Optimizing Diesel Fuel Supply Chain Operations for Hurricane Relief

Daniel Duque[a], Haoxiang Yang[b] and David P. Morton[a]

[a]Department of Industrial Engineering & Management Sciences, Northwestern University, Evanston, IL, USA; [b]Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM, USA

ARTICLE HISTORY
Compiled May 10, 2020

ABSTRACT
Hurricanes can cause severe property damage and casualties in coastal regions. Diesel fuel plays a crucial role in hurricane disaster relief. It is important to optimize fuel supply chain operations so that emergency demand for diesel can be mitigated in a timely manner. However, it can be challenging to estimate demand for fuel and make informed proactive and reactive decisions in the distribution process, accounting for the hurricane's path and severity. We develop predictive and prescriptive models to guide diesel fuel supply chain operations for hurricane disaster relief. We construct a model for estimating diesel fuel demand from historical weather forecasts and power outage data. This predictive model feeds into a prescriptive stochastic programming model implemented in a rolling-horizon fashion to dispatch tank trucks. This data-driven optimization tool provides a framework for decision support in preparation for approaching hurricanes, and our numerical results provide insights regarding key aspects of operations.

KEYWORDS
Diesel fuel supply chain; hurricane relief; stochastic programming

1. Introduction

Tropical storms and hurricanes formed in the Atlantic Ocean have been a major contributor to natural disasters on the east and gulf coasts of the United States and in Caribbean countries and territories, causing significant property damage and casualties. The Atlantic hurricane season is from early June through late November. There have been 43 major hurricanes since 2005, the year in which the costliest hurricane in history, Katrina, made landfall in Louisiana, and altogether they have caused more than 600 billion dollars of damage in the United States and the Caribbean. High winds, storm surge, and inland flooding from rainfall cause damage to infrastructure systems including electric power, transportation, and the liquid-fuel supply chain. For example, flooding from Hurricane Katrina led to breaches in the city's levee system and flooded all major roads in and out of New Orleans, Louisiana. Hurricane Sandy in 2012 caused power loss for more than 8 million customers. The hurricane season of 2017 is the costliest on record and included the effects of Hurricane Harvey hitting Houston, Texas and disrupting US petroleum refining capacity.

CONTACT D. P. Morton. Email: david.morton@northwestern.edu

We focus on the diesel fuel supply chain (DFSC) because diesel fuel is a critical resource for hurricane relief. Disruptions caused by natural disasters to a DFSC can have a detrimental impact on other aspects of disaster relief: transportation networks rely on the DFSC to provide fuel for vehicles used for evacuation and repair; and, critical infrastructure, such as hospitals, municipal water systems, and facilities for first responders, requires diesel fuel for power generation when electric power is lost.

In this paper, we propose a stochastic programming model of DFSC operations that is used within a rolling-horizon approach. We present a case study of DFSC operations in Florida for Hurricane Irma. We analyze weather data from the National Oceanic and Atmospheric Administration (NOAA) and other sources, and we model how hurricane winds affect the demand for diesel fuel. Using this relationship, we map a set of hurricane scenarios from the Global Ensemble Forecast System (GEFS) to scenarios for diesel fuel demand with a county-level resolution, which is used as a key input for the stochastic program. While there is a stream of literature on DFSC operations under normal conditions and disaster operations management of relief supplies that we discuss below, to our knowledge this is the first data-driven approach that specifically focuses on addressing DFSC operations for hurricane relief.

We first review relevant literature on hurricane relief and DFSC operations in Section 2. In Section 3, we describe in detail the two-stage stochastic program and the rolling-horizon approach we employ. Section 4 first develops our model for predicting demand for diesel fuel and then analyzes hypothetical DFSC operations during Hurricane Irma by presenting computational results and insights from our models, including comparisons with baseline alternatives. We conclude with remarks on potential extensions of our approach in Section 5.

2. Literature Review

Many previous studies have analyzed different aspects of hurricanes. The National Hurricane Center (NHC) uses multiple forecasting models, such as ECMWF [32], GFS [33], and UKMET [56], to prepare tropical cyclone advisories, which inform users regarding a hurricane's intensity and track. See the National Hurricane Center [52] for a list of the forecasting models that NHC uses. Although useful in providing guidance for emergency planning, these forecasting models do not fully characterize a hurricane's physical impact on infrastructure systems such as fuel supply chains, transportation networks, and electric power grids. Therefore, in addition to weather forecasts, hazard models for infrastructure damage are needed to support decision making for evacuation and recovery. In a hurricane, damage can result from high wind speeds, storm surge in coastal areas, and inland flooding [e.g., 19, 51, 61, 67]. Recent efforts [21, 26] combine output from multiple models of weather forecasts (WRF [64], CREST [68], and ADCIRC [48]) and provide a dynamic method to model the impact of wind speed and storm surge on evacuation.

Most of the previous literature on hurricane relief focuses on evacuation, rescue missions, and distribution of emergency supplies [e.g., 12, 69], instead of DFSCs. A DFSC is part of a larger supply chain for liquid fuels, which incorporates other petroleum products such as gasoline, heavy fuel oil, and jet fuel. A typical fuel supply chain starts from crude oil production sites. After extraction, and transportation to refineries by pipeline or ship, crude oil is processed to different petroleum products [14]. In the United States, the bulk of processed petroleum products are distributed through pipelines from refineries to terminals, while the rest is transported to terminals by

truck, vessel, or train. Terminals serve as large storage venues and satisfy the demand for fuel stations at which most end customers replenish diesel fuel. See Neiro and Pinto [53] and Sear [63] for detailed descriptions of the petroleum supply chain.

A spike in demand for diesel fuel is often observed during and after a hurricane. In the immediate aftermath, the electric power system may incur outages [e.g., 31]. While the power system is down, diesel can fuel backup power generation at crucial locations [7, 10]. Electricity is required to power: hospitals [6, 42, 47, 58]; chemical plants and temperature-sensitive laboratories, e.g., in universities [18]; communications facilities for evacuation coordination [40, 41]; and municipal water systems to treat and distribute drinking water [1, 2]. To pump fuel, service stations require electricity, and so the lack of diesel-fueled backup power can constrain transportation both for evacuation and distribution of disaster relief supplies. Some states, including Florida, have laws that require service stations at important locations, e.g., on an evacuation route, to have access to a backup generator [5]. After a hurricane, fuel terminals require electricity so that additives such as ethanol can be mixed with gasoline and other petroleum products, and so the lack of diesel fuel and generators can lead to cascading failures across multiple systems of infrastructure [8]. These demands for diesel can be coupled with compromised supplies, as refineries and ports for imported fuel may be shut down during a hurricane and take time to recover. Hurricane Harvey reduced overall US refining capacity by as much as 25-30% [25]. During Hurricane Irma, all ports in Florida were closed for up to five days, stopping the import of nearly all fuel, including diesel [8].

There is a rich history of applying mathematical models to optimize supply chain operations in the petroleum industry [e.g., 16, 23, 24]. Sear [63] analyzes different aspects of a fuel supply chain such as product types and the scope of demands and costs, and describes a linear programming framework for optimizing logistics. Neiro and Pinto [53] identify three key segments of the supply chain: processing units at refineries; petroleum and product tanks; and pipelines, each with its own sets of constraints, and they link them to form a large-scale deterministic fuel supply chain model. Escudero et al. [30] propose a two-stage stochastic linear programming model named CORO to incorporate uncertainty in demand, selling price, and supply cost. Dempster et al. [29] extend the CORO model to a multi-stage stochastic program. See Lima et al. [44] for a review of different models and their characteristics for fuel supply chain planning.

While a significant stream of literature addresses operations outside of disasters, there is little work directly addressing the DFSC for hurricane relief. Adhitya et al. [11] propose a heuristic to reschedule operations in the fuel supply chain after observing a disruption, such as delay of supply, resource unavailability, and changes in demand. Using a two-stage stochastic program, Beheshtian et al. [20] aim to design a resilient motor-fuel supply chain in preparation for a hurricane, which will induce uncertain damage. First-stage decisions enhance resilience of the supply network, and second-stage decisions replenish fuel in the residual network, accounting for damage. Suzuki [65] models constraints on the transportation problem of delivering relief resources accounting for limited fuel supply. Li et al. [43] develop a mixed-integer program for effective and equitable gasoline distribution after a natural disaster. They model surges in demand for gasoline and identify important service stations requiring backup generators, but their model does not quantify demand for diesel at those locations. To the best of our knowledge, there is no specific literature that models the DFSC for disaster relief and captures dependencies between the DFSC and other relevant infrastructure including transportation and electric power.

There is a broader literature on operations management for disasters. Altay and

Green [15] organize the literature by the phases of a disaster: mitigation, preparedness, response, and recovery. We focus on the latter three phases, with an eye towards proactive, dynamic approaches to the preparedness and response for the DFSC. At the intersection of preparedness and response, Rawls and Turquist [59] develop a two-stage stochastic program for prepositioning emergency supplies on a network, in which first-stage variables determine facility locations and inventories, and second-stage variables distribute resources, post-disaster. Uncertainty is modeled by scenarios for demand and network damage, with the goal of minimizing the expected cost of facility location, inventory, and distribution. Subsequent work advances this type of two-stage prepositioning model [13, 28, 54, 60], and further work examines slow on-set disasters [e.g., 49].

In our setting, less attention has been paid to models in which decisions are taken dynamically in short time frames. On that front, Pacheco and Batta [57] develop a dynamic model for prepositioning relief goods in preparation for a hurricane, wherein forecasts are updated every six hours, and decisions are revised in a rolling-horizon fashion as the hurricane progresses. Earlier work by Lodree and Taskin [45, 46] is similar in spirit in that Bayesian updates are made sequentially based on new hurricane forecasts to decide when and how much to order, for resources like flashlights, batteries, and generators. Once an order is placed, it is assumed to cover demand to the horizon.

The literature reviewed above largely focuses on the operation of one system, e.g., a power grid or transportation system. For our work, other systems and facilities rely on the DFSC. Therefore, also relevant to our work is literature that studies recovery from a disaster accounting for the interdependence of infrastructure systems. González et al. [34], for example, consider the design of an interdependent network problem to obtain a minimum-cost recovery strategy of a partially destroyed system of infrastructure networks, e.g., electric power, water distribution, and natural gas networks. Accounting for interdependence of the networks and geographical co-location, a mixed-integer programming model is proposed to determine reconstruction decisions of a network to minimize the total cost that includes reconstruction costs, disconnection costs, and flow costs.

Successfully gathering and processing real-time data for a proactive plan can be challenging for many of the approaches sketched above. Information about the status of infrastructure is often unavailable or incomplete, and there is limited research giving guidelines on how to approach these challenges. Yagci Sokat et al. [70] introduce a framework to evaluate real-time data in the context of humanitarian logistics, and Yagci Sokat et al. [71] propose a framework to impute missing data regarding road conditions and present its application in the 2010 earthquake in Haiti. Recent literature [39] also aims to improve understanding of critical events by incorporating multiple data sources including social media, which can be accessed and processed in real time.

