from hospitals in New Orleans during Hurricane Katrina could not be traced even after three months.

This paper is a step towards developing a comprehensive decision support tool that can assist hospital administrators in staging their patient evacuations in the event of a hurricane. The following list summarizes the research gaps and the contributions of our work.

- Existing mathematical models on hospital evacuation (Tayfur and Taaffe, 2009b; Bish et al., 2014) are not well-suited for hurricanes since they focus on moving all the patients in a hospital after a disaster. Instead, using a probabilistic description of hurricane evolution in track and intensity, we develop a risk and cost informed trade-off analysis, including the opportunity to move none or only a subset of patients. We incorporate the risks associated with wind speeds and flood depths for different hurricane scenarios and consider time-varying roadway conditions to estimate risk. Besides, cost objectives are also considered because substantial literature suggests that they play an important role in evacuation decisions.
- We propose a scenario tree-based stochastic optimization model which allows hospital administrators to stage evacuations over time and update their decisions depending on how a hurricane evolves. Also, hospital managers can use this model to develop emergency plans for conical storms, including for which hospitals' long-term mutual aid agreements with receiving hospitals and ambulance providers are warranted.
- We also demonstrate via an elaborate case study the value of this model in practice. In particular, we present a hypothetical optimization of evacuation decisions for a hospital in North Carolina during Hurricane Isabel. Our study uses data from multiple sources, such as the American Hospital Association (AHA) and the National Hurricane Center (NHC). It also utilizes outputs from previous research on scenario-ensemble-based hurricane forecasting and regional traffic network models that find optimal evacuation orders under predicted future conditions.

The remainder of this paper is organized as follows. In Section 2, we discuss literature on transportation issues during hurricanes and explore existing work on hospital evacuations. In Section 3, we provide some background on scenario trees and multi-stage stochastic programming and then propose an optimization formulation for minimizing a combination of cost and risk associated with evacuating a hospital before a hurricane's landfall. A case study with data from North Carolina's hospitals and Hurricane Isabel is then used to demonstrate the utility of the suggested formulation in Section 4. Finally, in Section 5, we conclude this study by highlighting some of the results. The proposed model addresses the evacuation of a single hospital. We leave the problem of evacuating multiple hospitals in a region as a topic for future research.

2 Literature Review

2.1 Hurricane Evacuation Modeling

Evacuation models have largely focused on moving a large number of travelers to safer locations under limited roadway capacity. Strategies in these models typically include timing mandatory evacuation orders, activating contraflow lanes, and providing proper route guidance. For detailed reviews on these topics, see Wolshon et al. (2005), Xiongfei et al. (2010), Murray-Tuite and Wolshon (2013), Bayram (2016), and Rambha et al. (2019).

To predict a transportation network's performance, it is essential to understand demand-side aspects, i.e., when people evacuate and their destination preferences. Estimating the probabilities with which individuals evacuate has been studied using random utility models involving independent variables such as household socio-demographic characteristics, source of evacuation notices, and storm severity (Hasan et al., 2011; Gudishala and Wilmot, 2012; Xu et al., 2016). Similar techniques have been used to predict the destination choices (which typically include homes of relatives, hotels, and shelters) of individuals and households (Cheng et al., 2008; Mesa-Arango et al., 2013). Using this information, one can estimate traffic conditions under some additional assumptions on the route choice behavior. For instance, one could assume that travelers are fully aware of the network conditions and act selfishly (resulting in a user equilibrium) (Hobeika and Kim, 1998) or follow centralized routing/departure time strategies, leading to a system optimum state (Chiu and Zheng, 2007; Ng and Waller, 2010). Other network loading methods that are based on myopic routing strategies (Sheffi et al., 1982), stochastic route choice models (Sadri et al., 2013), and multi-stage stochastic optimization using a cell transmission model (Karabuk and Manzour, 2019) have also been proposed. The methods proposed in this paper has similarities with Karabuk and Manzour (2019) in that we try to optimize both risk and cost of evacuating using a scenario tree-based framework in the presence of a disaster which has both track and intensity uncertainty. Their paper models the evacuation of a geographical region that is likely to be affected by a tornado and uses risk values of staying in buildings and traveling in vehicles to study the impact of lead times and information levels on evacuation decisions. In contrast, our model studies moving a vulnerable population of patients on a time-dependent network where the background traffic is derived from a similar regional-assignment model. Further, patient evacuations require fine-grained risk calculations since we have to factor in the type of patients and vehicles in which they are transported in addition to the destination choices and the associated risks from taking different routes to the receiving hospitals.

A key strategy to improve evacuation operations is to control the demand by intelligently timing mandatory evacuation orders (Sbayti and Mahmassani, 2006; Chen and Zhan, 2008; Bish and Sherali, 2013; Zhang et al., 2014). Recently, Yi et al. (2017) suggested a stochastic optimization model for this problem using scenario trees constructed from predictions of hurricane intensity and trajectory. Their model optimizes network-level travel time and risk while assuming that users choose minimum travel time routes. The risk component is constructed using risk factors, proposed by Apivatanagul et al. (2012), that depend on predicted storm surge and wind speeds. The supply-side components of a road network can also be made efficient using contraflow strategies. By reversing the direction of select roadway links, the capacity of routes leaving areas affected by a hurricane can be increased, and a speedy evacuation can be carried out (Cova and Johnson, 2003; Wolshon et al., 2005; Tuydes and Ziliaskopoulos, 2006; Kim et al., 2008; Xie and Turnquist, 2009; Xie et al., 2010; Wang and Wang, 2019). Empirical studies have indicated that the roadway capacities during emergencies are lower than those during routine operations since evacuees may be unfamiliar with different routes, and such factors must also be considered while designing supply-side control measures (Dixit and Wolshon, 2014).

From a long-term planning perspective, the time needed to evacuate can be indirectly influenced by identifying and building shelters at key locations. Bi-level optimization models by Sherali et al. (1991), Ng et al. (2010) and Li et al. (2012) address this problem with slightly different constraints and assumptions on route choices behavior of travelers. Other evacuation-related topics that have been explored include network design and scheduling of transit services for individuals without auto access (Bish, 2011; An et al., 2013; Kaisar et al., 2012), post-disaster logistics (Sheu, 2007, 2010; Sheu and Pan, 2014; Zhou et al., 2017), and measurement of the effectiveness of evacuations using clearance times and spatial and temporal risk factors (Han et al., 2007).

2.2 Hospital Evacuations

While the evacuation of general populations has been well studied, hospital evacuations have not received much attention in the literature. A survey of hospitals between 1971 and 1999 by Bagaria et al. (2009) revealed that hurricanes are the third most common reason for evacuating hospitals (preceded by internal fires and HazMat). Yet, clear decision rules do not exist for hospital evacuations during hurricanes. In recent years, post-Hurricane Katrina and Hurricane Sandy, many government and non-profit organizations have led efforts to produce guidebooks that provide key information and suggestions for use during an evacuation. A repository of such material can be found at the US Department of Health and Human Services (USDHHS, 2015) website. The Agency for Healthcare Research and Quality's decision guide for hospital evacuation (AHRQ, 2010) is another such resource which can be used for some "back of the envelope" calculations. While these documents are useful in understanding what to expect during an emergency, they are less specific when addressing crucial decisions such as when and whom to evacuate. In the following paragraphs, we discuss a few studies which have attempted to fill this gap.

Taaffe et al. (2006) proposed a discrete event simulation model for predicting the evacuation time for a hospital. The hospital housed three types of patients, and it was assumed that there were 50 patients of each type. Patients could be transported to three receiving hospitals located a hundred miles away using a fixed fleet of vans and ambulances. The model, programmed in Arena, simulated expected release times and staging times of patients while assuming that they depended on storm conditions. Assignments of patients to vehicles was done based on the patient type and vehicle capacity. The simulation framework was extended in Tayfur and Taaffe (2009b) by considering S-shaped curves and triangular distributions to model stochastic vehicle travel times and waiting times. Staffing constraints and rules for nurse availability at different times were also added. The authors then used OptQuest, a heuristic solver, to optimize the number of vehicles and nurses needed to evacuate patients at minimum cost (includes vehicle leasing, round trip, and staffing costs). The effect of the start time on cost and evacuation completion time was also analyzed.

For similar problem instances, Tayfur and Taaffe (2009a) suggested a mixed-integer programming model which minimizes operating and evacuation costs. The model provides the optimal number of patients of each type to be evacuated to different hospitals in different vehicles across time, the number of available vehicles, and the number of staging and transportation nurses. The cost structure for different patient and nurse types were assumed to be different. Stranding a patient was discouraged by imposing a high fixed cost. The optimal solution could be computed using CPLEX in only 3 out of 36 cases, and linear programming relaxation methods with rounding and local search heuristics were discussed. Childers et al. (2009) studied a simpler evacuation problem with two patient types and used a Markov decision process approach to determine which of the two patient types should be evacuated in each time step. The state in their model is the number of patients of each type and state transitions were due to the evacuation strategy as well as due to patient deaths. The objective of their formulation was to maximize the number of lives saved.

Bish et al. (2014) analyzed the evacuation problem using an optimization model in which the objective was to minimize patient risk instead of cost. The risk function was assumed to consist of two components—risk from staying (threat risk) and risk associated with transportation. Different threat risk functions—constant, linear, and exponential—were considered. The transportation risk, on the other hand, depended on the time required to stage and transport patients. Patients were classified into different categories depending on their condition, and the hospital could use a fixed fleet of different vehicles to transport them. Capacity constraints for different patient types at the receiving hospitals were also imposed. The travel times to and from different hospitals were assumed constant. The authors apply the formulation on a Level I Trauma Center in Virginia, and the robustness of the model inputs was examined. The current paper builds on Bish et al. (2014) and specifically addresses the uncertainty in hurricane trajectory and intensity and its influence on the optimal choices.

The following three key features distinguish our work from existing literature. First, previous research on this topic is agnostic to the type of disaster and over-simplifies the time-varying risk. In the case of hurricanes, storm path influences rainfall, which causes flooding and dictates the risk from transporting patients on roadways. The path and intensity also play a role in zonal evacuation orders and subsequently staffing costs and time-varying network congestion. Our paper explicitly models different kinds of anticipated risks by fusing data from multiple sources such as weather and regional-scale network loading models. Second, the decisions allowed in our work adapt to the current state of the storm due to the use of a scenario-tree-based optimization approach. During highly uncertain and dynamic disasters such as hurricanes, it may be in the best interest of hospital managers to revise their plans constantly instead of following an existing one. Our model provides this flexibility and, in turn, minimizes the resources needed for evacuation. Third, using a detailed case study, we discuss the data inputs to our model and illustrate several managerial questions that can be addressed by hospitals planning to evacuate during a hurricane.

3 Formulation

3.1 Background

While the trajectory and intensity of a hurricane are usually unknown, meteorological models can help narrow it down to a small set of scenarios (each of which describes the hurricane's path, central pressure, and wind velocity) called an *ensemble* (Skamarock et al., 2005). The scenarios in an ensemble are generated by perturbing the initial and boundary atmospheric conditions and selecting different physics/dynamics in the widely used Weather and Research Forecasting (WRF) model. More details on the construction of the ensemble can be found in (Blanton et al., 2020). These predictions can be made several days before a hurricane's landfall. For each scenario, it is possible to estimate the wind speeds, flood depths, and network travel times and formulate a different optimization model for minimizing the cost and/or risk of evacuation. However, the scenario that will occur is not known in advance, and hence it is essential to construct a policy of what actions should be taken as more and more information is learned about the eventual storm path and intensity.

In order to obtain such a well-hedged solution, a multi-stage stochastic optimization model with recourse is employed. This type of modeling requires uncertainty to be represented using *scenario trees* (Rockafellar and Wets, 1991; Watson and Woodruff, 2011). A scenario tree can be viewed as partitions of the state space (scenarios) that grow finer over time. At each time step, the partition divides the scenarios into sets (called *scenario bundles*) so that the scenarios within each set are indistinguishable from each other. For instance, at the beginning of the time horizon, the predicted hurricane trajectories are spatially very close to each other and hence cannot be distinguished. Over time, it would be possible to rule out a few of the scenarios in the ensemble. Scenario-based optimization also allows us to have richer sets of constraints that are valid along particular sample paths (Dupačová and Sladký, 2002; Birge and Louveaux, 2011). We will demonstrate this feature in the next section.

Mathematically, if $\mathcal{A}_t$ represents a partition of the scenarios at time t, then a set of scenarios $A \in \mathcal{A}_t$ (we will refer to this set A as a node) is essentially a union of sets in $\mathcal{A}_{t+1}$, and each set in $\mathcal{A}_{t+1}$ is a subset of exactly one set in $\mathcal{A}_t$. This method of representing uncertainty has a tree structure since the scenarios that are indistinguishable at time step t could branch into smaller sets of scenarios. The decision-maker is assumed to make non-anticipative choices, i.e., they take the same actions under all scenarios that are indistinguishable at any time step.

Suppose S denotes the set of scenarios. Let $z = (z_{st})_{s \in S, t \in T}$ be a vector of decision variables, where z_{st} indicates actions chosen in scenario s at time step t. These recourse variables allow the decision-maker to use the hurricane conditions at each time step t to identify the closest scenario s and take decisions prescribed by z_{st}. In the evacuation context, these variables are comprised of patients sent to the receiving hospitals and the number of vehicles of different types that carry these patients. Assuming that the probability with which scenario s occurs is $\Pr(s)$, an abstract version of the proposed model can be written as

$$\min \sum_{s \in S} \Pr(s) f_s(z) \tag{1}$$

Subject to

$$z_{st} \in Z_s \qquad \forall s \in S, t \in T \tag{2}$$

$$z_{st} = z_{s't} \qquad \forall\, t \in T, s, s' \in A, A \in \mathcal{A}_t \tag{3}$$

where $f_s(z)$ is the objective function component representing the cost and risk objectives for scenario s. Constraint (2) includes flow-balance and supply-demand type of constraints and should hold for each scenario. The non-anticipativity condition is modeled in (3). The complete formulation presented in Section 3.3 uses variables for every node of the scenario tree and hence implicitly models the non-anticipativity constraint.

To construct a scenario tree from an ensemble of hurricane trajectories, one could use a clustering-based technique proposed by Yang et al. (2017). This technique takes the ensemble generated from WRF models as input and uses hurricane attributes such as perpendicular and parallel distance between trajectories and Saffir-Simpson intensity at each time step to reduce them to a smaller set of scenarios and to generate scenario bundles containing scenarios that

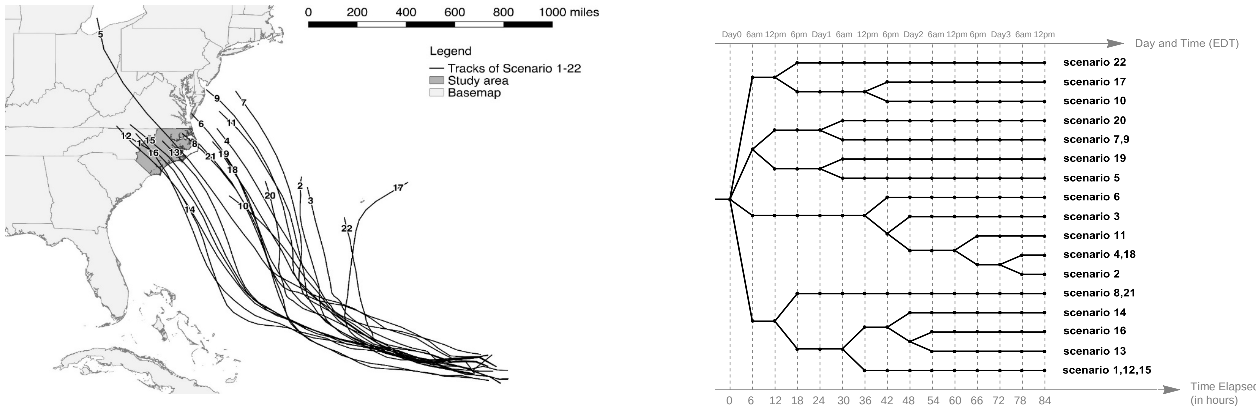

(a) Ensemble of 22 scenarios (b) Scenario Tree

Figure 1: Modeling uncertainty in hurricane path.

are indistinguishable. Figure 1a depicts an application of this technique to 22 scenarios of Hurricane Isabel (2003). The number of scenarios used in this paper is comparable with other papers that forecast weather and the effects of hurricanes using up to 20 scenarios on an average (Nasrollahi et al., 2012; Wanik et al., 2018; Villarini et al., 2019; Kowaleski and Evans, 2020). A scenario tree with one-hour time steps generated three-and-a-half days before the hurricane's landfall is shown in Figure 1b.

3.2 Model Inputs

In this section, we state the modeling assumptions and describe scenario tree parameters and other inputs needed for the stochastic program. We use capital letters to denote sets, Greek letters for parameters, and lower-case English alphabets for members of sets and decision variables. Superscripts are used specifically for nodes in the scenario tree and are not to be confused with exponents.

The sets used in the formulation are described in Table 1. We suppose that hospitals have a triage method to categorize patients into different types based on their condition and needs. Hospitals are also assumed to have a fixed fleet of vehicles of different types that are either privately owned or hired temporarily for evacuation purposes. Most commonly used vehicles include advanced life support (ALS) and basic life support (BLS) ambulances. In addition, hospitals may use buses or minivans to transport a larger number of non-critical patients. A set of hospitals that have a mutual agreement with the evacuating hospital for accepting patients in an emergency is assumed to be known beforehand. Decisions are taken at discrete intervals of time $T = \{0, 1, 2, \ldots, |T|\}$ within the duration left for evacuation. A node in a scenario tree is a collection of scenarios that are indistinguishable from each other. A node m is an ancestor of a node n if there exists a sample path from m to n. In other words, the scenarios in n are a subset of those in m. The set N represents the nodes in the scenario tree, and an overloaded version $N(n)$ is used to denote the set of ancestors of node n. The time step associated with node n is written using $t(n)$. For example, the scenario tree shown in Figure 2 has 5 scenarios, 10 nodes, 4 time steps, and 4 sample paths (1–2–4–7, 1–3–5–8, 1–3–5–9, and 1–3–6–10). The set of ancestors of node 8 is $N(8) = \{1, 3, 5\}$. We use $n \in r$ as shorthand for the set of nodes that belong to sample path r.

Table 1: Notation of sets used in the formulation.

Symbol	*Description*
P	Patient types (Emergency, ICU, Pediatric etc.)
V	Vehicle types (ALS and BLS ambulances, Buses etc.)
H	Hospitals to which patients can be shifted
N	Set of nodes in the scenario tree
T	Set of time steps/decision points
R	Set of sample paths in the scenario tree
$N(n)$	Set of all ancestor nodes of $n \in N$

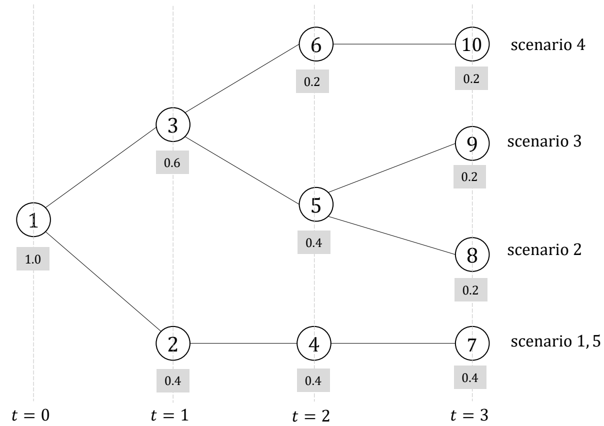

Figure 2: An example scenario tree.

The parameters that appear in the constraints of the formulation are shown in Table 2. We assume that the number of patients of each type is known and that there are limits on the intake for different patient types at all receiving hospitals. We also suppose that new patients are not admitted to the hospital during the time period of interest. Patient capacity and fleet size are also assumed to be known for each vehicle type. ALS and BLS ambulances can typically transfer 1-2 patients at a time, while the capacities of vans can be highly variable. Note that patients of certain types cannot sometimes be transported in certain kinds of vehicles. For instance, wheelchair patients cannot be accommodated in buses since they are not appropriately designed. Hence, we use indicators δ_{pv} to represent the feasibility of transporting a certain type of patient in a specific vehicle.

One of the key features of the proposed formulation has to do with modeling the traffic conditions. We assume that the travel times on network links are known for different departure times. For the case study presented in the paper, results from a dynamic traffic assignment model are borrowed from Yi et al. (2017), which is based on the formulation by Janson (1991) and Li et al. (2012), to estimate time-dependent link travel times. Using this information, time-dependent shortest paths (TDSP) are computed between the evacuating and receiving hospitals for each time step in T (see Chabini (1998) for details of the algorithm). The outputs are then used to define a binary parameter ζ_h^{mn} that tracks vehicles available for transportation. This parameter is set to 1, if a vehicle leaving at $t(m)$ to hospital h can return at or before $t(n)$. Travel times are not differentiated by vehicle type, but this assumption can be easily relaxed. To simplify the formulation and evacuation operations, we also assume that vehicles drop-off patients at a single receiving hospital in each trip and return to the hospital being evacuated to shift more patients. In other words, no vehicle routing problem (VRP) type constraints are imposed. While this assumption may not be significant for ALS and BLS ambulances whose capacity is low, it is somewhat restrictive for other vehicles with higher capacity.

Table 2: Constraint parameters.

Symbol	*Description*
α_p	No. of patients of type p to consider evacuating
β_v	No. of vehicles of type v available for evacuation
χ_v	Capacity of a vehicle of type v
κ_{ph}	No. of patients of type p that can be accommodated in hospital h
δ_{pv}	Binary parameter which is 1 if a patient of type p can be transported in a vehicle of type v and is 0 otherwise
ζ_h^{mn}	Binary parameter which is 1 if a vehicle that leaves the evacuating hospital at $t(m)$ can return at or before $t(n)$

The objective function of the proposed formulation minimizes a linear combination of the expected risk and expected cost of evacuation. To compute these expected values, we define the probability of observing a node n at $t(n)$, π^n, to be the sum of probabilities of scenarios belonging to it. In addition, we also define a parameter ρ^r to denote the probability of observing the set of indistinguishable scenarios along sample path r at the end of the time horizon (which is also the probability of observing the last node along the sample path r). For instance, in Figure 2, assuming that all scenarios are equally likely, the probability of observing nodes 4, 5, and 6 at $t = 2$ is 0.4, 0.4, and 0.2, respectively (see numbers in the grey boxes) because two scenarios belong to node 4 (Scenarios 1 and 5) and node 5 (Scenarios 2 and 3), and one scenario belongs to node 6 (Scenario 4). Likewise, the probability of sample path 1–2–4–7 is 0.4 because the sample path contains two indistinguishable scenarios.

In order to model patient risk, we consider two factors: risk due to evacuating λ^n_{pvh} and risk associated with sheltering-in-place μ^r_p. These risk parameters are probability measures and indicate the odds of an adverse event such as patient death. Several aspects of the evacuation problem such as patient and vehicle types, flooding, and routes used for travel influence these risk parameters, and a detailed description is postponed to Section 4 for better readability. The cost of evacuation, on the other hand, is assumed to include transportation and staffing expenditures. Transportation expenditures comprise fixed and per-mile costs, while staffing costs are indirectly modeled using overtime wage rates and the time spent by patients at the evacuating hospital. The overtime wage rates are assumed to come into effect only when the source hospital's zone gets an evacuation order.

Table 3: Objective function parameters.

Symbol	*Description*
π^n	Probability of observing node n
ρ^r	Probability of observing the terminal node along sample path r
ψ	Weight used to define a linear combination of components in the objective function
σ^n_p	Unit cost of attending to a patient of type p upto time step $t(n)$. We denote the cost of a stranded patient using σ_p
ω^n_{vh}	Unit cost of a vehicle v that is used to transfer patients to hospital h in time step $t(n)$
λ^n_{pvh}	Risk associated with evacuating patient of type p in vehicle v to hospital h at node n
μ^r_p	Risk of stranding patient of type p when sample path r is realized

3.3 Optimization Model

The variables used in the model are listed in Table 4. The decision-maker is assumed to be interested in x^n_{pvh}, the number of patients of type p that are shifted to hospital h in vehicle type v when at node n in the scenario tree, and y^n_{vh}, the number of vehicles of type v used to transport patients to hospital h at node n in the scenario tree. Vehicles can drop off patients and come back to the source hospital to transport more patients.