3. Modeling DFSC Operations for Hurricane Relief

Operations in a DFSC are driven by the supply and demand for diesel fuel. We assume a known supply, but a random demand, depending on the hurricane's path and intensity. To mitigate electric power outages, backup generators fueled by diesel restore power at critical locations. We map weather scenarios to power outages, which result in additional demand for diesel fuel, referred to as *surge demand* henceforth. As a hurricane develops, weather forecasts are periodically updated and released. We use these forecasts to estimate surge demand for diesel fuel. The dynamic nature of

the forecasts naturally yields a rolling-horizon approach that we detail in this section. Within this approach, we solve a two-stage stochastic program, which we describe in Section 3.1, in order to dispatch tank trunks to deliver diesel. Details on how we parameterize and employ the rolling-horizon approach are presented in Section 3.2. In Section 3.3 we close by discussing two deterministic alternatives to the two-stage model. Such alternatives are useful for two reasons: (i) a deterministic optimization model is relevant when the forecasts are only available in the form of point forecasts, as opposed to a collection of scenarios; and, (ii) a deterministic optimization model assuming perfect information can serve as an optimistic lower bound when minimizing shortfalls.

3.1. Two-stage Stochastic Program

Our two-stage stochastic program for DFSC operations, models supply at fuel-terminal nodes and demand at destination nodes, where we assume a known hourly *nominal* demand. In our implementation, the latter nodes aggregate demand at the level of a county, while supply at terminals is associated with a small number of specific counties that produce or receive exogenous diesel fuel. Fuel is distributed via tank trucks, which transport diesel from supply nodes to demand nodes, although other transportation modes such as rail or pipelines could also be incorporated. We assume that the number of tank trucks is fixed and that the schedule of supply arriving at the terminals is known.

We model the flow of fuel throughout the network, as well as the flow of loaded and empty trucks, to characterize the system's ability to transport diesel. When the hurricane arrives, surge demand for diesel is driven by outages in the electric power system. Because the impact of the hurricane is uncertain, the *surge* demand is random. We characterize this surge with scenarios that are built using wind forecasts and historical data of outages, as detailed in Section 4. Due to the demand surge, and potential resulting supply deficit, significant diesel fuel shortages may occur. During the span of the planning period, a small percentage of nominal demand has the highest priority (e.g., for trucks to transport diesel), but otherwise surge demand has priority.

We consider a two-stage stochastic program that models the inventory dynamics of diesel fuel as well as the routing of tank trucks over time on a network, again with the geographic resolution of counties. Counties with supply are denoted by set $\mathcal{I}_s$, counties with demand by set $\mathcal{I}_d$, and we create two nodes for each county that has both supply and demand. The nodes of the entire network are represented by all counties, with supply node duplicates where appropriate, and is denoted by $\mathcal{I} = \mathcal{I}_s \cup \mathcal{I}_d$.

Time is discretized into periods of one hour length with all relevant times being integer values. In the two-stage model, the first stage comprises time periods $\mathcal{T}_f$ and the second stage comprises time periods $\mathcal{T}_s$. Naturally, the first time period in $\mathcal{T}_s$ is one unit of time ahead of the last time period in $\mathcal{T}_f$. While the entire time horizon of interest is specified by $\{1, 2, \ldots, T-1, T\}$, with an eye to the rolling-horizon approach that we describe in Section 3.2, we let t_f be the first time period in $\mathcal{T}_f$, and similarly, let t_s be the first time period in $\mathcal{T}_s$. Hence, $\mathcal{T}_f = \{t_f, t_f + 1, t_f + 2, \ldots, t_s - 1\}$ and $\mathcal{T}_s = \{t_s, t_s + 1, t_s + 2, \ldots, H\}$, where time $t_f \geq 1$ to time $H \leq T$ is the planning period for an instance of the model we specify here, and where we define $\mathcal{T} = \mathcal{T}_f \cup \mathcal{T}_s$.

For a given set of time periods, the model describing the inventory dynamics and in-transit tank trucks is represented recursively as a function of past inventory states and routing decisions. For a specific time period t, regardless of whether $t \in \mathcal{T}_f$ or

$t \in \mathcal{T}_s$, state variables describing the system are shown in Table 1.

$I^d_{i,t}$	inventory of diesel fuel at demand node $i \in \mathcal{I}_d$ at time t (barrels)
$I^s_{i,t}$	inventory of diesel fuel at supply node $i \in \mathcal{I}_s$ at time t (barrels)
$r_{i,t,\ell}$	number of loaded tank trucks scheduled to arrive at demand node $i \in \mathcal{I}_d$ at time $t + \ell$
$g_{i,t,\ell}$	number of empty tank trucks scheduled to arrive at supply node $i \in \mathcal{I}_s$ at time $t + \ell$

Table 1. State variables for the diesel fuel supply chain model

Table 2 summarizes additional decision variables for controls and recourse actions. The flow of both loaded trucks (x-variables) and empty trucks (y-variables) is specified via four indices: the origin node, i, the time period of departure, t, the destination node, j, and the arrival time period, t'. Three types of such movements are allowed. A loaded tank truck can deliver from $i \in \mathcal{I}_s$ to $j \in \mathcal{I}_d$, i.e., from a supply node to a demand node, and only full truck-loads can be delivered. An empty truck must immediately return from $i \in \mathcal{I}_d$ to some $j \in \mathcal{I}_s$, which need not be its original supply node. Finally, an empty tank truck can relocate from $i \in \mathcal{I}_s$ to $j \in \mathcal{I}_s$. Travel time is assumed to be deterministic, and we define sets $\mathcal{A}^x_{i,t}$ and $\mathcal{A}^y_{i,t}$ that contain all destination, arrival-time pairs, (j, t'), accessible from node i with start time t, for loaded and empty trucks respectively. With $j = i$ we allow an empty truck to be idle at a supply node for one period of time, i.e., $(i, t+1) \in \mathcal{A}^y_{i,t}$.

$x_{i,t,j,t'}$	number of loaded trucks that depart from node $i \in \mathcal{I}_s$ at time $t \in \mathcal{T}$ and arrive at node $j \in \mathcal{I}_d$ at time $t' \in \mathcal{T}$
$y_{i,t,j,t'}$	number of empty trucks that depart from node $i \in \mathcal{I}$ at time $t \in \mathcal{T}$ and arrive at node $j \in \mathcal{I}_s$ at time $t' \in \mathcal{T}$
$z^n_{i,t}$	nominal shortage at node $i \in \mathcal{I}_d$ at time $t \in \mathcal{T}$ (barrels)
$z^e_{i,t}$	surge shortage at node $i \in \mathcal{I}_d$ at time $t \in \mathcal{T}$ (barrels)

Table 2. Additional decision variables for the diesel fuel supply model

Table 3 summarizes the parameters of the model including both deterministic parameters (capacity, supply, nominal demand, and demand point forecast) and random variables (surge demand).

b	capacity per truck (barrels)
$s_{i,t}$	additional supply of diesel fuel at node $i \in \mathcal{I}_s$ at time $t \in \mathcal{T}$
$d^n_{i,t}$	nominal demand (in barrels) for diesel fuel at node $i \in \mathcal{I}_d$ at time $t \in \mathcal{T}$ (barrels)
$\hat{d}^e_{i,t}$	surge demand point forecast for diesel fuel at node $i \in \mathcal{I}_d$ at time $t \in \mathcal{T}$ (barrels)
$\tilde{d}^e_{i,t}$	random surge demand for diesel fuel at node $i \in \mathcal{I}_d$ at time $t \in \mathcal{T}$ (barrels)
$\delta_{\max}$	upper bound on the transit time for a tank truck between two nodes

Table 3. Parameters for the diesel fuel supply model

The mathematical formulation of our two-stage stochastic program is shown in model (1). The cost functions of surge shortage and nominal shortage at time period t are denoted by $c^e_t(z^e_{\cdot,t})$ and $c^n_t(z^n_{\cdot,t})$, respectively. The "$\cdot$" notation in the subscript represents the collection of z variables with the same t but different i indices. We defer the explanation of how we parameterize these cost functions to Section 4.3. Again looking ahead to the rolling-horizon approach, the model is parameterized by the input that specifies the values of state variables due to previously taken decisions. To

distinguish current decision variables and fixed values from previous time periods, we modify the latter symbols by adding the modifier "$\widehat{\ }$" and include the appropriate time index in the formulation. When subscript indices are omitted the term corresponds to the appropriately dimensioned vector; e.g., $\hat{I}^d$ as a vector of fixed inventories for $\hat{I}^d_{i,t}$, $i \in \mathcal{I}_d$, with relevant time periods specified in the formulation.

$$Q_{\mathcal{T}_f}(\hat{I}^d, \hat{I}^s, \hat{r}, \hat{g}, \hat{d}^e) =$$

$$\min \quad \sum_{t \in \mathcal{T}_f} \left[c^e_t(z^e_{\cdot,t}) + c^n_t(z^n_{\cdot,t}) \right] + \mathbb{E}[Q_{\mathcal{T}_s}(I^d, I^s, r, g, \tilde{d}^e)] \tag{1a}$$

$$\text{s.t.} \quad I^d_{i,t} = I^d_{i,t-1} + b\, r_{i,t-1,1} - d^n_{i,t} - \hat{d}^e_{i,t} + z^n_{i,t} + z^e_{i,t} \qquad \forall i \in \mathcal{I}_d, t \in \mathcal{T}_f \tag{1b}$$

$$I^s_{i,t} = I^s_{i,t-1} + s_{i,t} - b \sum_{(j,t') \in \mathcal{A}^x_{i,t}} x_{i,t,j,t'} \qquad \forall i \in \mathcal{I}_s, t \in \mathcal{T}_f \tag{1c}$$

$$r_{i,t-1,1} = \sum_{(j,t') \in \mathcal{A}^y_{i,t}} y_{i,t,j,t'} \qquad \forall i \in \mathcal{I}_d, t \in \mathcal{T}_f \tag{1d}$$

$$\sum_{(j,t') \in \mathcal{A}^x_{i,t}} x_{i,t,j,t'} + \sum_{(j,t') \in \mathcal{A}^y_{i,t}} y_{i,t,j,t'} = g_{i,t-1,1} \qquad \forall i \in \mathcal{I}_s, t \in \mathcal{T}_f \tag{1e}$$

$$r_{j,t,\ell} = r_{j,t-1,\ell+1} + \sum_{i \in \mathcal{I}_s : (j,t+\ell) \in \mathcal{A}^x_{i,t}} x_{i,t,j,t+\ell} \qquad \forall j \in \mathcal{I}_d, t \in \mathcal{T}_f, \ell = 1, 2, \ldots, \delta_{\max} \tag{1f}$$

$$g_{j,t,\ell} = g_{j,t-1,\ell+1} + \sum_{i \in \mathcal{I}_d : (j,t+\ell) \in \mathcal{A}^y_{i,t}} y_{i,t,j,t+\ell} \qquad \forall j \in \mathcal{I}_s, t \in \mathcal{T}_f, \ell = 1, 2, \ldots, \delta_{\max} \tag{1g}$$