Table 4: Variables used in the formulation.

Symbol	*Description*
x^n_{pvh}	No. of p type patients transported in vehicle v to hospital h in time step $t(n)$ at node n in the scenario tree
y^n_{vh}	No. of vehicles of type v starting to hospital h in time step $t(n)$ at node n in the scenario tree

The objective minimizes a weighted sum of the expected risk and the expected cost of evacuation. The risk portion of the objective, weighted using $\psi \in [0, 1]$, has two components as mentioned earlier. The first component, containing $\sum_{n\in N} \pi^n(\sum_{p\in P} \sum_{v\in V} \sum_{h\in H} \lambda^n_{pvh} x^n_{pvh})$, denotes expected number of adverse events that result from shifting patients.

$$\min\ \psi \left\{ \sum_{n\in N} \pi^n \left(\sum_{p\in P} \sum_{v\in V} \sum_{h\in H} \lambda^n_{pvh} x^n_{pvh} \right) + \sum_{r\in R} \rho^r \left(\sum_{p\in P} \mu^r_p \Big(\alpha_p - \sum_{v\in V} \sum_{h\in H} \sum_{n\in r} x^n_{pvh} \Big) \right) \right\}$$
$$+ (1-\psi) \left\{ \sum_{n\in N} \pi^n \left(\sum_{p\in P} \sum_{v\in V} \sum_{h\in H} \sigma^n_p x^n_{pvh} + \sum_{v\in V} \sum_{h\in H} \omega^n_{vh} y^n_{vh} \right) + \sum_{r\in R} \rho^r \left(\sum_{p\in P} \sigma_p \Big(\alpha_p - \sum_{v\in V} \sum_{h\in H} \sum_{n\in r} x^n_{pvh} \Big) \right) \right\} \quad (4)$$

Subject to

$$\sum_{v\in V}\sum_{h\in H}\sum_{n\in r} x^n_{pvh} \leq \alpha_p \qquad \forall p\in P, r\in R \tag{5}$$

$$\sum_{v\in V}\sum_{n\in r} x^n_{pvh} \leq \kappa_{ph} \qquad \forall p\in P, h\in H, r\in R \tag{6}$$

$$\sum_{p\in P} \delta_{pv} x^n_{pvh} \leq \chi_v y^n_{vh} \qquad \forall v\in V, h\in H, n\in N \tag{7}$$

$$\sum_{h\in H} y^n_{vh} \leq \beta_v - \sum_{h\in H}\sum_{m\in N(n)} y^m_{vh} + \sum_{h\in H}\sum_{m\in N(n)} \zeta^{mn}_h y^m_{vh} \qquad \forall v\in V, n\in N \tag{8}$$

$$x^n_{pvh} \leq M\delta_{pv} \qquad \forall p\in P, v\in V, h\in H, n\in N \tag{9}$$

$$x^n_{pvh} \in \mathbb{Z}_+ \qquad \forall p\in P, v\in V, h\in H, n\in N \tag{10}$$

$$y^n_{vh} \in \mathbb{Z}_+ \qquad \forall v\in v, h\in H, n\in N \tag{11}$$

The second component, constructed using $\sum_{r\in R}\rho^r(\sum_{p\in P}\mu^r_p(\alpha_p - \sum_{v\in V}\sum_{h\in H}\sum_{n\in r} x^n_{pvh}))$, represents the expected number of adverse events due to sheltering-in-place. Note that $\alpha_p - \sum_{v\in V}\sum_{h\in H}\sum_{n\in r} x^n_{pvh}$ in the above expression is the number of patients of type p that remain at the hospital over the course of the storm. Notice also that the model does not simplify the decision to either complete evacuation of patients of type p or no evacuation. The number evacuated is the subject of the optimization. The risk of sheltering in place is computed for each sample path and is multiplied by the probability of observing the sample path (i.e., ρ^r). A similar procedure is used to construct the expected cost using two components—vehicle operating costs and staffing costs. The expected cost is weighted using $(1-\psi)$.

Constraint (5) ensures that the sum of the patients transported in different time steps is less than the total number of patients of each type along each sample path. Constraint (6) requires that the receiving hospitals have adequate capacity for all patient types. The capacity and conservation constraints for different vehicle types are represented in (7) and (8) respectively. In (7), the number of patients being evacuated in a particular vehicle v is constrained to be less than or equal to $\chi_v y^n_{vh}$, the total capacity of all departing vehicles of type v. Inequality (8) states that the number of vehicles of type v available for use at node n is less than or equal to fleet size of v type vehicles minus the number of vehicles that left in previous time steps (i.e., $\sum_{h\in H}\sum_{m\in N(n)} y^m_{vh}$) plus the number of vehicles that returned at or before the current time step (i.e., $\sum_{h\in H}\sum_{m\in N(n)} \zeta^{mn}_h y^m_{vh}$). Recall that vehicles were assumed to transport patients to a single hospital during each trip. We also use (9) to ensure that if a patient of type p cannot be transported in a vehicle of type v then x^n_{pvh} must be zero. The value of M in this inequality is set to a sufficiently large constant such as α_p. Finally, (10) and (11) restrict the x and y variables to take integral values.

Note that by using a variable for each node in the scenario tree, we ensure that similar decisions are taken at all scenarios which are indistinguishable from each other (non-anticipativity condition). One could also reformulate the problem using a variable for each scenario and explicitly state the non-anticipativity constraints. Such a model will have far more variables and constraints than the current formulation.

4 Demonstration

4.1 Data

In this section, we discuss a case study involving a hypothetical evacuation of Vidant Medical Center (formerly known as Pitt County Memorial Hospital), a Level I trauma center located in Pitt County, North Carolina, during Hurricane Isabel that made landfall on 18 September 2003. Nine neighboring general medical and surgical hospitals were chosen for receiving patients evacuated from Vidant Medical Center (see blue markers in Figure 3). The hospitals in the case study were chosen based on their geographical locations and the optimal mandatory evacuation orders for different areas as computed in Yi et al. (2017). While Vidant Medical Center is located in a zone receiving evacuation orders for some hurricane scenarios, the receiving hospitals are in zones that did not receive an evacuation order and hence were not under any threat. Evacuation operations were allowed three-and-a-half days before landfall, and decisions were assumed to be taken at one-hour intervals, resulting in 84 decision stages. The scenario clusters are, however, generated at six-hour intervals using the algorithm by Yang et al. (2017) since weather advisories, such as those issued by the NHC, are often provided at a six-hour time granularity (see the vertical grid lines in Figures 1b and 7).

Table 5: List of receiving hospitals.

S. No.	*Hospital Name*	*Number of Beds*	*Census*	*Available Beds*	*Round-Trip Dist.*
1	Cape Fear Valley Medical Center	660	590	70	227
2	Duke University Hospital	820	676	144	215
3	Halifax Regional Medical Center	142	78	64	162
4	Johnston UNC Health Care	142	92	50	129
5	Nash Health Care Systems	303	186	117	89.4
6	Rex Healthcare	660	496	164	170
7	Sampson Regional Medical Center	105	50	55	169
8	WakeMed Raleigh Campus	683	537	146	155
9	Wayne Memorial Hospital	274	137	137	77.7

Wilson Medical Center, indicated by the green marker in Figure 3, is used later in Section 4.2.5 to demonstrate long-term planning applications of the model. Hospital data on the number of beds and census (which denotes the average occupancy per day) were collected from the American Hospital Association guide (AHA, 2015a) and are shown in Table 5 along with round trip distances (in miles) to the evacuating hospital. Vidant Medical Center is reported to have 909 beds with an average occupancy of 713. For the purpose of the case study, we consider a worst-case scenario by assuming that the evacuating hospital is 100% full and that the receiving hospitals can only accommodate the average number of available beds shown in Table 5.

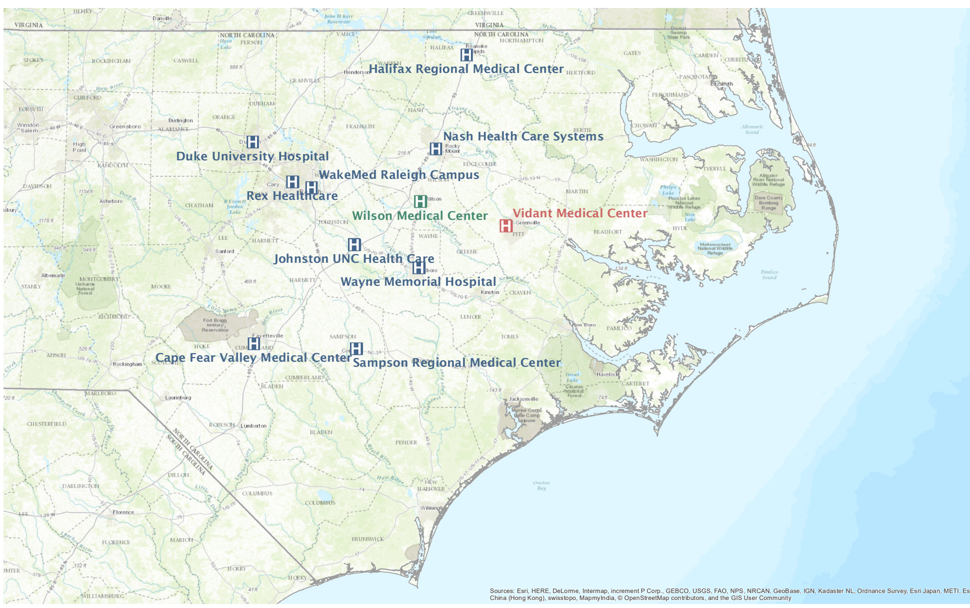

Figure 3: Hospital locations.

The evacuating hospital is assumed to have a fleet comprising four kinds of vehicles as shown in Table 6. Two types of ALS ambulances—ALS I and ALS II, which differ in their level of service, are supposed to be available for evacuation operations. ALS II ambulances offer more advanced health care when compared to ALS I ambulances and are relatively expensive. Vehicle operating costs are assumed to include per-trip costs that cover driver wages and emergency personnel and a per-mile cost of \$16.36 that accounts for fuel consumption. Round-trip distances to the receiving hospitals were used to compute approximate fuel costs. (One could use the lengths of TDSPs for more accurate estimates.) The rates for ambulances were retrieved from an emergency medical service (EMS) website of a county in North Carolina (`http://www.wakegov.com/ems/patient/Pages/feeandpatientcost.aspx`). Due to lack of information, expenses for renting a van are assumed to be the same as that of using an ALS II ambulance. (The base rates from the website were multiplied by a factor of 1.5 since vehicle operating costs are typically higher during a hurricane and because most trips involve traveling to another town or city.) Patient treatment costs were included in the objective for different scenarios only if the source hospital's zone is expected to receive mandatory evacuation orders. We used evacuation order predictions from Yi et al. (2017) for this purpose. According to AHA (2015b), hospitals in North Carolina incur an expenditure of $1,837.51 per patient per day. Since staff are paid overtime wages during an emergency, we suppose that the extra cost for sheltering a patient is 50% of $1,837.51 per day. (Similar assumptions on evacuation costs were made in past studies (Tayfur and Taaffe, 2009a).) The expenditure terms in the formulation depend on patient types, but we used average statistics since this data was unavailable.

Table 6: Fleet characteristics.

S. No.	*Vehicle Type*	*Number available*	*Capacity*	*Rate*
1	BLS	6	2	773.95
2	ALS I	4	1	919.06
3	ALS II	2	1	1330.23
4	Van	1	10	1330.23

The optimal travel times for vehicle trips under different scenarios were computed using TDSP algorithms on a roadway network. The network consists of U.S. interstate and state highways and has 3,596 nodes and 9,986 links. The travel time on a link is assumed to be time-dependent and varies every 15 minutes. Link travel times predicted by Yi et al. (2017) were used as input data, and the nodes closest (measured using the haversine formula) to the hospitals were used as a proxy for trip origins and destinations. Optimal travel times are then used to compute risks associated with transporting patients and in estimating ζ_h^{mn}. For this purpose, we first solve a one-to-all TDSP problem from the source hospital for different departure times using a label-correcting method. The optimal round trip times are then obtained by again using a one-to-all TDSP algorithm with the receiving hospitals as origins.