$$z^n_{i,t} \le d^n_{i,t} \qquad \forall i \in \mathcal{I}_d, t \in \mathcal{T}_f \tag{1h}$$

$$z^e_{i,t} \le \hat{d}^e_{i,t} \qquad \forall i \in \mathcal{I}_d, t \in \mathcal{T}_f \tag{1i}$$

$$I^d_{i,t_f-1} = \hat{I}^d_{i,t_f-1} \qquad \forall i \in \mathcal{I}_d \tag{1j}$$

$$I^s_{i,t_f-1} = \hat{I}^s_{i,t_f-1} \qquad \forall i \in \mathcal{I}_s \tag{1k}$$

$$r_{i,t_f-1,1} = \hat{r}_{i,t_f-1,1} \qquad \forall i \in \mathcal{I}_d \tag{1l}$$

$$g_{i,t_f-1,1} = \hat{g}_{i,t_f-1,1} \qquad \forall i \in \mathcal{I}_s \tag{1m}$$

$$I^d_{i,t} \ge 0 \quad \forall i \in \mathcal{I}_d, t \in \mathcal{T}_f \tag{1n}$$

$$I^s_{i,t} \ge 0 \quad \forall i \in \mathcal{I}_s, t \in \mathcal{T}_f \tag{1o}$$

$$r_{i,t,l} \ge 0 \qquad \forall i \in \mathcal{I}_d, t \in \mathcal{T}_f, l = 1, 2, \ldots, \delta_{\max} \tag{1p}$$

$$g_{i,t,l} \ge 0 \qquad \forall i \in \mathcal{I}_s, t \in \mathcal{T}_f, l = 1, 2, \ldots, \delta_{\max} \tag{1q}$$

$$x_{i,t,j,t'} \ge 0, \quad \forall i \in \mathcal{I}_s, t \in \mathcal{T}_f, (j,t') \in \mathcal{A}^x_{i,t} \tag{1r}$$

$$y_{i,t,j,t'} \ge 0, \quad \forall i \in \mathcal{I}_d \cup \mathcal{I}_s, t \in \mathcal{T}_f, (j,t') \in \mathcal{A}^y_{i,t} \tag{1s}$$

$$z^n_{i,t}, z^e_{i,t} \ge 0 \quad \forall i \in \mathcal{I}_d, t \in \mathcal{T}_f. \tag{1t}$$

Here, $Q_{\mathcal{T}_s}(\cdot)$ has identical form as $Q_{\mathcal{T}_f}(\cdot)$ except that: (i) $\mathcal{T}_s$ replaces $\mathcal{T}_f$ throughout model (1); (ii) t_s replaces t_f in constraints (1j)-(1m); (iii) the expectation term in (1a) is eliminated; and (iv) as indicated in $Q_{\mathcal{T}_s}(\cdot)$'s argument, realizations of $\tilde{d}^e_{it}$ replace parameter $\hat{d}^e_{i,t}$ in constraints (1b) and (1i).

We minimize the sum of penalty costs from nominal and surge shortage and the expected value of the second stage cost, as shown in the objective function (1a). Constraints (1b) and (1c) model the inventory of fuel at demand and supply nodes, respectively. Constraint (1d) models the balance of trucks at demand nodes by requiring

that the number of loaded trucks arriving at a particular node equals the number of empty trucks leaving from that node at time t. Constraint (1e) follows a similar logic but for supply nodes, and allows repositioning of an empty tank truck from one supply node to another and idling at the same supply node. Constraints (1f)–(1g) update the state variables of loaded and empty trucks according to the number of trucks already in transit and the new departures at time t. Constraints (1h)-(1i) prohibit improper mixing of nominal and surge demands with their shortage variables. Constraints (1j)–(1m) specify boundary conditions, by initializing the state variables inherited from $\hat{I}^d, \hat{I}^s, \hat{r}, \hat{g}$. Constraints (1n)–(1t) require all variables to be non-negative.

3.2. Rolling-horizon Approach

We embed model (1) in a rolling-horizon approach—see Algorithm 1—over the entire time span of interest of an arriving hurricane and its aftermath, $\{1, 2, \ldots, T-1, T\}$. This set specifies both the hourly temporal resolution and *problem* horizon. Given this, to define a model instance (1) for the rolling-horizon procedure, we require an initial time period for decisions, t_f, the length of the first stage, F, and the *model* horizon, H. In this way, $\mathcal{T}_f = \{t_f, t_f+1, \ldots, t_f+F-1\}$ and $\mathcal{T}_s = \{t_f+F, t_f+F+1, \ldots, t_f+H-1\}$.

As a practical matter, deploying operational decisions will require advance notice, implying that if the model is proposing decisions that start at time t_f, it will use forecasts issued some time before t_f, say $t_f - N$. The last element in the rolling-horizon construct is the number of time periods to roll forward, which we denote $R \in \{1, 2, \ldots, F\}$. This means that the minimum length of time to roll forward is one hour and the maximum is the length of the first-stage set $\mathcal{T}_f$. The following summarizes the parameters required to run model (1) in our rolling-horizon approach:

- T: problem time horizon
- H: model time horizon
- F: length of the first stage
- R: length of time to roll forward
- N: length of time in advance that decisions should be made.

Rolling forward in time essentially means that control variables (x and y in our model) are fixed for all trucks departing at time $t \in \mathcal{T}_f$ such that $t_f \leq t \leq t_f + R - 1$. In addition, we assume demand realizations for these time periods are known upon fixing these variables. Of course, these realizations of surge demand will differ from our forecasts, and hence rolling forward requires updating the actual values of the state variables, in particular I^d. Having fixed the control variables, and having accounted for the realized surge demand, the new values of the state variables are computed because they serve as input to the next instance of model (1), which starts at time $t_f + R$. To formalize this procedure, let $\mathcal{T}_R = \{1, R+1, 2R+1, \ldots\}$ be the set of initial time periods at which the model is solved. Algorithm 1 summarizes the procedure described in this section.

After Algorithm 1 initializes state variables and time-indexed sets for the current model instance, in step 5 we obtain the point forecast of the surge demand for $t \in \mathcal{T}_f$ and scenarios for surge demand for $t \in \mathcal{T}_s$. We discuss this further in the next section. Solving model (1) in step 6 determines the dispatch plan for time periods $t \in \mathcal{T}_r$, a set defined in step 7. Re-solving the model in step 10 is simply a convenient way of computing the state variables at the last period in $\mathcal{T}_r$. Step 11 saves this for input to the next instance of model (1) with the next start time in $\mathcal{T}_R$. To understand the way

Algorithm 1 Rolling-horizon approach

Require: I_0^d, I_0^s, r_0, g_0, initial inventories of fuel and trucks; $d^n, \hat{d}^e, \tilde{d}^e$, nominal and surge demand.
Ensure: `void`

Let $\hat{I}^d = I_0^d$, $\hat{I}^s = I_0^s$, $\hat{r} = r_0$, $\hat{g} = g_0$
for $t_f \in \mathcal{T}_R$ **do**
  Let $\mathcal{T}_f = \{t_f, t_f + 1, \ldots, t_f + F - 1\}$
  Let $\mathcal{T}_s = \{t_f + F, t_f + F + 1, \ldots, t_f + H - 1\}$
  Retrieve point forecast $\hat{d}^e$ and scenarios $\tilde{d}^e$ issued at time $t_f - N$ for $\mathcal{T}_f$ and $\mathcal{T}_s$
  Solve model (1) with input $(\hat{I}^d, \hat{I}^s, \hat{r}, \hat{g}, \hat{d}^e)$ and the scenarios of $\tilde{d}^e$
  Let $\mathcal{T}_r = \{t_f, t_f + 1, \ldots, t_f + R - 1\}$ ▷ Start rolling forward
  Fixed control variable x and y for all tank truck departures in $t \in \mathcal{T}_r$
  Update surge demand, replacing $\hat{d}^e$ with actual realizations for all $t \in \mathcal{T}_r$
  Re-solve model (1) to compute state variables at time $t_o = t_f + R - 1$
  $\hat{I}^d = I_{t_o}^d$, $\hat{I}^s = I_{t_o}^s$, $\hat{r} = r_{t_o}$, $\hat{g} = g_{t_o}$ ▷ End rolling forward
end for
return `void`

in which we roll ahead and fix variables, it is helpful to consider the two extreme cases of $R = 1$ and $R = F$. The former only fixes the decisions made in the current initial time period, t_f, and the latter fixes those in all first-stage time periods, i.e., through period $t_f + F - 1$.

3.3. Deterministic Alternative Models

Model (1) is a relatively compact representation of a two-stage stochastic program that facilitates different models as special cases. For a model horizon, H, our choice of F determines adaptability to different sample paths. If $F = H$ then the first stage spans the entire planning horizon, and there is no second stage. Despite losing resolution in the demand realizations, this model is useful when the only available data come from a single point forecast, and therefore the expectation in (1a) is eliminated. Note also that a high-quality point forecast reduces the value of a two-stage model since the scenarios provide little additional information to compute a policy.

The same setup (with $F = H$) also serves to evaluate ideal cases, not for the purpose of deriving an operations policy, but rather to obtain an optimistic bound on what is possible. If instead of a point forecast we assume perfect information of future demand throughout the entire span of the hurricane (problem horizon T), such a model serves as a benchmark to assess both the two-stage model based on demand scenarios and the deterministic model based on point forecasts.

4. Case Study: Hurricane Irma

In this section we detail a case study, applying our approach to data available for Hurricane Irma, over the period from August 30, 2017 to September 12, 2017. Hurricane Irma caused severe damage in the Caribbean and the State of Florida, although we limit our focus to Florida. We first provide an overview of the DFSC in Florida and of Hurricane Irma. We detail the process of mapping the severity of the hurricane to

the demand for diesel fuel in Section 4.1 and of constructing hurricane scenarios from GEFS data in Section 4.2. In Section 4.3 we describe how the cost function in (1) prioritizes and penalizes shortfalls in satisfying demand for diesel fuel. With the procedure for estimating all parameters for the stochastic programming model of Section 3 in place, we present the computational results and summarize insights in Section 4.5.

Florida does not produce crude oil or have refining capability of significance. More than 95% of Florida's refined petroleum products are transported to the state through ships or pipelines, and there is rarely export of fuel from the state. During Hurricane Irma, because only Florida was hit by the hurricane, the refining capacity of other Gulf States (Texas, Mississippi, Louisiana, and Alabama) was unaffected, and the supply of diesel fuel quickly resumed after the hurricane. Thus, this case study allows us to isolate the effects of disruption to the *demand* for diesel relative to the nominal supply chain.

Florida's daily consumption of refined petroleum products is about 550 thousand barrels [3], of which diesel is about 27% (150 thousand barrels). We assume that the demand is approximately equal to the supply because only a small amount of diesel is typically stored in inventory. Among all the end-use sectors for diesel, on-highway use (motor vehicle consumption) makes up around 70% of the demand. Other major use sectors are: vessel bunkering (9%), off-highway (6%, mainly construction use), farming (5%), and commercial (5%) [4]. However, during a hurricane, additional demand arises to power backup generators for critical infrastructure, while transportation of that fuel may be interrupted. The typical daily peak electric load for the entire State of Florida is approximately 27.2 GW. We assume that the hurricane can damage a portion of the power system, and present the impact of various levels of damage to the DFSC in Section 4.5.