To classify patients into different groups based on their condition, we suppose that the evacuating hospital employs a triage system. Triage systems are widely used in the military, emergency departments, and mass casualty disasters to prioritize the treatment of patients. Patient classification is typically done on a 3- to 5-level system using triage scores computed from flow charts and algorithms that take into account comorbidity of patients, nature of the injury, and clinical descriptors such as Glassgow coma scale, systolic blood pressure, and respiratory rate (Champion et al., 1989; Iserson and Moskop, 2007; Gilboy et al., 2012). They are also sometimes described using colors or symbols (Kennedy et al., 1996). The World Medical Association (WMA, 2006) recommends a five-level triage for disaster victims in which level 1 patients are the most critical, and level 4 are the least critical. Level 5 is the expectant category, used for dead patients or are likely to die even if treated. Several other triage systems such as the Simple Triage and Rapid Treatment (START), Canadian Triage and Acuity Scale (CTAS), Abbreviated Injury Scale (AIS), and Emergency Severity Index (ESI) are also commonly used in practice. These systems vary both in terms of their sensitivity and the average time required for classification (Lerner et al., 2008).

Hospital evacuations are somewhat different from other scenarios requiring a triage system since doctors and nurses have access to detailed medical records of patients at the time of evacuation. Kelen et al. (2006) suggested a 5-level "reverse triage" or a disposition classification system to discharge patients from the hospital. Although patients in their study were discharged to create additional capacity, the same triage method may be used in the event of a hurricane. However, public data on patient demand for different classes of patients were unavailable. As a proxy, we used the proportions of different patient types admitted to emergency departments in the US in 2009 (Hing and Bhuiya, 2012, see Figure 5). Patient types, proportions, and demand for each type are shown in Table 7. We also assumed that the capacity at the receiving hospitals for different patient types is also similarly distributed.

Table 7: Triage classification and patient demand.

S. No.	*Patient Type*	*Percentage*	*Demand*
1	Minimum	7	64
2	Low	35	336
3	Moderate	41	391
4	High	10	100
5	Very High	2	18

Binary input data δ_{pv} is used to indicate if a type-p patient can be transferred in a vehicle of type v, as shown in Table 8. For instance, an ALS II ambulance with the highest level of medical care can be used to transfer all types of patients, whereas a van can only transfer patients whose risk level is 'minimum' or 'low.'

Table 8: Patient vehicle feasibility.

	BLS	ALS I	ALS II	Van
Minimum	1	1	1	1
Low	1	1	1	1
Moderate	1	1	1	0
High	0	1	1	0
Very High	0	0	1	0

The risk associated with evacuating λ^n_{pvh} is assumed to be due to (1) depleting medical resources, (2) flooding of the source hospital in periods prior to the departure of the patient, (3) limited medical care in ambulances, and (4) flooding of roadway links during transportation. We assume that these events are mutually independent for ease of representation and from a data availability perspective. Suppose a binary variable ϕ^n is set to 1 if the zone containing the source hospital receives a mandatory evacuation order at node n at time $t(n)$ and is 0 otherwise. Note also that an order once issued is never retracted, and hence if the ϕ value of a node in the scenario tree is set to 1, then all downstream nodes also receive an evacuation order. As mentioned earlier, the values for this parameter are obtained from the outputs of previous work by Yi et al. (2017).

Let γ^n_p denote the probability of an adverse event at node n for a patient of type p due to depleting resources such as power, staff, and equipment. This type of risk is considered only if $\phi^n = 1$ and is modeled as a function of the time elapsed since the issuance of the order. For this purpose, we use a linear function such as the one defined in Bish et al. (2014). The usage of predefined functional forms for measuring risk is somewhat limiting but one could use more sophisticated risk measures based on past data if available. If $\phi^n = 1$, and if l denotes the ancestor of node n at which the mandatory evacuation order was first issued, then we assume $\gamma^n_{\text{min}} = 0.02(t(n) - t(l))/|T|$, where $t(n) - t(l)$ is the number of time steps elapsed since the issuance of a mandatory evacuation order. We further assumed that $\gamma^n_{\text{low}} = 2\gamma^n_{\text{min}}$, $\gamma^n_{\text{moderate}} = 3\gamma^n_{\text{min}}$, and so on. Hence, if a patient of type p is evacuated at node n at time $t(n)$, the probability that an adverse event did not occur in preceding time steps is given by $\prod_{m\in N(n)}(1 - \gamma^m_p)^{\phi^m}$.

The second risk component ξ^n denotes the probability of an adverse event due to flooding of the source hospital. To obtain these parameters, we use outputs from a hydrological model Coupled Routing and Excess Storage (CREST) (Wang et al., 2011) and a hurricane storm surge model ADvanced CirCulation (ADCIRC) (Westerink et al., 2008; Dietrich et al., 2011) and read off the corresponding risk due to flood depth and wind speed from Figure 4 which is taken from Apivatanagul et al. (2012). These two risk measures are combined to obtain ξ^n, which is then used to calculate $\prod_{m\in N(n)}(1 - \xi^m)$, the probability of non-occurrence of an adverse event due to flooding of the source hospital. The risk values for 'shelter' in Figure 4 are used as a proxy for the risk due to winds and flooding at the source hospital.

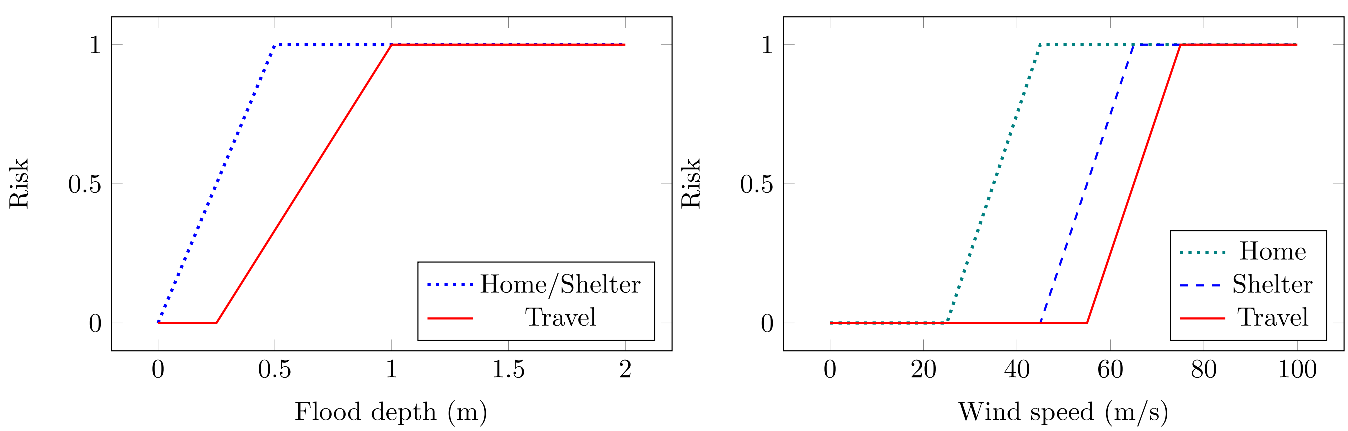

Figure 4: Risk factors as a function of flood depth (left) and wind speed (right).

Patients are at additional risk during transportation because the medical services in ambulances do not match those provided at a hospital. To model this type of risk, we follow a method similar to Bish et al. (2014) in which the transportation risk φ^n_{pvh} of a patient of type p moved using vehicle v to hospital h at node n at time $t(n)$ is a function of the travel time. Specifically, if κ_{pv} is the probability of an adverse event for different vehicle types carrying different patients in one TDSP time interval (i.e., 15 minutes) and τ_h^n is the number of such time intervals in the shortest trip for a vehicle headed to hospital h at node n, then $\varphi_{pvh}^n = 1 - (1 - \kappa_{pv})^{\tau_h^n}$. The values of κ_{pv} for 'minimum' risk patients for BLS, ALS I, ALS II, and Van are assumed to equal 4E-05, 2E-05, 2E-06, and 2E-04 respectively. Just like the γ values, we then assumed that $\kappa_{\text{low},v} = 2\kappa_{\text{min},v}$, $\kappa_{\text{moderate},v} = 3\kappa_{\text{min},v}$, and so on.

Finally, moving patients during a hurricane also carries risk due to flooding of highway links. Previously mentioned wind speed and storm surge models also provide the risk on each link in the road network and using the shortest path information, one can compute the risk due to flooding while traveling ϑ_h^v. For instance, if P_h^n denotes the set of links belonging to the shortest path to hospital h when departing at node n in the scenario tree, then $\vartheta_h^n = 1 - \prod_{(i,j)\in P_h^n}(1 - \varsigma_{ij})$, where ς_{ij} is the flooding risk on highway link (i, j). Using these four risk components, the risk of evacuating a patient of type p in a vehicle of type v to hospital h at time step $t(n)$ at node n in the scenario tree is given by (12).

$$\lambda_{pvh}^n = 1 - \left(\prod_{m\in N(n)} (1 - \gamma_p^m)^{\phi^m} \cdot \prod_{m\in N(n)} (1 - \xi^m) \cdot (1 - \varphi_{pvh}^n) \cdot (1 - \vartheta_h^n) \right) \tag{12}$$

The risk associated with sheltering-in-place is constructed similarly except that it just involves the first two components of the evacuation risk. Under this assumption, the risk of not evacuating a type p patient when sample path r is realized is given by

$$\mu_p^r = 1 - \left(\prod_{n\in r} (1 - \gamma_p^n)^{\phi^n} \cdot \prod_{n\in r} (1 - \xi^n) \right) \tag{13}$$

Although the risk due to lack of resources would continue to increase over time, we assume that it remains constant after the end of the modeling time horizon.

4.2 Results

The optimization model (4)-(11) was solved using the data described in Section 4.1 and the CPLEX Callable Library in Java on a Windows machine with an Intel Core i7-6700 CPU @ 3.40GHz, 16 GB RAM, and 8 MB L3 Cache. The problem instance has a total of 195,480 integer variables and 85,920 constraints.

4.2.1 Efficiency Frontier

We first solved linear programming (LP) relaxations of the problem with different weights ψ to produce an efficiency frontier with risk on the horizontal axis and cost on the vertical axis (see Figure 5). The ψ values associated with the solutions on the efficiency frontier creates a partition on the interval $[0, 1]$ and one can use the following bisection method to enumerate solutions on the frontier. First, LP solutions for $\psi = 0$ and $\psi = 1$ are evaluated. The problem is then re-solved setting $\psi = 0.5$, and if the objective components do not coincide with the previous solutions, two sub-intervals $[0, 0.5]$ and $[0.5, 1]$ are created, and two new LPs are solved using their mid-points, i.e., $\psi = 0.25$ and 0.75, and the process is continued. If the LP objectives for a ψ value used in an intermediate iteration coincides with that obtained from its left or right endpoint, we refrain from partitioning the associated sub-interval. For example, if $\psi = 0.75$ yielded the same LP solution as that of $\psi = 0.5$, then the interval $[0.5, 0.75]$ is ignored. When no further partitions are created, the procedure enumerates all points on the efficiency frontier, but we terminated this process after discovering about 300 solutions.

Although an efficiency frontier of non-dominated integer solutions can be constructed using the methods suggested by Klein and Hannan (1982) and Sylva and Crema (2004), solving the LP relaxed versions allows hospital administrators to quickly approximate the risks and costs for different weights (each LP takes a few seconds to find the optimal solution), thereby allowing them to select an appropriate weight and resolve the problem as an integer program (IP) for further analysis. For illustration purposes, we picked three weights 0, 0.999878, and 1, the IP solutions for which are represented using red triangles in Figure 5. The problem instance with $\psi = 0$ minimizes cost while $\psi = 1$ minimizes risk. The percentage gap between the LP relaxation and the IP objectives for the three instances was about 0.3%. Although the integrality gap was small, many of the optimal decision variables in the LP relaxation were fractional, and rounding them could lead to infeasibilities. Hence, we solved the IPs using CPLEX's MIP solver.

The IPs took about 10 minutes to find solutions with a gap of 0.02%. In the remaining subsections, we take a closer look at the IP solutions to these three instances.