Infrastructure damage during a hurricane is mainly caused by high winds, storm surge, and inland flooding [55]. Hurricane Irma made landfall as a Category 4 hurricane in the Florida Keys on September 10, 2017, and struck southwestern Florida at Category 3 intensity. The center of Irma passed through the western part of Florida; see Figure 1. According to the NHC [22], peak sustained wind speeds were 115 kt in the Florida Keys and about 100 kt in Southwest Florida, near Marco Island. The maximum sustained wind recorded inland across South Florida averaged about 60 kt. For portions of the Lower Florida Keys, the storm surge produced inundation levels of 5 to 8 ft. As Irma progressed north, the maximum sustained wind speed decreased to about 55 kt on the east coast and 40 kt on the west, even though the hurricane's track was along the western coast of the state. This contrast is caused by Irma's large wind field, and similar effects were observed for storm surge: while Irma led to a maximum inundation level of 3-5 ft between Naples and Fort Myers, storm surge flooding of up to 6 ft occurred in Miami-Dade County. In addition, heavy rain was brought to much of the state by Hurricane Irma, "rainfall totals of 10 to 15 inches were common across the peninsula and the Keys" [22].

4.1. From Weather Data to Diesel Fuel Demand

An important aspect of our analysis is to use data that decision makers can obtain in real time, as a hurricane approaches. Such resources include hurricane advisories and weather forecasts from the NHC. However, only having weather data is not enough to estimate the impact of a hurricane. Our analysis, linking the weather forecast data and the demand for diesel fuel, helps answer the following questions, keeping in mind

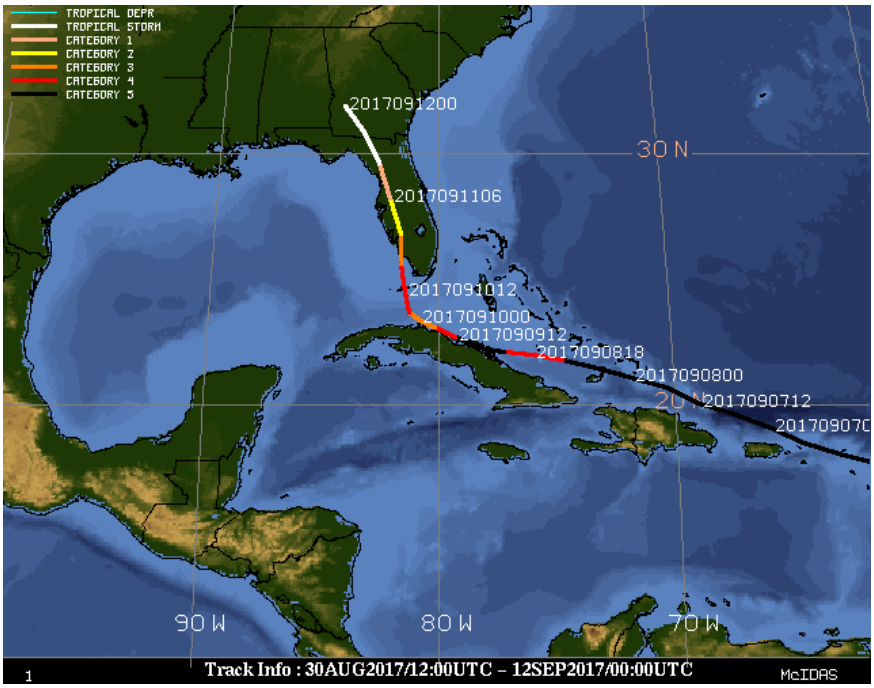

Figure 1. Track of Hurricane Irma: the intensity in Saffir-Simpson Hurricane Wind Scale is marked in different colors and the time stamps are displayed in white text. Source: [17].

that diesel is mainly used for emergency power generation for hurricane relief:

(1) Given the weather conditions (e.g., wind speeds) at a county, what is the expected power loss?
(2) How is the power loss mapped to the diesel demand?
(3) How does the power loss change in the recovery phase after the hurricane passes?

All data collected fall between September 9, 2017 and September 25, 2017, which contains the time period during which Irma passed through the state. Florida has 67 counties and all the data are collected at the county level. We obtained the outage report from Florida Division of Emergency Management [31]. The report is generated every three hours and details the number of customers without power. We make the following assumptions to help transform the outage data to the diesel demand:

- The power demand of the State of Florida is $\bar{D} = 27.2$ GW.
- The power load of a county is proportional to its population. We denote the population fraction (relative to Florida's population) of county i as f_i.
- The outage percentage of county i at time period t is denoted $o_{i,t}$. Therefore, the lost load of county i at time t is $\bar{D} f_i o_{i,t}$.
- The conversion constant from diesel fuel to electric energy is $\rho = 1688.6$ barrel/GWh [9].
- The fulfillment rate of lost load is $\alpha \in [0, 1]$, which specifies the fraction of total load that "should" be fulfilled. We vary parameter α as we test the DFSC's capability with different fulfillment rates.

With the above assumptions, the amount of diesel required (in barrels) to meet emer-

gency generation load in county i within one hour, starting at time t, is:

$$d_{i,t} = \alpha\rho\bar{D}f_i o_{i,t}. \tag{2}$$

While there is no direct forecast of the outage percentage, $o_{i,t}$, we can attempt to map weather conditions to outage levels. We obtain actual weather data from the Local Climatological Data (LCD) and the weather forecast data from the National Digital Forecast Database (NDFD), both maintained by NOAA. LCD data in Florida is collected at 42 locations, each a weather station in a distinct county, usually at an airport, with sustained wind speed, gust speed, pressure, and precipitation recorded hourly. NDFD forecasts are made every three hours at a grid of points for the following 168 hours (7 days), and data for the area between the grid points is interpolated by NOAA algorithms. We collect the NDFD data for every county in Florida and the data contain temperature, sustained wind speed, gust speed, precipitation level, humidity, probability of surface wind speed exceeding 34 kt, 50 kt, and 64 kt.

We also collect storm surge data from the Probabilistic Tropical Storm Surge (P-Surge) model by NOAA, and high water mark (HWM) data from the United States Geological Survey (USGS) [66] to estimate storm surge and resulting flooding levels. The P-Surge model estimates the probability that storm surge exceeds a certain height for the same grid of points as NDFD data. For a given time, we use the expected storm surge height at the closest point. HWM data record the height of water marks left by flooding on buildings or trees, which can show the peak flooding level.

For each county we can plot how the outage percentage changed as the hurricane progressed, as shown in Figure 2 for Miami-Dade County. From such plots we note:

Observation 1 The outage percentage grows from zero before reaching the peak in a synchronized manner as the wind speed and the storm surge height (not shown) increases;

Observation 2 After its peak, the outage percentage does not appear correlated with the weather predictors (gust speed, storm surge, etc.);

Observation 3 After the peak, the outage percentage decreases in a roughly linear fashion with the time elapsed since the peak, perhaps due to maintenance efforts.

Based on these observations, we build a quantitative model with two parts: (i) a predictive model to map from the maximum level of a weather condition (e.g., wind speed) to the peak outage percentage; and (ii) a model to estimate the recovery rate from the outage.

There has been research about predicting hurricane power outages [e.g., 35, 37, 38]. We use a similar statistical analysis approach, but we do not include geographic features such as the land cover and soil variables because our outages are aggregated by county. Our predictors, in vector x, include the maximum gust speed, the maximum storm surge height, and the highest recorded water mark height, because damage to the power system is determined by the worst weather condition; the response variable, y, is the percentage of power outage. We develop a logistic regression model as:

$$y_i = \left(1 + e^{-(\beta^\top x_i + \beta_0)}\right)^{-1} + \epsilon_i, \qquad i \in \mathcal{I}, \tag{3}$$

where i again indexes counties. Before fitting the regression model, we add one artificial data point $x = 0, y = 0$ to represent the absence of a hazard. Our fitted results show that the storm surge height and the flooding height are not statistically significant—for

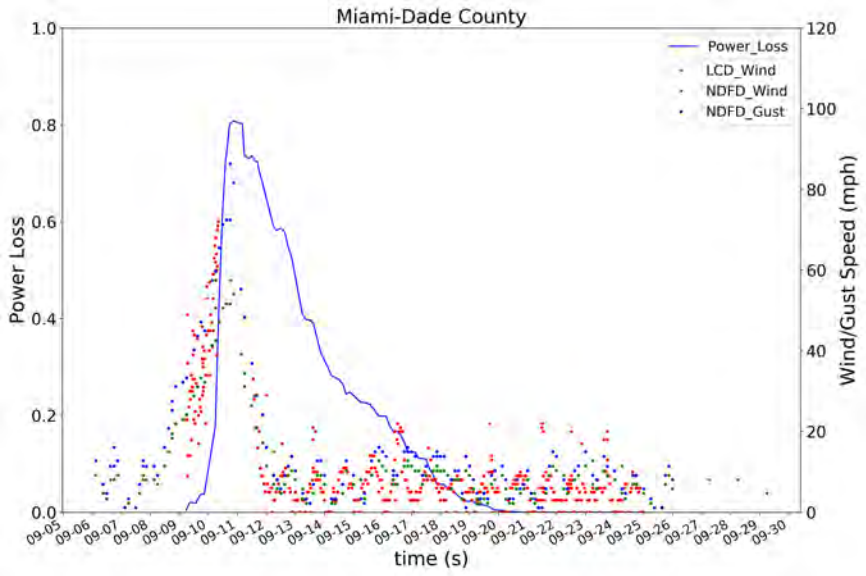

Figure 2. Power outage percentage ($o_{i,t}$), LCD gust speed, and NDFD 3-hour forecast gust speed for Miami-Dade County between September 9, 2017 and September 25, 2017.

Hurricane Irma—even when we add an auxiliary variable indicating coastal counties, which is similar to the result in Guikema et al. [35]. If we keep the maximum gust speed as the only predictor, we obtain the estimated parameters as $\beta = 0.0889$, and $\beta_0 = -6.388$. Figure 3 shows that the nonlinear regression model (blue curve) fit to the data (black dots) across the 67 counties of Florida. For this regression model we obtain $R^2 = 0.523$.

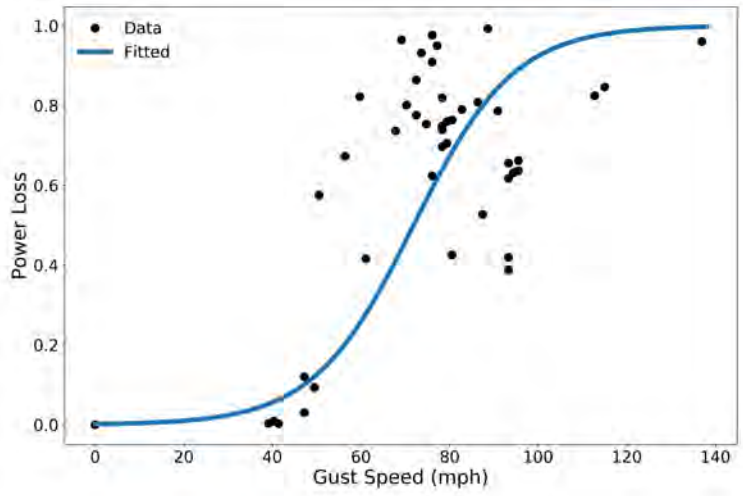

Figure 3. Logistic regression model between the maximum gust speed and the percentage of power loss in the county.

Power outages decrease during the recovery process after the hurricane passes. We approximate the recovery process using a linear function of the elapsed time since the

peak outage. Figure 4 plots the peak power loss versus the time to recover to "normal operations," defined here as less than 1% outage, across Florida's 67 counties. The fitted coefficients of the linear regression model are $\bar{\beta} = 171.59$ and $\bar{\beta}_0 = -7.10$, with $R^2 = 0.6661$. The goal is to estimate how long it takes to return to normal operations so that diesel fuel is no longer needed for emergency generation demand, given the peak level of power outage. For example, the estimated recovery time from an 80% outage is $0.8 \times 171.59 - 7.10 = 131.17$ hours. In addition, since we assume a linear recovery rate (1.7159 hours to recover 1% of the outage), we can estimate the outage percentage given the elapsed time since the hurricane peak. In summary, given a time series forecasting gust speeds at county i, w_{it}, $t = 1, 2, \ldots$, and the time index corresponding to the maximum gust speed denoted as $\hat{t}$, we can estimate the power outage rate $o_{i,t}$ using the following:

$$o_{i,t} = \begin{cases} \left(1 + e^{-(\beta^\top w_{it} + \beta_0)}\right)^{-1} & t \leq \hat{t} \\ \bar{\beta}^{-1}(t - \hat{t} - \bar{\beta}_0) & t > \hat{t}. \end{cases} \tag{4}$$

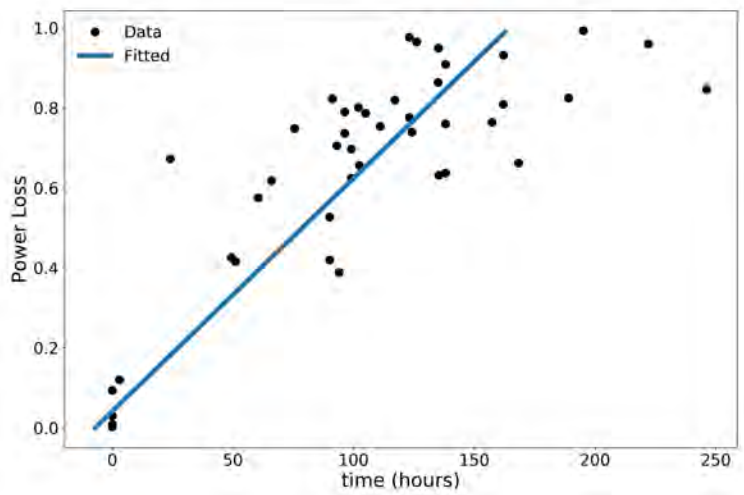

Figure 4. Estimated linear relationship between the recovery time and the maximum power outage level.

4.2. Construction of Hurricane Scenarios

With the quantitative models of Section 4.1, given a time series of forecast wind speeds, we can estimate the power outage at corresponding times, and this section describes how we do so using the GEFS database. GEFS is an ensemble of 21 separate weather forecasting models. A forecast is made every three hours and contains 16 days of weather information, including the wind gust speed. GFS forecasts are for a grid of points with 1° resolution, with estimates for other points again available by NOAA's interpolation e.

We assume equal probabilities for each of these 21 scenarios. When a GEFS forecast is made, we obtain 21 scenarios of wind gust speeds for a county for the next 16 days.

Each scenario for these speeds is then mapped to a time series of diesel demand using equations (2) and (4). Figure 5 illustrates the comparison between the estimated diesel demand from two GEFS scenarios, the mapped demand from the NDFD gust speed forecasts (available for a shorter horizon), and the diesel demand calculated according to actual outage data.

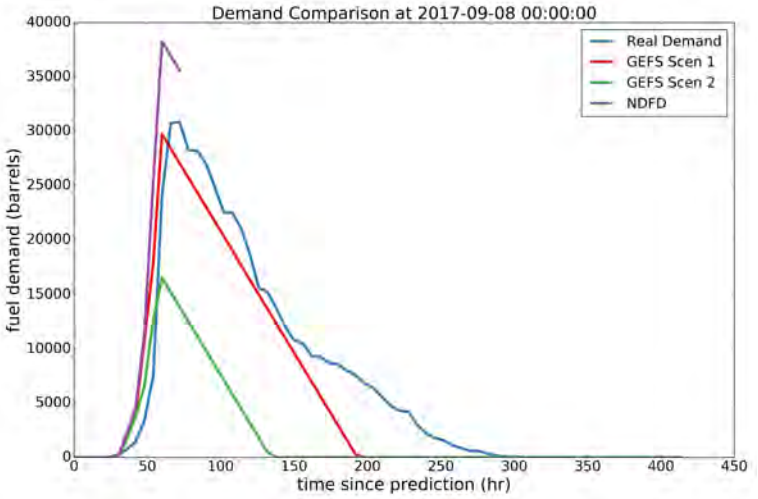

Figure 5. Comparison between the diesel demand profile based on actual outages, the mapped diesel demand from two forecast gust speed scenarios, and the mapped diesel demand from NDFD forecast for Miami-Dade County at 12pm EST September 8th, 2017. Fulfillment rate $\alpha = 100\%$.

4.3. Shortage Cost Setups

An increasing piece-wise linear convex cost function is often used to model consequences of growing severity as unsatisfied demand grows (see, e.g., Schütz et al. [62]). We model penalty costs for failing to satisfy nominal demand, c_t^n, and surge demand, c_t^e, as piece-wise linear convex functions that are parameterized with K^n and K^e pieces, respectively:

$$c_t^n(z_{\cdot,t}^n) = \max_{k=1,\dots,K^n} \left\{ a_k^n + m_k^n \sum_{i \in \mathcal{I}_d} z_{i,t}^n \right\} \qquad \forall t \in \mathcal{T} \tag{5a}$$

$$c_t^e(z_{\cdot,t}^e) = \sum_{i \in \mathcal{I}_d} \left[\max_{k=1,\dots,K^e} \left\{ a_{k,i}^e + m_{k,i}^e z_{i,t}^e \right\} \right] \qquad \forall t \in \mathcal{T}. \tag{5b}$$

Here coefficients $m_{k,i}^e, a_{k,i}^e$ denote the slope and intercept for county-specific surge shortage while coefficients m_k^n, a_k^n are analogous but for nominal shortage. These functions are incorporated in model (1) by adding auxiliary variables and constraints akin to the right-hand side of constraints (5a) and (5b). Our choice for the penalty functions is motivated by the fact that the effects of nominal and surge shortages at different levels. Because trucks run on diesel fuel, if the total nominal shortage is too high, we do not have enough diesel fuel to dispatch trucks and the transportation network cannot function. Therefore, in equation (5a), the nominal shortage cost is calculated in terms of the aggregated nominal shortage across the DFSC. On the other hand,

surge shortage cost described in equation (5b) is calculated at each county due to the cost of not powering local facilities.

While we distinguish nominal and surge demand in global and local ways, as just described, our piece-wise linear functions have a total of four pieces and slopes with $K^e = K^n = 2$, ordered in decreasing steepness: (i) nominal demand, 90-100%; (ii) surge demand, 50%-100%; (iii) surge demand, 0%-50%; and (iv) nominal demand, 0-90%. Thus the most important demand to satisfy is global nominal demand, until "only" 90% remains unmet. Then, we should satisfy surge demand in each county until 50% remains unmet, and then again until surge demand is fully satisfied. The remaining fuel is then used to satisfy nominal demand. We assume 10% of the nominal demand is needed for key transportation across the state, including for tank trucks, and this is why the first piece of the nominal demand has the steepest shortfall cost. After that, we prioritize surge demand at critical infrastructure such as hospitals, municipal water systems, and facilities for first responders. Therefore, we set a high penalty cost for the surge shortage, with two levels of priority. Table 4 summarizes the values for a^n_k, m^n_k, $a^e_{k,i}$, and $m^e_{k,i}$ for a given time period t.

Piece k	Surge		Nominal	
	$m^e_{k,i}$	$a^e_{k,i}$	m^n_k	a^n_k
1	2	0	1	0
2	5	$(m^e_{2,i} - m^e_{1,i}) \cdot 0.5 \cdot \hat{d}^e_{i,t}$	10	$(m^n_2 - m^n_1) \cdot 0.9 \cdot \sum_{i \in \mathcal{I}_d} d^n_{i,t}$

Table 4. Parameters used to define penalty functions for failing to satisfy demand for diesel fuel. See equation (5) and model (1).

4.4. Solution Stability

Model (1) can have multiple optimal solutions, particularly in the early part of the planning period, because multiple dispatch plans achieve zero shortage before the hurricane arrives and demand surges. Under the rolling-horizon approach, small changes in the solution in early time periods can lead to differences as large as 5% in the expected shortage cost as we replicate experiments. To obtain a stable solution, we incorporate a tie-breaking mechanism by small adjustments in the objective function:

- for variables $I^d_{i,t}$, we use a per-barrel reward for inventory of fuel at demand sites with a coefficient of $-(1.001)^t f_i |\mathcal{T}_f|^{-1}$;
- for variables x and y, we use a penalty of $10^{-4} \cdot \tau_{ij}$ to discourage long-distance transport, where τ_{ij} is the travel time i from j embedded in sets $\mathcal{A}^x$ and $\mathcal{A}^y$.

Adding these tie-breaking coefficients does not degrade the total penalty cost for models with a long horizon, but helps significantly in stabilizing solutions in the rolling-horizon procedure. We employ these objective function terms in all experiments in the following section.

4.5. Experimental Results

We present experimental results of our DFSC optimization model in this section. We compare shortfall penalties for surge and nominal demand obtained by running Algorithm 1 with the stochastic programming model (1), using deterministic alternatives as benchmarks, as described in Section 3.3. We also test our model with controlled

parameters, H, F, R, N and α (see Sections 3.2 and 4.1) in order to evaluate their effect on the model and obtain operational insights. All results that we report use either a point forecast or a set of scenarios to dispatch tank trucks, but use actual power outages from Irma, and the associated demand for diesel, to compute shortages.

We use a single thread on a 2.4 GHz Intel Xeon core with 8 GB of RAM to run all tests. All models are constructed using the Gurobi Python API and solved by Gurobi 8.1.1 [36] with default parameters. For all tests, solving both the optimization problem (1) and all deterministic alternatives takes less than one hour, which is the unit time period. Thus, the computational performance does not affect the execution of the rolling-horizon approach, and so we focus our presentation on the value of our model and analysis of model parameters, rather than the required computational effort.