Figure 6 shows a breakdown of the objective function components. The left panel indicates the expected risk associated with evacuating $\sum_{n\in N}\pi^n(\sum_{p\in P}\sum_{v\in V}\sum_{h\in H}\lambda^n_{pvh}x^n_{pvh})$ and sheltering-in-place $\sum_{r\in R}\rho^r(\sum_{p\in P}\mu^r_p(\alpha_p - \sum_{v\in V}\sum_{h\in H}\sum_{n\in r}x^n_{pvh}))$ and the right panel shows the expected transportation $\sum_{n\in N}\pi^n(\sum_{p\in P}\sum_{v\in V}\sum_{h\in H}\sigma^n_p x^n_{pvh} + \sum_{v\in V}\sum_{h\in H}\omega^n_{vh}y^n_{vh})$ and staffing costs $\sum_{r\in R}\rho^r(\sum_{p\in P}\sigma_p(\alpha_p - \sum_{v\in V}\sum_{h\in H}\sum_{n\in r}x^n_{pvh}))$.

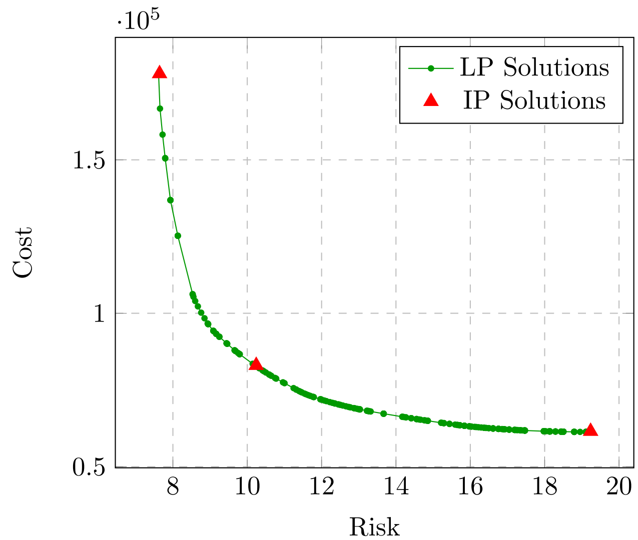

Figure 5: Efficiency Frontier of LP solutions.

When $\psi = 0$, the objective is to minimize cost, in which case it is cheaper not to evacuate many patients. Thus, the risk of sheltering-in-place and staffing cost is high. As ψ increases, minimizing risk becomes important, which prompts more evacuations, resulting in higher evacuation risk and transportation cost but lower shelter-in-place risk and staffing cost.

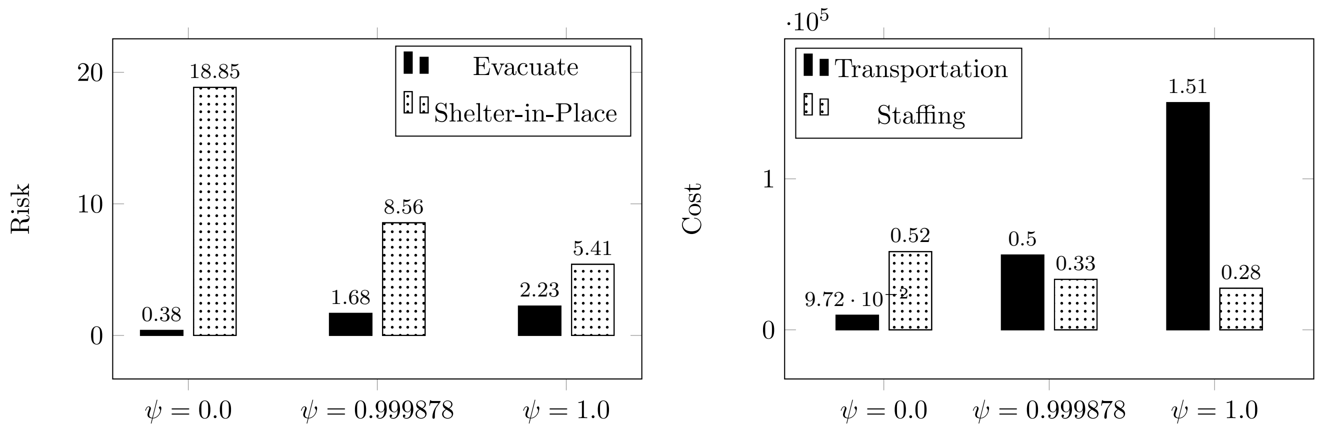

Figure 6: Risk (left) and cost in dollars (right) components of optimal IP solutions.

4.2.2 Model Performance under Different Scenarios

The number of patients evacuated along each sample path for the three weights are shown in Figure 7. According to Yi et al. (2017), the highlighted portion (along sample path terminating in Scenario 5) indicates the nodes at which the zone containing the source hospital was given an evacuation order.

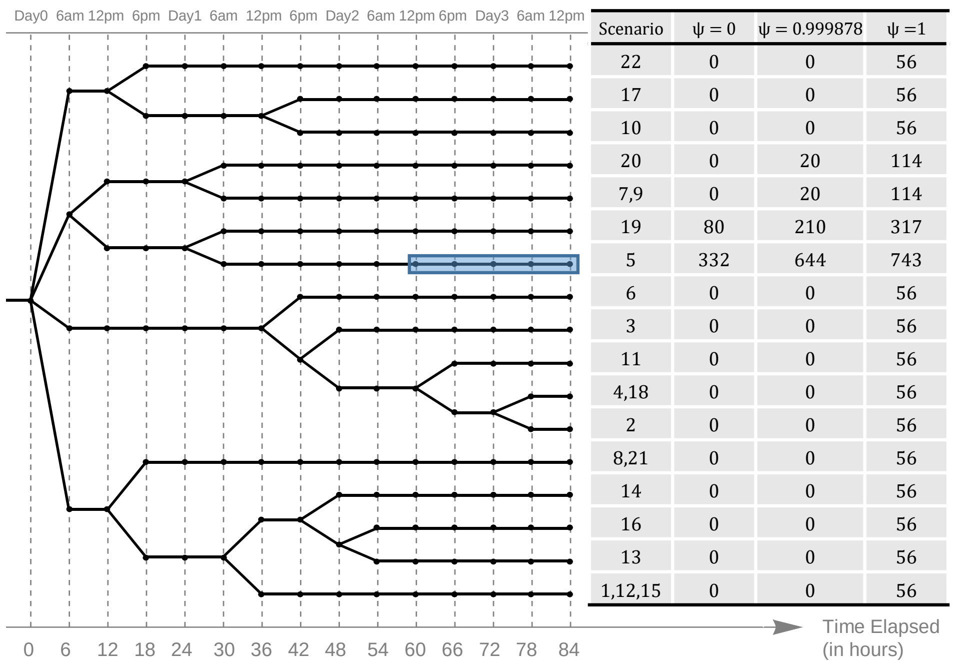

Scenario	ψ = 0	ψ = 0.999878	ψ =1
22	0	0	56
17	0	0	56
10	0	0	56
20	0	20	114
7,9	0	20	114
19	80	210	317
5	332	644	743
6	0	0	56
3	0	0	56
11	0	0	56
4,18	0	0	56
2	0	0	56
8,21	0	0	56
14	0	0	56
16	0	0	56
13	0	0	56
1,12,15	0	0	56

Figure 7: Number of patients evacuated along different sample paths.

Since staffing costs enter the objective only at these highlighted nodes, when ψ is zero (i.e., when the cost is being minimized), no patients are evacuated along most sample paths. When minimizing risk, the sample path ending in Scenario 5 witnessed the maximum number of evacuations followed by sample paths that contain Scenario 5 during early time periods. For instance, Scenario 19 cannot be distinguished from Scenario 5 until time step 30, and hence when $\psi = 1$, 317 patients are evacuated even if Scenario 19 were to occur. For similar reasons, patients are preemptively evacuated along sample paths terminating in Scenarios 7, 9, and 20, but the evacuation operations are terminated as soon as it becomes clear that these scenarios will not occur. Also, notice that only 743 of the 909 patients are evacuated because of limited vehicle capacity and higher evacuation risk during later time periods.

Table 9: Objective function components along sample paths.

Scenario	$\psi = 0$		$\psi = 0.999878$		$\psi = 1$	
	Risk	*Cost ('000)*	*Risk*	*Cost ('000)*	*Risk*	*Cost ('000)*
22	0	0	0	0	0.03	78.50
17	0	0	0	0	0.03	78.50
10	0	0	0	0	0.03	78.50
20	0	0	0.03	6.47	0.07	158.96
7,9	0	0	0.03	6.47	0.07	158.96
19	0.08	22.49	0.14	227.96	0.19	398.32
5	423.11	1331.04	225.10	1578.90	167.11	1707.88
6	0	0	0	0	0.03	78.50
3	0	0	0	0	0.03	78.50
11	0	0	0	0	0.03	78.50
4,18	0	0	0	0	0.03	78.50
2	0	0	0	0	0.03	78.50
8,21	0	0	0	0	0.03	78.50
14	0	0	0	0	0.03	78.50
16	0	0	0	0	0.03	78.50
13	0	0	0	0	0.03	78.50
1,12,15	0	0	0	0	0.03	78.50
Expected Value	19.24	61.52	10.24	83.01	7.64	178.08

The risk and cost (in thousands of dollars) components of the objective function across sample paths are shown in Table 9. The values indicate the performance of the optimal solutions (that are designed in response to the stochasticity in the hurricane trajectory) for each of the 22 scenarios and the variance in the outcomes. It is interesting to note that if Scenario 5 occurred, one could drastically reduce the risk or the number of casualties by just spending about 10-20% more. However, setting ψ closer to 1 results in unnecessary risk and expenditure if the hurricane does not affect the source hospital.

4.2.3 Distribution of Patients across Hospitals

Table 10 shows the distribution of patients admitted at different receiving hospitals in expectation across all sample paths and for the sample path terminating in Scenario 5.

Table 10: Patient admissions at receiving hospitals.

Hospital	*Beds*	*Dist.*	*Expected no. of patients*			*Sample path of Scenario 5*		
			$\psi = 0$	$\psi = 0.999878$	$\psi = 1$	$\psi = 0$	$\psi = 0.999878$	$\psi = 1$
Cape Fear Valley Medical Center	70	227	0.00	0.00	1.27	0	0	28
Duke University Hospital	144	215	0.00	0.00	0.18	0	0	4
Halifax Regional Medical Center	64	162	0.91	2.91	3.32	20	64	64
Johnston UNC Health Care	50	129	1.82	4.09	44.41	20	50	50
Nash Healthcare Systems	117	89.4	4.64	5.77	13.45	102	117	117
Rex Healthcare	164	170	0.00	3.91	14.50	0	82	155
Sampson Regional Medical Center	55	169	0.45	2.32	2.45	10	49	51
WakeMed Raleigh Campus	146	155	2.73	9.95	11.05	60	145	137
Wayne Memorial Hospital	137	77.7	8.18	12.59	16.36	120	137	137
Total			18.73	41.55	107.00	332	644	743

As expected, hospitals that are closer are almost filled while those that are far away are not fully utilized since the transportation risk is a function of travel time, flood depths, and wind speeds along highway links. Hence, when the goal of the problem is to minimize risk, only a few patients are sent to Cape Fear Valley Medical Center and Duke University Hospital. These types of results are extremely useful because they allow administrators to inform receiving hospitals of the number of patients that may be expected. Such information allows the receiving hospitals to ensure adequate medical resources and, if necessary, also decompress for creating sufficient in-patient capacity. Notice that there is a wide difference in the results for the expected and the extreme case (Scenario 5). Thus, it is also essential for hospital administrators to relay updated estimates from time to time based on the scenarios that are unlikely to occur.

4.2.4 Patient Types and Evacuation Strategies

The results also provide insights into how patients of different types must be prioritized and evacuated. It is often unclear if the most critical patients must be evacuated first since they would get better care at the source hospital than during transportation. The plots in Figure 8 indicate the optimal fraction of patients of each type that are evacuated over time for three choices of weights ψ under Scenario 5. Suppose r_5 denotes the sample path containing Scenario 5.