4.5.1. Comparison with Deterministic Benchmarks

We compare three deterministic alternatives to our stochastic program (denoted "GEFS"), to examine the value of modeling uncertainty and how much improvement can be made under a perfect hurricane forecast. As described in Section 3.3, as an alternative we can run Algorithm 1 with deterministic demand profiles instead of 21 scenarios. The ways of doing so are via NDFD, average GEFS (denoted "GAVG"), and perfect information (denoted "PI"). For NDFD, we use the point forecast for wind gust speed from NDFD as the deterministic forecast; for GAVG, we use the mean of 21 scenarios; for PI, we use the actual demand data, i.e., based on actual outages. In addition, we solve a deterministic optimization problem with demand information for the entire problem horizon T, denoted "BEST." While PI essentially provides a benchmark on best practice, *given* a set of H, F, R and N, BEST gives the best achievable practice with the distinction being that PI can only see to the model horizon H while BEST has access to perfect information to the problem horizon T.

Figure 6 shows surge and nominal shortages and the total demand across time. These results assume $\alpha = 0.5$ of the demand must be satisfied and uses rolling-horizon parameters of $H = 96, F = 72, R = 24$ and $N = 24$ hours. The colored lines specify shortages for the four alternatives, and the black line indicates total demand, accounting for α. The cumulative shortage and penalty results at the end of the problem time horizon are summarized in Table 5.

Alternative	Surge Shortage	Nominal Shortage	Total Shortage	Surge Penalty	Nominal Penalty	Total Penalty
GEFS	161.4	1012.4	1173.8	370,309.9	1,012,352.2	1,382,662.0
NDFD	469.0	705.3	1174.3	1,227,241.5	705,335.0	1,932,576.5
GAVG	204.2	969.6	1173.8	489,425.8	969,608.0	1,459,033.8
PI	1.9	1171.6	1173.5	3825.1	1,171,584.0	1,175,409.1
BEST	0.0	1173.5	1173.5	0.0	1,173,496.6	1,173,496.6

Table 5. Cumulative surge, nominal and total shortages and their penalties at the end of time horizon for different alternatives, with parameters $\alpha = 0.5$ and $H = 96, F = 72, R = 24$ and $N = 24$ hours. The unit for demand data is thousand of barrels.

If more diesel is delivered to a county than its demand, it goes unused, but most diesel is consumed. Thus we observe from Figure 6 and Table 5 that the total shortages for different alternatives are similar. The table and figure show that the alternatives differ in their treatment of surge demand and nominal demand. Given the cost structure in Section 4.3, after satisfying the first 10% of nominal demand, surge demand should be prioritized over nominal demand.

PI yields a total penalty cost 0.1% higher than BEST as PI makes decisions looking only 96 hours into the future. Together PI and BEST achieve the lowest penalty costs by incurring larger nominal shortages early in the time horizon so that accumulated inventory is better used when the hurricane arrives. This suggests that knowing the demand ahead of time can save at least 15% of costs compared to the next best alternative, GEFS.

Putting aside a perfect weather forecast, GEFS performs better than GAVG and NDFD, which shows the value of incorporating multiple scenarios. GEFS achieves 161.4 thousand barrels of surge shortage and a total cost of about 1.38×10^6, which is 28.5% and 5.2% lower than respective total costs of NDFD and GAVG. Thus, there is value in hedging against multiple possible hurricane paths.

Alternative	Surge Shortage	Nominal Shortage	Total Shortage	Surge Penalty	Nominal Penalty	Total Penalty
GEFS	142.7	1031.0	1173.8	312,191.3	1,031,032.5	1,343,223.9
NDFD	575.2	599.1	1174.2	1,675,772.6	599,072.6	2,274,845.3
GAVG	157.9	1015.9	1173.8	356,436.9	1,015,872.9	1,372,309.8
PI	0.0	1173.5	1173.5	0.0	1,173,496.6	1,173,496.6

Table 6. Cumulative surge, nominal and total shortages and their penalties at the end of time horizon for different alternatives, with parameters $\alpha = 0.5$ and $H = 96, F = 72, R = 12$ and $N = 12$ hours. The unit for demand data is thousand of barrels.

We repeat the same analysis for the same horizon ($H = 96$) and first-stage length ($F = 72$), but a shorter rolling period ($R = 12$) and a shorter span to make decisions in advance ($N = 12$), with results in Table 6 and Figure 7. With more frequent updates of our decisions, and more up-to-date forecasts, GEFS, GAVG, and PI all improve over the results of Table 5 while NDFD degrades, surprisingly. What is consistent is that NDFD's results are inferior to those of GAVG and GEFS. These two test instances are also consistent with the general trend of relative costs of these alternatives:

$$\text{PI} < \text{GEFS} < \text{GAVG} < \text{NDFD}.$$

The combination of the shorter rolling period and more up-to-date forecasts translates into a more informed and adaptive response. The model with perfect information reduces the surge shortage to zero, and GEFS and GAVG decrease surge shortage. Comparing the two figures, we see that the nominal shortage curve shifts to the left and the surge shortage is lower in Figure 7, as the optimization starts moving fuel for nominal demand earlier to accumulate inventory in preparation for the upcoming surge demand. For NDFD, this increased adaptability does not translate to a decreased surge shortage or a decreased total cost, potentially because the NDFD forecast is significantly more variable than GEFS over time.

4.5.2. Analysis of Model Parameters

The above comparison between two tests with different R and N motivates us to explore how different parameters affect performance. We run multiple tests, with deviations from the base case of $H = 96, F = 24, R = 24, N = 24, \alpha = 0.5$ to examine the impact of each of H, F, R, N and α on the shortage and the penalty cost.

Table 7 shows the results of the parametric tests, with the results split into six groups. The first group shows that an increased first-stage length decreases the surge shortage and total cost up to a limit, i.e., some stochastic hedging brings value. The second and third sets of results show value in simultaneously increasing the model

(H,F,R,N,α)	Surge Shortage	Nominal Shortage	Total Shortage	Surge Penalty	Nominal Penalty	Total Penalty
(96,24,24,24,0.5)	197.8	976.0	1173.8	459,564.7	975,999.3	1,435,564.0
(96,48,24,24,0.5)	172.7	1001.1	1173.8	391,341.4	1,001,080.2	1,392,421.6
(96,72,24,24,0.5)	161.4	1012.4	1173.8	370,309.9	1,012,352.2	1,382,662.0
(96,96,24,24,0.5)	204.2	969.6	1173.8	489,425.8	969,608.0	1,459,033.8
(96,24,24,24,0.5)	197.8	976.0	1173.8	459,564.7	975,999.3	1,435,564.0
(120,48,24,24,0.5)	158.4	1015.4	1173.8	363,113.7	1,015,368.6	1,378,482.3
(144,72,24,24,0.5)	153.3	1020.5	1173.8	347,197.0	1,020,467.1	1,367,664.1
(96,24,24,24,0.5)	197.8	976.0	1173.8	459,564.7	975,999.3	1,435,564.0
(72,24,24,24,0.5)	201.7	972.1	1173.8	470,461.3	972,090.6	1,442,552.0
(48,24,24,24,0.5)	328.3	845.5	1173.8	751,868.0	845,507.9	1,597,375.9
(96,24,24,24,0.5)	197.8	976.0	1173.8	459,564.7	975,999.3	1,435,564.0
(96,24,12,24,0.5)	176.4	997.3	1173.6	390,568.3	997,265.8	1,387,834.2
(96,24,6,24,0.5)	159.4	1014.2	1173.7	341,405.3	1,014,229.9	1,355,635.2
(96,24,24,24,0.5)	197.8	976.0	1173.8	459,564.7	975,999.3	1,435,564.0
(96,24,24,12,0.5)	201.5	972.4	1173.9	421,761.8	972,370.9	1,394,132.7
(96,24,24,6,0.5)	142.6	1030.9	1173.5	316,316.1	1,030,869.4	1,347,185.5
(96,24,24,24,1.0)	1112.0	1280.2	2392.2	2,626,870.5	1,280,238.1	3,907,108.7
(96,24,24,24,0.5)	197.8	976.0	1173.8	459,564.7	975,999.3	1,435,564.0
(96,24,24,24,0.2)	20.4	422.3	442.7	45,723.6	422,341.4	468,065.0

Table 7. Cumulative surge, nominal and total shortages and their penalties at the end of time horizon for different sets of parameters, with parameters $\alpha = 0.5, H = 96, F = 72, R = 24$ and $N = 24$ serving as the base case. The unit for demand data is thousand of barrels.

horizon H and the length of the first stage F. The fourth and fifth sets suggest that as we shorten the time periods to roll forward, R, or the delay in the forecast, N, the total cost and the surge shortage decreases, because it allows more flexible and adaptive planning. These quantify the value of pursuing a more flexible and quick-responsive operation.

Finally, we inspect how the shortage and total cost are affected by different fulfillment rates. When the supply is fixed, an increased fulfillment rate increases demand and leads to larger shortages and higher costs. However, the relationship is not linear in that the surge shortage, total shortage, and shortage costs grow more quickly as the fulfillment rate grows. This suggests that it is key to prioritize critical loads and that there would be benefits to local self-reliance either by reducing loads or via storm-resilient microgrids.

5. Conclusions

In this paper, we formulated and solved a model for diesel fuel supply chain operations under a disruption caused by a hurricane. It is key to be proactive in planning before the hurricane arrives, accounting for its uncertain path and intensity. We utilized weather forecasts from NOAA to construct hurricane scenarios and mapped the weather data to demand for diesel fuel, which is mainly used in emergency power generation and transportation of fuel. We proposed a rolling-horizon approach: at each decision point, we solved a two-stage stochastic program with the GEFS-informed hurricane scenarios, committed to the dispatch of tank trucks until the next time a decision is made, and a new stochastic program was solved with updated scenarios.

We created a case study for Hurricane Irma, considering realistic settings in modeling diesel demand from accessible public data. Computational results justify the value

of our stochastic programming model, and sensitivity analysis on model parameters provide useful insights about valuable improvements for preparedness.

In the context of Hurricane Irma, only wind speed was predictive of power outages, but more generally we expect further weather predictors, such as rainfall and storm surge, to add predictive value. Our rolling-horizon approach can be extended to a multi-stage stochastic program, which poses more challenging issues in how to construct a reasonable model that can be decomposed and solved efficiently. Moreover, it is important to coordinate the DFSC and other infrastructure networks for hurricane relief. In the future, we aim to augment our model, and rolling-horizon algorithm, with interfaces to port operations [27], transportation networks [50], and electricity networks [72] to better inform decision support for hurricane response and relief.

Acknowledgments

This work was supported by the U.S. Department of Homeland Security under Grant Award 2017-ST-061-QA0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security. We thank the Center for Nonlinear Studies at Los Alamos National Laboratory for partially supporting Haoxiang Yang's work. The authors thank Lauren Davis, Pitu Mirchandani, and Giulia Pedrielli for multiple conversations that contributed to our work. And, the authors thanks Craig S. Gordon, Laura Laybourn, and Eric Rollison for suggestions that helped frame the analysis.

References

[1] Is your water or wastewater system prepared? What you need to know about generators. Technical Report EPA 901-F-09-027, United States Environmental Protection Agency New England, September 2009.