The x-axis represents time in hours and the y-axis indicates the cumulative fraction of evacuated patients of different types, i.e., $\frac{1}{\alpha_p}(\sum_{n\in r_5}\sum_{v\in V}\sum_{h\in H} x^n_{pvh})$. Each plot contains the fraction evacuated across time for the five patient types assumed in Table 7. When $\psi = 0$, the risk does not feature in the objective function, and hence the most critical patients are never evacuated. For $\psi = 0.999878$ and $\psi = 1$, a higher fraction of critical patients are evacuated when compared with less critical patients. More interestingly, notice the start times and the rate at which different patients are evacuated. It appears that once we start evacuating critical patients, it is optimal to prioritize their departures over the others.

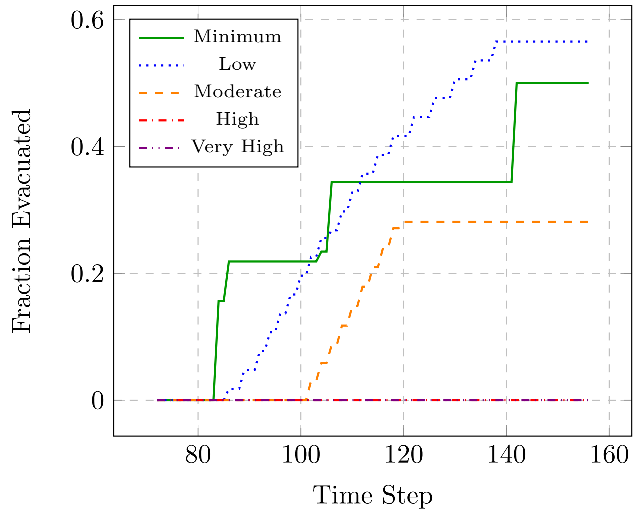

(a) $\psi = 0$

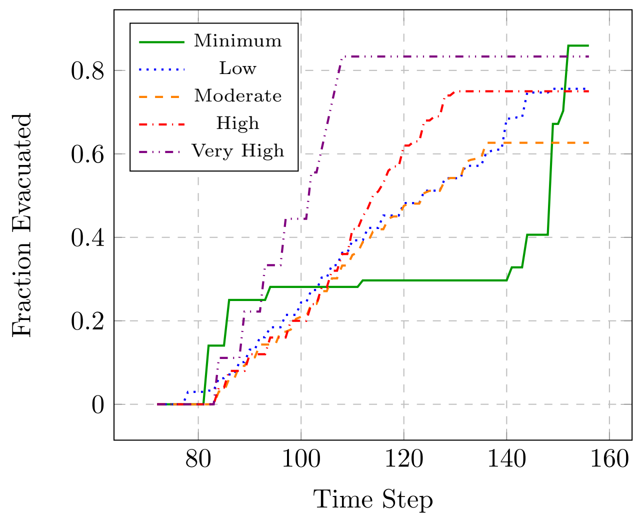

(b) $\psi = 0.999878$

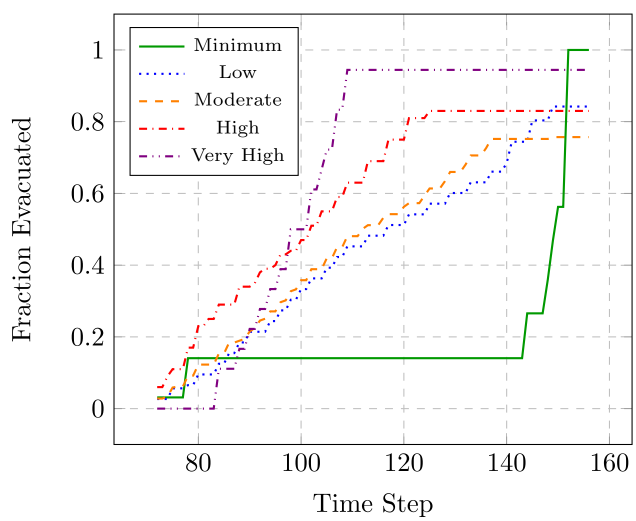

(c) $\psi = 1$

Figure 8: Patient evacuation over time for different weights.

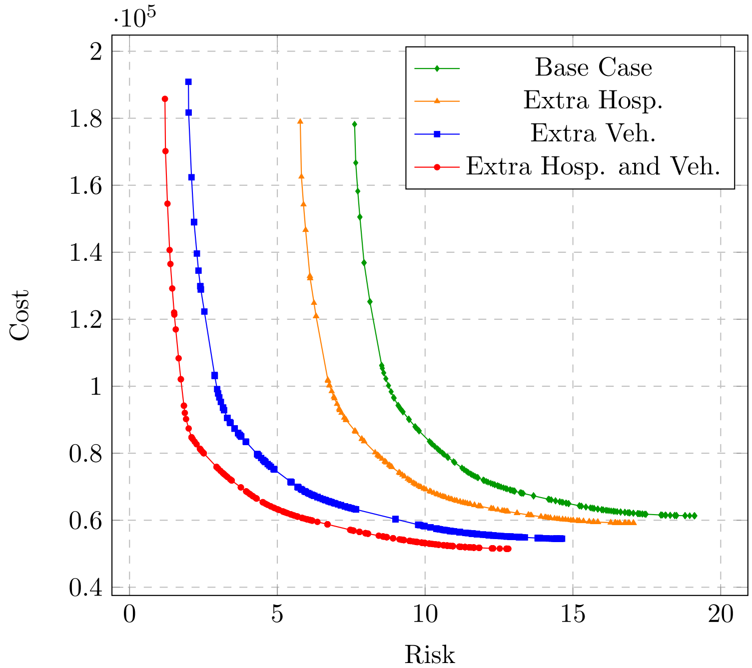

Figure 9: Efficiency frontiers for different cases.

4.2.5 Long-term Planning Applications

The model proposed in this paper can not only be used to evacuate patients before and during a hurricane but can also be used to make long-term strategic decisions related to agreements with receiving hospitals and vehicle contracts. To demonstrate this, we constructed efficiency frontiers of LP relaxations (Figure 9) for the following four cases:

(a) A base case scenario that was discussed so far.

(b) An agreement to send patients is made with a new hospital—Wilson Medical Center (indicated by the green marker in Figure 3), which has a capacity of 117 beds.

(c) The vehicle fleet in the base case is expanded by adding 2 BLS ambulances and 1 van, but the receiving hospitals are kept the same.

(d) An extra receiving hospital and extra vehicles are available during evacuation (i.e., both (b) and (c)).

Since more resources are available in Cases (b), (c), and (d), the efficiency frontier shifts closer to the origin. Further, as we did not include any fixed cost for making mutual agreements, adding a new hospital in Case (b) gives solutions that lower the risk and the cost for all weights ψ. On the other hand, expanding the fleet size has direct implications on vehicle operating costs, and hence the least risky solutions are slightly more expensive than the base case. Thus, in this example, one could prioritize increasing fleet size over finding other partner hospitals to tie-up with. Similar analyses can be conducted to determine the type of vehicles to add to the current fleet.

4.2.6 Value of Information

In this section, we compare the multi-stage stochastic optimization framework with deterministic counterparts in which (a) the model parameters are replaced by their expected values, which we will refer to as the expected value (EV) problem and (b) estimate the expected value of perfect information (EVPI). In both cases, we define the decision variables at each time step instead of defining them at each node in the scenario tree and accordingly modify (4)-(11). Specifically, we use x^t_{pvh} and y^t_{vh} to represent the number of p type patients transported in vehicle v to hospital h at time t and the number of vehicles of type v heading to hospital h at t respectively.

In the EV problem, we replace the node-dependent model parameters with a parameter for each time step. For example, instead of λ^n_{pvh}, the risk associated with evacuating a patient at node n in the scenario tree, we use λ^t_{pvh} that represents the risk due to evacuating at time t and is computed by averaging the λ^n_{pvn} values across the nodes in the scenario tree at time t. In this model, at time t, decisions x^t_{pvh} and y^t_{vh} are oblivious to new information about the hurricane scenarios. The results of the EV problem are summarized in Table 11 and one can conclude that the EV model performs poorly since its cost and risk estimates are very different from that of the stochastic optimization model. Since these solutions do not adapt to new information, they may grossly underestimate or overestimate the need for evacuation. For example, if the cost is minimized, i.e., when $\psi = 0$, no patients are evacuated even if Scenario 5 were to occur. Similarly, when minimizing risk, i.e., if $\psi = 1$, a total of 743 patients are evacuated over the entire time horizon, which is unnecessary for most scenarios.

Table 11: Optimal solutions to the multi-stage EV and a single-stage evacuation problem.

Weight	$\psi = 0$	$\psi = 0.999878$	$\psi = 1$
Single-stage evacuation			
Risk	28.93	28.64	28.02
Cost ('000)	75.92	77.69	103.76
Total patients evacuated	0	10	28
Multi-stage EV Problem			
Risk	28.93	25.87	8.09
Cost ('000)	75.92	96.31	1129.42
Total patients evacuated	0	110	743
Multi-stage Stochastic Optimization Model			
Risk	19.24	10.24	7.64
Cost ('000)	61.52	83.01	178.08
Expected patients evacuated	19	42	107

We also ran the integer programs using a single decision epoch (set to three-and-a-half days before landfall) to illustrate the multi-stage model's advantages. When the ψ values is 0, the results are similar to the EV problem because it was optimal to evacuate nobody (see Table 11). However, when minimizing risk, the single-stage model evacuates only 28 patients and results in nearly 20 extra expected casualties than the EV or the stochastic optimization solution. If Scenario 5 occurred, this solution produces a risk component of 616, which is nearly four times the optimal risk of the stochastic optimization solution (refer to Table 9). The cost of this single decision epoch model is lower since vehicles do not make multiple trips.

To estimate the EVPI, we first computed wait-and-see solutions by solving a deterministic integer program for each of the 22 scenarios. The decision variables for each of these problems are again x^t_{pvh} and y^t_{vh}. Since they vary across scenarios, we can write them as x^t_{pvhs} and y^t_{vhs} to indicate the solutions for scenario s, but we skip adding another subscript for brevity. These problems provide the optimal sequence of decisions that a hospital manager could have taken in hindsight under each scenario. Assuming complete knowledge of a scenario lets us estimate the flood depths, wind speeds, and travel times on highway links for different departure times exactly, and these estimates can be used to derive the integer program parameters. For instance, ζ^{mn}_h in (8) is replaced with $\zeta^{tt'}_h$ and is set to 1 if a vehicle leaving hospital h at t can return at or before t' under the given scenario. The objective function of this wait-and-see problem is shown in (14).

$$\min \ \psi\left\{\sum_{p\in P}\sum_{v\in V}\sum_{h\in H}\sum_{t\in T}\lambda^t_{pvh}x^t_{pvh} + \sum_{p\in P}\mu_p\Big(\alpha_p - \sum_{v\in V}\sum_{h\in H}\sum_{t\in T}x^t_{pvh}\Big)\right\}$$
$$+ (1-\psi)\left\{\sum_{p\in P}\sum_{v\in V}\sum_{h\in H}\sum_{t\in T}\sigma^t_p x^t_{pvh} + \sum_{v\in V}\sum_{h\in H}\sum_{t\in T}\omega^t_{vh}y^t_{vh} + \sum_{p\in P}\sigma_p\Big(\alpha_p - \sum_{v\in V}\sum_{h\in H}\sum_{t\in T}x^t_{pvh}\Big)\right\} \quad (14)$$

As mentioned earlier, for the scenario tree used in this paper, Yi et al. (2017) recommends a mandatory evacuation order only for Scenario 5, and in all other cases, we found that the optimal wait-and-see solutions were to not evacuate any patients. For Scenario 5, the optimal cost and risk estimates and the number of patients evacuated are shown in Table 12.

Table 12: Wait-and-see solutions for Scenario 5.

	$\psi = 0$	$\psi = 0.999878$	$\psi = 1.0$
Risk	313.73	107.31	97.99
Cost ('000)	1212.85	1636.41	1759.44
Expected patients evacuated	496	822	847

Scenario 5 represents the most severe outcome as the hurricane passes close to the evacuating hospital (see Figure 1a and 3). Since perfect information is assumed, the objective estimates of the wait-and-see solution are lower than that predicted for Scenario 5 under the stochastic optimization model (see Table 9). The EVPI is the difference in the optimal objective values of the stochastic model and the average of the wait-and-see solutions across all the 22 scenarios and is summarized in Table 13.