[2] Power resilience: Guide for water and wastewater utilities. Technical Report EPA 800–R–15–004, United States Environmental Protection Agency, December 2015.

[3] Hurricane Irma: Infrastructure impact assessment: Florida refined petroleum product availability 1130 EDT September 8, 2017. Technical report, National Protection and Programs Directorate, Office of Cyber and Infrastructure Analysis, September 2017.

[4] Florida sales of distillate fuel oil by end use, December 2017. URL `https://www.eia.gov/dnav/pet/pet_cons_821dst_dcu_SFL_a.htm`.

[5] The 2017 Florida statutes, Title XXXIII, Chapter 526, March 2018. URL `http://www.leg.state.fl.us/statutes/index.cfm?App_mode=Display_Statute&URL=0500-0599/0526/Sections/0526.143.html`.

[6] Role of diesel generators in the healthcare industry, 2018. URL `http://www.dieselserviceandsupply.com/Generators_Healthcare.aspx`. Accessed: 2018-03-05.

[7] Diesel generators provide power in critical times, March 2018. URL `http://www.dieselserviceandsupply.com/Hurricanes_Backup_Power.aspx`. Accessed: 2018-03-05.

[8] Hurricane Irma's effect on Florida's fuel distribution system and recommended improvements. Technical report, The Florida Department of Transportation, January 2018. URL `http://www.fdot.gov/info/CO/news/newsreleases/020118_FDOT-Fuel-Report.pdf`.

[9] Approximate diesel fuel consumption chart, 2019. URL `https://www.`

dieselserviceandsupply.com/temp/Fuel_Consumption_Chart.pdf. Accessed: 2019-08-21.

[10] C. Abbey, D. Cornforth, N. Hatziargyriou, K. Hirose, A. Kwasinski, E. Kyriakides, G. Platt, L. Reyes, and S. Suryanarayanan. Powering through the storm: microgrids operation for more efficient disaster recovery. *IEEE Power and Energy Magazine*, 12(3): 67–76, 2014.

[11] A. Adhitya, R. Srinivasan, and I. A. Karimi. Heuristic rescheduling of crude oil operations to manage abnormal supply chain events. *AIChE Journal*, 53(2):397–422, 2007.

[12] A. M. Afshar and A. Haghani. Heuristic framework for optimizing hurricane evacuation operations. *Transportation Research Record*, 2089(1):9–17, 2008.

[13] D. Alem, A. Clark, and A. Moreno. Stochastic network models for logistics planning in disaster relief. *European Journal of Operational Research*, 255(1):187–206, 2016.

[14] G. Alfke, W. W. Irion, and O. S. Neuwirth. Oil refining. *Ullmann's Encyclopedia of Industrial Chemistry*, 2007.

[15] N. Altay and W. G. Green. OR/MS research in disaster operations management. *European Journal of Operational Research*, 175(1):475–493, 2006.

[16] J. S. Aronofsky and A. C. Williams. The use of linear programming and mathematical models in under-ground oil production. *Management Science*, 8(4):394–407, 1962.

[17] Scott Bachmeier. Hurricane Irma storm-track. `http://tropic.ssec.wisc.edu/storm_archive/2017/storms/11L/11L.html`, 2017. Accessed: 2019-08-18.

[18] F. Bajak and R. Dunklin. Explosions rock flood-crippled chemical plant near Houston, Sptember 2017. URL `https://apnews.com/43288ccafa394ae59fd8a548b23a29a1`.

[19] J. D. Bales. Effects of Hurricane Floyd inland flooding, September–October 1999, on tributaries to Pamlico Sound, North Carolina. *Estuaries*, 26(5):1319–1328, Oct 2003.

[20] A. Beheshtian, K. P. Donaghy, R. R. Geddes, and O. M. Rouhani. Planning resilient motor-fuel supply chain. *International Journal of Disaster Risk Reduction*, 24:312–325, 2017.

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

[22] J. P. Cangialosi, A. S. Latto, and R. Berg. National Hurricane Center tropical cyclone report: Hurricane Irma. Research report AL112017, National Hurricane Center, 2018. URL `https://www.nhc.noaa.gov/data/tcr/AL112017_Irma.pdf`.

[23] A. Charnes, W. W. Cooper, and B. Mellon. Blending aviation gasolines–a study in programming interdependent activities in an integrated oil company. *Econometrica*, 20 (2):135–159, 1952.

[24] A. Charnes, W. W. Cooper, and B. Mellon. A model for programming and sensitivity analysis in an integrated oil company. *Econometrica*, 22(2):193–217, 1954.

[25] D. Cullen. Fuel market impact of Hurricanes Harvey & Irma, September 2017. URL `http://www.breakthroughfuel.com/blog/fuel-market-impact-of-hurricanes-harvey-irma-advisor-pulse/`.

[26] R. A. Davidson, L. K. Nozick, T. Wachtendorf, B. Blanton, B. Colle, R. L. Kolar, S. DeYoung, K. M. Dresback, W. Yi, K. Yang, and N. Leonardo. An integrated scenario ensemble-based framework for hurricane evacuation modeling: Part 1—decision support system. *Risk Analysis*, 40(1):97–116, 2020.

[27] L. B. Davis and D. Edwards. Determining optimal fuel strategies under uncertainty. Research report, North Carolina A&T State University, 2020.

[28] L. B. Davis, F. Samanlioglu, X. Qu, and S. Root. Inventory planning and coordination in disaster relief efforts. *International Journal of Production Economics*, 141(2):561–573, 2013.

[29] M. A. H. Dempster, N. H. Pedron, E. A. Medova, J. E. Scott, and A. Sembos. Planning logistics operations in the oil industry. *Journal of the Operational Research Society*, 51 (11):1271–1288, 2000.

[30] L. F. Escudero, F. J. Quintana, and J. Salmerón. CORO, a modeling and an algorithmic framework for oil supply, transformation and distribution optimization under uncertainty. *European Journal of Operational Research*, 114(3):638–656, 1999.

[31] Florida Division of Emergency Management. Hurricane Irma: Power outage data, September 2017. URL `http://floridadisaster.org/info/outage_reports/irma/`.

[32] European Centre for Medium-Range Weather Forecasts. Modelling and prediction, March 2019. URL `https://www.ecmwf.int/en/research/modelling-and-prediction`.

[33] Global Climate and Weather Modeling Branch. The GFS atmospheric model. Technical report, National Weather Service, 2003. URL `https://www.emc.ncep.noaa.gov/officenotes/newernotes/on442.pdf`.

[34] A. D. González, L. Dueñas Osorio, M. Sánchez-Silva, and A. Medaglia. The interdependent network design problem for optimal infrastructure system restoration. *Computer-Aided Civil and Infrastructure Engineering*, 31(5):334–350, 2016.

[35] S. D. Guikema, R. Nateghi, S. M. Quiring, A. Staid, A. C. Reilly, and M. Gao. Predicting hurricane power outages to support storm response planning. *IEEE Access*, 2:1364–1373, 2014.

[36] Gurobi Optimization, Inc. *Gurobi Optimizer Reference Manual*, 2016. URL `http://www.gurobi.com`.

[37] S. R. Han, S. D. Guikema, and S. M. Quiring. Improving the predictive accuracy of hurricane power outage forecasts using generalized additive models. *Risk Analysis: An International Journal*, 29(10):1443–1453, 2009.

[38] S. R. Han, S. D. Guikema, S. M. Quiring, K. H. Lee, D. Rosowsky, and R. A. Davidson. Estimating the spatial distribution of power outages during hurricanes in the Gulf Coast region. *Reliability Engineering & System Safety*, 94(2):199–210, 2009.

[39] M. L. Itria, M. Kocsis-Magyar, A. Ceccarelli, P. Lollini, G. Giunta, and A. Bondavalli. Identification of critical situations via event processing and event trust analysis. *Knowledge and Information Systems*, 52(1):147–178, 2017.

[40] A. Kwasinski. Hurricane Sandy effects on communication systems. Technical Report PR-AK-0112-2012, The University of Texas at Austin, 2012.

[41] A. Kwasinski. Lessons from field damage assessments about communication networks power supply and infrastructure performance during natural disasters with a focus on Hurricane Sandy. In *FCC Workshop on Network Resiliency*, 2013.

[42] D. R. Levinson. Hospital emergency preparedness and response during superstorm Sandy. Technical Report OEI-06-13-00260, Department of Health and Human Services, 2014.

[43] X. Li, R. Batta, and C. Kwon. Effective and equitable supply of gasoline to impacted areas in the aftermath of a natural disaster. *Socio-Economic Planning Sciences*, 57:25–34, 2017.

[44] C. Lima, S. Relvas, and A. P. F. D. Barbosa-Póvoa. Downstream oil supply chain management: a critical review and future directions. *Computers & Chemical Engineering*, 92: 78–92, 2016.

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

[46] E. J. Lodree and S. Taskin. A Bayesian decision model with hurricane forecast updates for emergency supplies inventory management. *Journal of the Operational Research Society*, 62(6):1098–1108, 2011.

[47] N. Lorenzi. Critical features of emergency power generators, September 2015. URL `https://www.hfmmagazine.com/articles/1712-critical-features-of-emergency-power-generators`.

[48] R. A. Luettich, J. J. Westerink Jr., and N. W. Scheffner. ADCIRC: an advanced three-dimensional circulation model for shelves coasts and estuaries, report 1: theory and methodology of ADCIRC-2DDI and ADCIRC-3DL. Technical Report DRP-92-6, U.S. Army Engineers Waterways Experiment Station, 1992.