Table 13: Risk and cost components of the expected value of perfect information.

	$\psi = 0$		$\psi = 0.999878$		$\psi = 1$	
	Risk	*Cost ('000)*	*Risk*	*Cost ('000)*	*Risk*	*Cost ('000)*
Stochastic optimization objective	19.24	61.52	10.24	83.01	7.64	178.08
Expected wait-and-see solution	14.26	55.13	4.88	74.38	4.45	79.97
EVPI	4.98	6.39	5.36	8.63	3.19	98.11

The EVPI indicates the advantage we derive from knowing the future. For instance, when $\psi = 1$, we could reduce the expected number of casualties by 3.19 and save \$98,110 on an average had we known the storm path and characteristics accurately.

5 Conclusions

5.1 Summary

In this paper, we addressed the problem of evacuating hospitals during hurricanes using a multi-stage stochastic program, which identifies the optimal number of patients of different types who must be evacuated at different time periods to different hospitals and the vehicles used to transport them. The objective of the problem was to minimize a linear combination of risk and cost associated with evacuation. The risk component was assumed to comprise risk from sheltering-in-place and evacuating, which was influenced by flood depth, wind speeds, and time-dependent network congestion. The cost component, on the other hand, was assumed to be due to transportation and staffing. Since a hurricane's evolution is subject to considerable uncertainty, we used a scenario-tree-based approach for representing uncertainty. Predicted hurricane trajectories were used to construct a scenario tree, which formed the basis for a stochastic optimization problem. Comparisons with deterministic versions of the problem were also made to illustrate the utility of an adaptive decision-making framework.

5.2 Societal Benefits and Managerial Implications

Hurricanes trigger large-scale evacuations in the US every year. Often, during such disasters, evacuation of vulnerable populations such as hospital patients is a complex problem because shifting them is an expensive affair, and more importantly, moving them and not moving them can put them at risk. As evidenced by past experiences, some of which have been summarized in Section 1, evacuating a hospital during a hurricane is an extremely challenging problem, and one must weigh the risks associated with evacuating and sheltering-in-place. Evacuation is also an expensive affair since many scarce resources have to be pooled on short notice, and discontinuing regular operations has a high opportunity cost. Thus, hospital managers must possess a high degree of situational awareness for a successful evacuation (Downey et al., 2013). Most past facilities have relied on instincts to evacuate their patients rather than follow a structured decision-making process (McGinty et al., 2017). It is very difficult to an individual or group of individuals to perform the complex trade-offs required and therefore some form of decision-support is needed. Our paper fills this gap using a system of systems approach that fuses data from multiple domains such as climate, health, and transportation. The proposed solution is data- and technology-dependent and uses hurricane forecasts and network conditions, both of which have a significant bearing on the risk and cost of evacuating.

A case study featuring Hurricane Isabel and hospital statistics from North Carolina was used to answer several questions that would otherwise be difficult for hospital managers to address based on experience in the stressful time as exists when a hurricane is approaching. These operational and planning problems include prioritizing patients for evacuation, assigning them to different vehicles and receiving hospitals, updating decisions as more information on the hurricane becomes available, determining ideal vehicle composition and fleet size, and identifying a set of receiving hospitals. Using our model, hospital administrators will be able to make arrangements that adapt to dynamic weather conditions and control the trade-off between the cost of evacuating patients and the risks involved. They will also be able to get insights into why specific staging sequences are recommended by the model, which will enable them to make informed choices and better explain the proposed evacuation procedures to patients and staff.

In the future, to make the proposed models practical, it could be run as an application using specific queries that provide them real-time responses on whom to evacuate and how. For better model fidelity, it is paramount that good data from several federal and local agencies also be available in real-time. In the US, organizations such as FEMA started a data-sharing platform called OpenFEMA since Hurricane Sandy (OpenFEMA, 2020). Likewise, post-COVID, several open datasets are also available on hospitals and their capacities (HIFLD, 2021; OpenDataDC, 2021). The emergency evacuation applications for hospitals could use such standardized APIs, making it convenient for large-scale deployments and hospital staff training.

Although our empirical analysis was performed in the context of Atlantic hurricanes and data from hospitals in the US, the framework can be extended in its current form to other regions in the world that are affected by tropical storms. For example, agencies such as India Meteorological Department and Japan Meteorological Agency also provide advisories and forecasts similar to NHC for cyclones and typhoons originating in the Indian and Pacific oceans. Every year, these storms trigger massive evacuations in densely populated coastal cities of Asia, many of which also contain large hospitals and critical care facilities.

5.3 Limitations and Future Research

Further research is warranted in at least a couple of areas. First, characterizing the risk of travel by patient category and vehicle type is very useful. We have adopted a structure similar to other authors; however, additional research is necessary to formalize how to represent these risks. Second, we focus on a single hospital for the evacuation decisions. In large storms, often, there is a collection of hospitals making similar choices. Further, there may be a fixed fleet of vehicles that are to be shared among the hospitals or a large number of scenarios, which makes it difficult to get optimal solutions to integer programs quickly. In such cases, we could explore scenario-based decomposition heuristics such as progressive hedging (Rockafellar and Wets, 1991; Watson and Woodruff, 2011). Understanding how to make these decisions is substantially more complicated.

Acknowledgments

The authors would like to acknowledge the support of NSF (Grant Nos. 0826832 and 1331269) for conducting this research. Comments from three reviewers and the guest editor are also greatly appreciated.

References

Adalja, A. A., M. Watson, N. Bouri, K. Minton, R. C. Morhard, and E. S. Toner (2014). Absorbing citywide patient surge during Hurricane Sandy: A case study in accommodating multiple hospital evacuations. *Annals of Emergency Medicine 64*(1), 66 – 73.e1.

AHA (2015a). AHA guide to the health care field. Health Forum LLC.

AHA (2015b). AHA hospital statistics. Health Forum LLC.

AHRQ (2010). Hospital evacuation decision guide. https://archive.ahrq.gov/prep/hospevacguide/. Accessed: 2016-12-23.

An, S., N. Cui, X. Li, and Y. Ouyang (2013). Location planning for transit-based evacuation under the risk of service disruptions. *Transportation Research Part B: Methodological 54*, 1–16.

Apivatanagul, P., R. A. Davidson, and L. K. Nozick (2012). Bi-level optimization for risk-based regional hurricane evacuation planning. *Natural Hazards 60*(2), 567–588.

Bagaria, J., C. Heggie, J. Abrahams, and V. Murray (2009). Evacuation and sheltering of hospitals in emergencies: a review of international experience. *Prehospital and disaster medicine 24*(5), 461–467.

Bayram, V. (2016). Optimization models for large scale network evacuation planning and management: A literature review. *Surveys in Operations Research and Management Science 21*(2), 63–84.

Birge, J. R. and F. Louveaux (2011). *Introduction to stochastic programming.* Springer Science & Business Media.

Bish, D. R. (2011). Planning for a bus-based evacuation. *OR spectrum 33*(3), 629–654.

Bish, D. R., E. Agca, and R. Glick (2014). Decision support for hospital evacuation and emergency response. *Annals of Operations Research 221*(1), 89–106.

Bish, D. R. and H. D. Sherali (2013). Aggregate-level demand management in evacuation planning. *European Journal of Operational Research 224*(1), 79–92.

Blanton, B., K. Dresback, B. Colle, R. Kolar, H. Vergara, Y. Hong, N. Leonardo, R. Davidson, L. Nozick, and T. Wachtendorf (2020). An integrated scenario ensemble-based framework for hurricane evacuation modeling: Part 2—Hazard modeling. *Risk analysis 40*(1), 117–133.

Chabini, I. (1998). Discrete dynamic shortest path problems in transportation applications: Complexity and algorithms with optimal run time. *Transportation Research Record: Journal of the Transportation Research Board 1645*, 170–175.

Champion, H. R., W. J. Sacco, W. S. Copes, D. S. Gann, T. A. Gennarelli, and M. E. Flanagan (1989). A revision of the trauma score. *Journal of Trauma and Acute Care Surgery 29*(5), 623–629.

Chen, X. and F. B. Zhan (2008). Agent-based modelling and simulation of urban evacuation: relative effectiveness of simultaneous and staged evacuation strategies. *Journal of the Operational Research Society 59*(1), 25–33.

Cheng, G., C. G. Wilmot, and E. J. Baker (2008). A destination choice model for hurricane evacuation. In *Proceedings of the 87th Annual Meeting Transportation Research Board, Washington, DC, USA*, pp. 13–17.

Childers, A., G. Visagamurthy, and K. Taaffe (2009). Prioritizing patients for evacuation from a health-care facility. *Transportation Research Record: Journal of the Transportation Research Board 2137*, 38–45.

Chiu, Y.-C. and H. Zheng (2007). Real-time mobilization decisions for multi-priority emergency response resources and evacuation groups: model formulation and solution. *Transportation Research Part E: Logistics and Transportation Review 43*(6), 710–736.

Cova, T. J. and J. P. Johnson (2003). A network flow model for lane-based evacuation routing. *Transportation research part A: Policy and Practice 37*(7), 579–604.

Dietrich, J., M. Zijlema, J. Westerink, L. Holthuijsen, C. Dawson, R. A. Luettich, R. Jensen, J. Smith, G. Stelling, and G. Stone (2011). Modeling hurricane waves and storm surge using integrally-coupled, scalable computations. *Coastal Engineering 58*(1), 45–65.

Dixit, V. and B. Wolshon (2014). Evacuation traffic dynamics. *Transportation research part C: emerging technologies 49*, 114–125.

Dosa, D. M., N. Grossman, T. Wetle, and V. Mor (2007). To evacuate or not to evacuate: Lessons learned from Louisiana nursing home administrators following Hurricanes Katrina and Rita. *Journal of the American Medical Directors Association 8*(3), 142 – 149.

Downey, E. L., K. Andress, and C. H. Schultz (2013). External factors impacting hospital evacuations caused by Hurricane Rita: the role of situational awareness. *Prehospital and disaster medicine 28*(03), 264–271.

Dupačová, J. and K. Sladký (2002). Comparison of multistage stochastic programs with recourse and stochastic dynamic programs with discrete time. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik 82*(11-12), 753–765.

Franco, C., E. Toner, R. Waldhorn, B. Maldin, T. O'Toole, and T. V. Inglesby (2006). Systemic collapse: medical care in the aftermath of Hurricane Katrina. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science 4*(2), 135–146.

French, E. D., M. L. Sole, and J. F. Byers (2002). A comparison of nurses' needs/concerns and hospital disaster plans following Florida's Hurricane Floyd. *Journal of Emergency Nursing 28*(2), 111 – 117.

Gibbs, L. and C. Holloway (2013). Hurricane Sandy after action: report and recommendations to Mayor Michael R. Bloomberg. *Hurricane Sandy After Action: Report and Recommendations to Mayor Michael R. Bloomberg, The City of New York 36.*

Gilboy, N., P. Tanabe, D. Travers, and A. M. Rosenau (2012). Emergency severity index (ESI): a triage tool for emergency department care, version 4. *Implementation handbook*, 12–0014.

Gray, B. H. and K. Hebert (2007). Hospitals in Hurricane Katrina: challenges facing custodial institutions in a disaster. *Journal of Health Care for the Poor and Underserved 18*(2), 283–298.

Gudishala, R. and C. Wilmot (2012). Comparison of time-dependent sequential logit and nested logit for modeling hurricane evacuation demand. *Transportation Research Record: Journal of the Transportation Research Board* (2312), 134–140.

Han, L. D., F. Yuan, and T. Urbanik (2007). What is an effective evacuation operation? *Journal of urban planning and development 133*(1), 3–8.

Hasan, S., S. Ukkusuri, H. Gladwin, and P. Murray-Tuite (2011). Behavioral model to understand household-level hurricane evacuation decision making. *Journal of Transportation Engineering 137*(5), 341–348.

HIFLD (2021). Homeland Infrastructure Foundation-Level Data. https://hifld-geoplatform.opendata.arcgis.com/datasets/hospitals. Accessed: 2021-02-27.

Hing, E. and F. A. Bhuiya (2012). *Wait Time for Treatment in Hospital Emergency Departments, 2009.* US Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Health Statistics.

Hobeika, A. G. and C. Kim (1998, May). Comparison of traffic assignments in evacuation modeling. *IEEE Transactions on Engineering Management 45*(2), 192–198.

Iserson, K. V. and J. C. Moskop (2007). Triage in medicine, part I: concept, history, and types. *Annals of emergency medicine 49*(3), 275–281.

Janson, B. N. (1991). Dynamic traffic assignment for urban road networks. *Transportation Research Part B: Methodological 25*(2-3), 143–161.

Kaisar, E. I., L. Hess, and A. B. Portal Palomo (2012). An emergency evacuation planning model for special needs populations using public transit systems. *Journal of Public Transportation 15*(2), 3.

Karabuk, S. and H. Manzour (2019). A multi-stage stochastic program for evacuation management under tornado track uncertainty. *Transportation research part E: logistics and transportation review 124*, 128–151.

Kelen, G. D., C. K. Kraus, M. L. McCarthy, E. Bass, E. B. Hsu, G. Li, J. J. Scheulen, J. B. Shahan, J. D. Brill, and G. B. Green (2006). Inpatient disposition classification for the creation of hospital surge capacity: a multiphase study. *The Lancet 368*(9551), 1984–1990.

Kennedy, K., R. V. Aghababian, L. Gans, and C. P. Lewis (1996). Triage: techniques and applications in decision-making. *Annals of Emergency Medicine 28*(2), 136–144.

Kim, S., S. Shekhar, and M. Min (2008). Contraflow transportation network reconfiguration for evacuation route planning. *IEEE Transactions on Knowledge and Data Engineering 20*(8), 1115–1129.

Klein, D. and E. Hannan (1982). An algorithm for the multiple objective integer linear programming problem. *European Journal of Operational Research 9*(4), 378 – 385.

Kowaleski, A. M. and J. L. Evans (2020). Use of multiensemble track clustering to inform medium-range tropical cyclone forecasts. *Weather and Forecasting 35*(4), 1407–1426.

Landsea, C. (2015). Atlantic oceanographic & meteorological laboratory. http://www.aoml.noaa.gov/hrd/tcfaq/E11.html. Accessed: 2017-01-06.

Lerner, E. B., R. B. Schwartz, P. L. Coule, E. S. Weinstein, D. C. Cone, R. C. Hunt, S. M. Sasser, J. M. Liu, N. G. Nudell, I. S. Wedmore, et al. (2008). Mass casualty triage: an evaluation of the data and development of a proposed national guideline. *Disaster medicine and public health preparedness 2*(S1), S25–S34.

Li, A. C., L. Nozick, N. Xu, and R. Davidson (2012). Shelter location and transportation planning under hurricane conditions. *Transportation Research Part E: Logistics and Transportation Review 48*(4), 715–729.

McGinty, M. D., T. A. Burke, B. Resnick, D. J. Barnett, K. C. Smith, and L. Rutkow (2017). Decision processes and determinants of hospital evacuation and shelter-in-place during Hurricane Sandy. *Journal of Public Health Management and Practice 23*(1), 29–36.

Mesa-Arango, R., S. Hasan, S. V. Ukkusuri, and P. Murray-Tuite (2013). Household-level model for hurricane evacuation destination type choice using Hurricane Ivan data. *Natural Hazards Review 14*(1), 11–20.

Murray-Tuite, P. and B. Wolshon (2013). Evacuation transportation modeling: An overview of research, development, and practice. *Transportation Research Part C: Emerging Technologies 27*, 25 – 45.

Nasrollahi, N., A. AghaKouchak, J. Li, X. Gao, K. Hsu, and S. Sorooshian (2012). Assessing the impacts of different WRF precipitation physics in hurricane simulations. *Weather and Forecasting 27*(4), 1003–1016.

Ng, M., J. Park, and S. T. Waller (2010). A hybrid bilevel model for the optimal shelter assignment in emergency evacuations. *Computer-Aided Civil and Infrastructure Engineering 25*(8), 547–556.

Ng, M. and S. T. Waller (2010). Reliable evacuation planning via demand inflation and supply deflation. *Transportation Research Part E: Logistics and Transportation Review 46*(6), 1086–1094.

OpenDataDC (2021). Definitive Healthcare: USA Hospital Beds. https://opendata.dc.gov/datasets/1044bb19da8d4dbfb6a96eb1b4ebf629_0. Accessed: 2021-02-27.

OpenFEMA (2020). Openfema Data Sets. https://www.fema.gov/about/openfema/data-sets. Accessed: 2021-02-25.

Rambha, T., E. Jafari, and S. D. Boyles (2019). Transportation network issues in evacuations. In K. K. Stephens (Ed.), *New Media in Times of Crisis*, Chapter 7, pp. 144–161. Routledge.

Ricci, K. A., A. R. Griffin, K. C. Heslin, D. Kranke, and A. Dobalian (2015). Evacuate or shelter-in-place? the role of corporate memory and political environment in hospital-evacuation decision making. *Prehospital and disaster medicine 30*(3), 233–238.

Rockafellar, R. T. and R. J.-B. Wets (1991). Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research 16*(1), 119–147.

Sadri, A. M., S. V. Ukkusuri, P. Murray-Tuite, and H. Gladwin (2013). How to evacuate: model for understanding the routing strategies during hurricane evacuation. *Journal of transportation engineering 140*(1), 61–69.

Sbayti, H. and H. Mahmassani (2006). Optimal scheduling of evacuation operations. *Transportation Research Record: Journal of the Transportation Research Board 1964*, 238–246.

Sexton, K. H., L. M. Alperin, and J. D. Stobo (2007). Lessons from Hurricane Rita: The University of Texas Medical Branch Hospital's evacuation. *Academic medicine 82*(8), 792–796.

Sheffi, Y., H. Mahmassani, and W. B. Powell (1982). A transportation network evacuation model. *Transportation Research Part A: General 16*(3), 209 – 218.

Sherali, H. D., T. B. Carter, and A. G. Hobeika (1991). A location-allocation model and algorithm for evacuation planning under hurricane/flood conditions. *Transportation Research Part B: Methodological 25*(6), 439–452.

Sheu, J.-B. (2007). An emergency logistics distribution approach for quick response to urgent relief demand in disasters. *Transportation Research Part E: Logistics and Transportation Review 43*(6), 687–709.

Sheu, J.-B. (2010). Dynamic relief-demand management for emergency logistics operations under large-scale disasters. *Transportation Research Part E: Logistics and Transportation Review 46*(1), 1–17.

Sheu, J.-B. and C. Pan (2014). A method for designing centralized emergency supply network to respond to large-scale natural disasters. *Transportation research part B: methodological 67*, 284–305.

Skamarock, W. C., J. B. Klemp, J. Dudhia, D. O. Gill, D. M. Barker, W. Wang, and J. G. Powers (2005). A description of the advanced research WRF version 2. Technical report, DTIC Document.

Soika, K. (2006). Evacuation planning in Texas before and after Hurricane Rita. *House Research Organization, Texas House of Representatives, Interim News 79*(2), 1–8.

Sylva, J. and A. Crema (2004). A method for finding the set of non-dominated vectors for multiple objective integer linear programs. *European Journal of Operational Research 158*(1), 46 – 55.

Taaffe, K., M. Johnson, and D. Steinmann (2006). Improving hospital evacuation planning using simulation. In *Proceedings of the 38th conference on Winter simulation*, pp. 509–515. Winter Simulation Conference.

Tayfur, E. and K. Taaffe (2009a). A model for allocating resources during hospital evacuations. *Computers & Industrial Engineering 57*(4), 1313 – 1323.

Tayfur, E. and K. Taaffe (2009b). Simulating hospital evacuation—the influence of traffic and evacuation time windows. *Journal of Simulation 3*(4), 220–234.

Tuydes, H. and A. Ziliaskopoulos (2006). Tabu-based heuristic approach for optimization of network evacuation contraflow. *Transportation Research Record: Journal of the Transportation Research Board* (1964), 157–168.

USDHHS (2015). Topic collection: Healthcare facility evacuation sheltering. https://asprtracie.hhs.gov/technical-resources/57/Healthcare-Facility-Evacuation-Sheltering/56. Accessed: 2016-12-30.

Villarini, G., B. Luitel, G. A. Vecchi, and J. Ghosh (2019). Multi-model ensemble forecasting of North Atlantic tropical cyclone activity. *Climate Dynamics 53*(12), 7461–7477.

Wang, J., Y. Hong, L. Li, J. J. Gourley, S. I. Khan, K. K. Yilmaz, R. F. Adler, F. S. Policelli, S. Habib, D. Irwn, et al. (2011). The coupled routing and excess storage (CREST) distributed hydrological model. *Hydrological sciences journal 56*(1), 84–98.

Wang, Y. and J. Wang (2019). Integrated reconfiguration of both supply and demand for evacuation planning. *Transportation research part E: logistics and transportation review 130*, 82–94.

Wanik, D., E. Anagnostou, M. Astitha, B. Hartman, G. Lackmann, J. Yang, D. Cerrai, J. He, and M. Frediani (2018). A case study on power outage impacts from future Hurricane Sandy scenarios. *Journal of Applied Meteorology and Climatology 57*(1), 51–79.

Watson, J.-P. and D. L. Woodruff (2011). Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems. *Computational Management Science 8*(4), 355–370.

Westerink, J. J., R. A. Luettich, J. C. Feyen, J. H. Atkinson, C. Dawson, H. J. Roberts, M. D. Powell, J. P. Dunion, E. J. Kubatko, and H. Pourtaheri (2008). A basin-to channel-scale unstructured grid hurricane storm surge model applied to southern Louisiana. *Monthly Weather Review 136*(3), 833–864.

WMA (2006). Wma statement on medical ethics in the event of disasters. http://www.wma.net/en/30publications/10policies/d7/. Accessed: 2017-01-25.

Wolshon, B., E. Urbina, C. Wilmot, and M. Levitan (2005). Review of policies and practices for hurricane evacuation. I: Transportation planning, preparedness, and response. *Natural hazards review 6*(3), 129–142.

Wolshon, B., E. Urbina Hamilton, M. Levitan, and C. Wilmot (2005). Review of policies and practices for hurricane evacuation. II: Traffic operations, management, and control. *Natural Hazards Review 6*(3), 143–161.

Xie, C., D.-Y. Lin, and S. T. Waller (2010). A dynamic evacuation network optimization problem with lane reversal and crossing elimination strategies. *Transportation research Part E: logistics and transportation review 46*(3), 295–316.

Xie, C. and M. Turnquist (2009). Integrated evacuation network optimization and emergency vehicle assignment. *Transportation Research Record: Journal of the Transportation Research Board* (2091), 79–90.

Xiongfei, Z., S. Qixin, H. Rachel, and R. Bin (2010). Network emergency evacuation modeling: A literature review. In *Optoelectronics and Image Processing (ICOIP), 2010 International Conference on*, Volume 2, pp. 30–34. IEEE.

Xu, K., R. A. Davidson, L. K. Nozick, T. Wachtendorf, and S. E. DeYoung (2016). Hurricane evacuation demand models with a focus on use for prediction in future events. *Transportation Research Part A: Policy and Practice 87*, 90–101.

Yang, K., R. A. Davidson, L. K. Nozick, B. Blanton, and B. Colle (2017). Scenario-based hazard trees for depicting resolution of hurricane uncertainty over time. *Natural Hazards Review*.

Yi, W., L. Nozick, R. Davidson, B. Blanton, and B. Colle (2017). Optimization of the issuance of evacuation orders under evolving hurricane conditions. *Transportation Research Part B: Methodological 95*, 285 – 304.

Zachria, A. and B. Patel (2006). Deaths related to Hurricane Rita and mass evacuation. *Chest Journal 130*(4, Supplement), 124S–c–1 – 124S–c–2.

Zhang, Z., K. Spansel, and B. Wolshon (2014). Effect of phased evacuations in megaregion highway networks. *Transportation Research Record: Journal of the Transportation Research Board 2459*, 101–109.

Zhou, Y., J. Liu, Y. Zhang, and X. Gan (2017). A multi-objective evolutionary algorithm for multi-period dynamic emergency resource scheduling problems. *Transportation Research Part E: Logistics and Transportation Review 99*, 77–95.