[49] M. Merakl ı and S. Küçükyavuz. Risk aversion to parameter uncertainty in markov decision processes with an application to slow-onset disaster relief. *IISE Transactions*, pages 1–21,

2019. URL https://doi.org/10.1080/24725854.2019.1674464.
[50] P. Mirchandani and K. G. Ayu. A data-driven probabilistic simulation model and visualization for hurricane. Research report, Arizona State University, 2020.
[51] National Hurricane Center. Hurricane preparedness: Hazards, July 2019. URL https://www.nhc.noaa.gov/prepare/hazards.php.
[52] National Hurricane Center. NHC track and intensity models, March 2019. URL https://www.nhc.noaa.gov/modelsummary.shtml.
[53] S. M. S. Neiro and J. M. Pinto. A general modeling framework for the operational planning of petroleum supply chains. *Computers & Chemical Engineering*, 28(6-7):871–896, 2004.
[54] N. Noyan. Risk-averse two-stage stochastic programming with an application to disaster management. *Computers & Operations Research*, 39(3):541–559, 2012.
[55] U.S. Department of Homeland Security. Foundational methodology to support infrastructure decision analysis: Initial development efforts and test case. Research report, U.S. Department of Homeland Security, 2009.
[56] Met Office. Met Office numerical weather prediction models, March 2019. URL https://www.metoffice.gov.uk/research/modelling-systems/unified-model/weather-forecasting.
[57] G. G. Pacheco and R. Batta. Forecast-driven model for prepositioning supplies in preparation for a foreseen hurricane. *Journal of the Operational Research Society*, 67:98–131, 2016.
[58] T. Powell, D. Hanfling, and L. O. Gostin. Emergency preparedness and public health: the lessons of Hurricane Sandy. *Journal of the American Medical Association*, 308(24): 2569–2570, 2012.
[59] C. G. Rawls and M. A. Turquist. Pre-positioning of emergency supplies for disaster response. *Transportation Research Part B: Methodological*, 44:521–534, 2010.
[60] M. Rezaei-Malek, R. Tavakkoli-Moghaddam, B. Zahiri, and A. Bozorgi-Amiri. An interactive approach for designing a robust disaster relief logistics network with perishable commodities. *Computers & Industrial Engineering*, 94:201–215, 2016.
[61] I. N. Robertson, H. R. Riggs, S. C. Yim, and Y. L. Young. Lessons from Hurricane Katrina storm surge on bridges and buildings. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 133(6):463–483, 2007.
[62] P. Schütz, L. Stougie, and A. Tomasgard. Stochastic facility location with general long-run costs and convex short-run costs. *Computers and Operations Research*, 35(9):2988–3000, 2008.
[63] T. N. Sear. Logistics planning in the downstream oil industry. *Journal of the Operational Research Society*, 44(1):9–17, 1993.
[64] W. C. Skamarock, J. B. Klemp, J. Dudhia, D. O. Gill, D. M. Barker, W. Wang, and J. G. Powers. A description of the advanced research wrf version 2. Technical Report NCAR/TN-468+ STR, National Center For Atmospheric Research, 2005.
[65] Y. Suzuki. Disaster-relief logistics with limited fuel supply. *Journal of Business Logistics*, 33(2):145–157, 2012.
[66] The United States Geological Survey. USGS flood event viewer: providing hurricane and flood response data, September 2017. URL https://stn.wim.usgs.gov/fev/#IrmaSeptember2017.
[67] P. J. Vickery, J. Lin, P. F. Skerlj, L. A. Twisdale, and K. Huang. HAZUS-MH hurricane model methodology. I: hurricane hazard, terrain, and wind load modeling. *Natural Hazards Review*, 7(2):82–93, 2006.
[68] J. Wang, Y. Hong, L. Li, J. J. Gourley, S. Khan, K. K. Yilmaz, R. F. Adler, F. S. Policelli, S. Habib, D. Irwn, A. S. Limaye, T. Korme, and L. Okello. The coupled routing and excess storage (CREST) distributed hydrological model. *Hydrological Sciences Journal*, 56(1): 84–98, 2011.
[69] M. J. Widener and M. W. Horner. A hierarchical approach to modeling hurricane disaster relief goods distribution. *Journal of Transport Geography*, 19(4):821–828, 2011.
[70] K. Yagci Sokat, R. Zhou, I. S. Dolinskaya, K. Smilowitz, and J. Chan. Capturing real-time

data in disaster response logistics. *Journal of Operations and Supply Chain Management*, 9(1):23–54, 2016.

[71] K. Yagci Sokat, I. S. Dolinskaya, K. Smilowitz, and R. Bank. Incomplete information imputation in limited data environments with application to disaster response. *European Journal of Operational Research*, 269(2):466–485, 2018.

[72] H. Yang and H. Nagarajan. Optimal power flow in distribution networks under stochastic N-1 disruptions. In *2020 Power Systems Computation Conference (PSCC) Proceedings*, 2020.

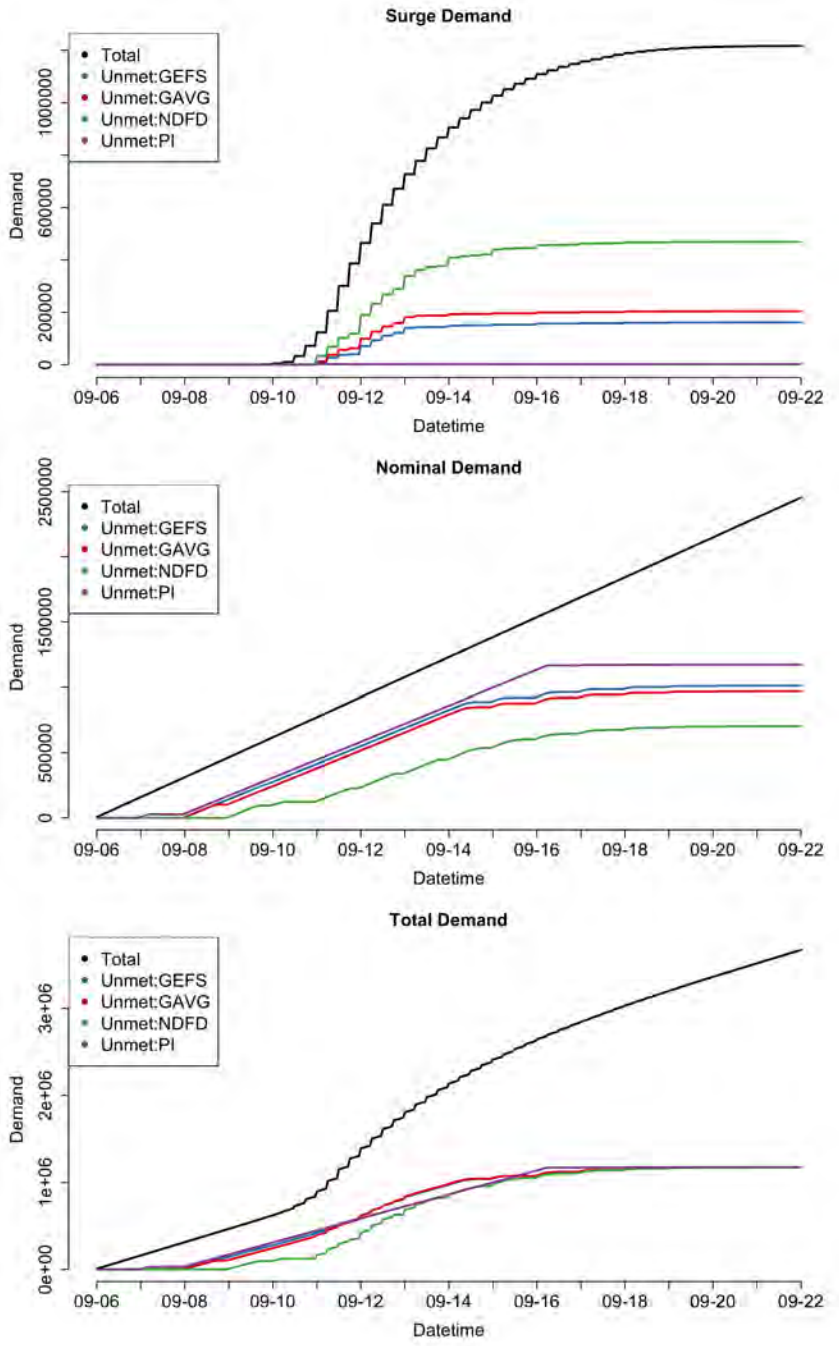

Figure 6. Surge (top), nominal (middle), and total (bottom) demands are shown in black. Unmet demand under the stochastic program (GEFS), two deterministic models (GAVG and NDFD) as well as under perfect information (PI) are also shown. All values are in barrels of diesel fuel for parameters $\alpha = 0.5$ and $H = 96, F = 72, R = 24$ and $N = 24$ hours.

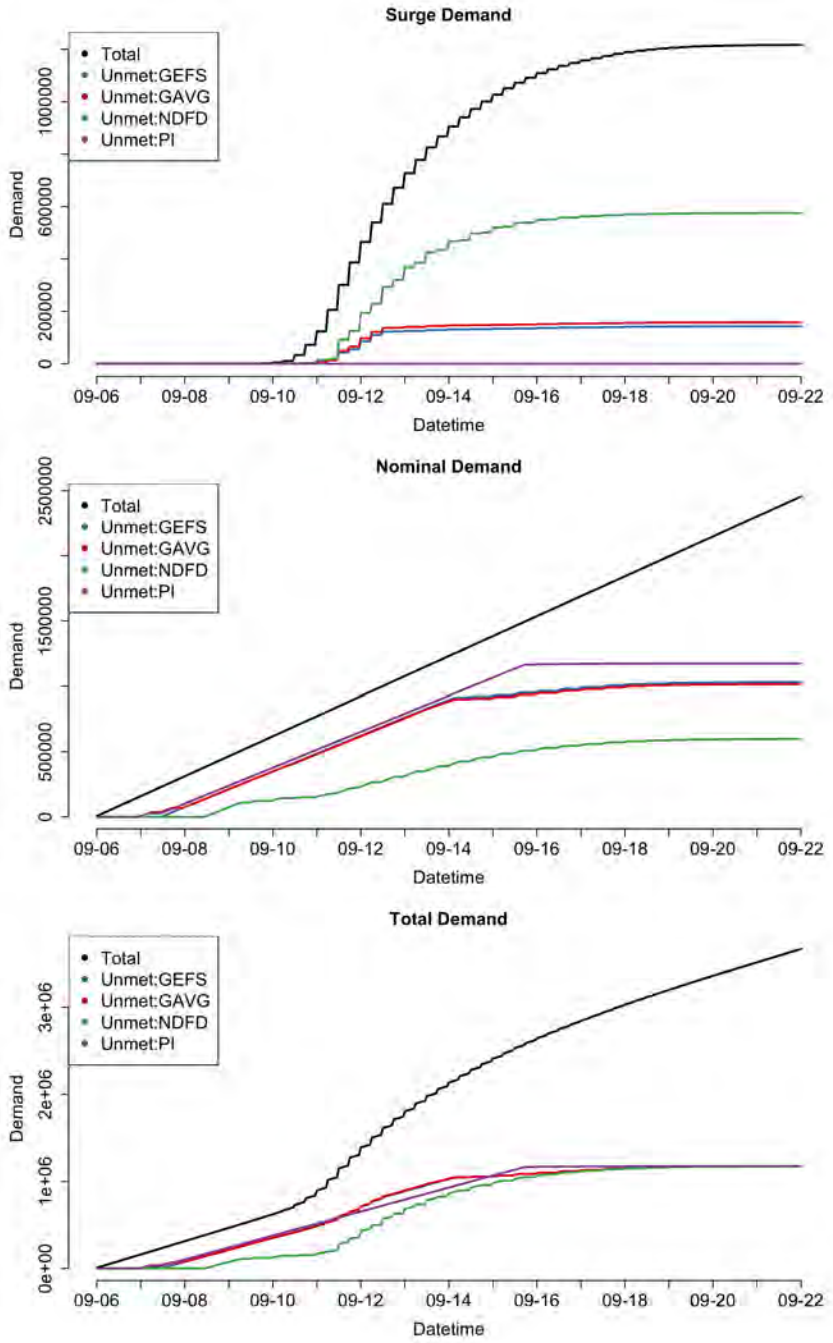

Figure 7. This figure follows the same format as Figure 6 except that Algorithm 1 now uses parameters $\alpha = 0.5$ and $H = 96, F = 72, R = 12$ and $N = 12$ hours. Decreasing R and N tends to increase the nimbleness of the strategy.