

Multistage Stochastic Programming with a Random Number of Stages: Applications in Hurricane Disaster Relief Logistics Planning

Murwan Siddig

*Deutsche Post Chair - Optimization of Distribution Networks, RWTH Aachen University,
siddig@dpo.rwth-aachen.de*

Yongjia Song*

Department of Industrial Engineering, Clemson University, yongjis@clemson.edu

** Denotes corresponding author*

Abstract

We consider a logistics planning problem of prepositioning relief commodities in preparation for an impending hurricane landfall. We model the problem as a multi-period network flow problem where the objective is to minimize the total expected logistics cost of operating the network to meet the demand for relief commodities. We assume that the hurricane's attributes evolve over time according to a Markov chain model, and the demand quantity at each demand point is calculated based on the hurricane's attributes (intensity and location) at the terminal stage, which corresponds to the hurricane's landfall. We introduce a fully adaptive multi-stage stochastic programming (MSP) model that allows the decision-maker to adapt their logistics decisions over time according to the evolution of the hurricane's attributes. In addition, we develop a novel extension of the standard MSP model to address the challenge of having a random number of stages in the planning horizon due to the uncertain landfall time of the hurricane. We benchmark the performance of the adaptive decision policy given by the MSP models with alternative decision policies, including a static policy, a rolling-horizon policy, a wait-and-see policy, and a decision-tree-based policy, all based on two-stage stochastic programming models. Our numerical results and sensitivity analyses provide key insights into the value of MSP in the hurricane disaster relief logistics planning problem.

Keywords: OR in disaster relief, stochastic programming, hurricanes, logistics planning

1. Introduction

According to the National Oceanic and Atmospheric Administration and National Centers for Environmental Information, the frequency and costs of weather and climate disasters have increased in the United States (US) over the last few decades. As of this writing, the US sustained 338 weather and climate disasters since 1980. Altogether, these disasters led to significant loss of human life and damages to property and infrastructure that exceeded 2.295 trillion dollars. Hurricanes, also known generically as tropical cyclones, are the most destructive and costly among these disaster events.

The distribution of the expenses is dominated by hurricane losses, which total around 1,194.4 billion dollars worth of damage. Hurricanes are also responsible for the highest number of deaths, which totals a staggering 6,864 fatalities (NOAA NCEI, 2022).

Despite their significant impacts, hurricanes can be detected a few days before they make landfall. Indeed, starting from approximately five days before a hurricane’s predicted landfall time, the National Hurricane Center (NHC) of the US provides forecast advisories that contain information about the hurricane’s attributes, including the predicted trajectory, forward speed, intensity, and potentially affected areas. Nevertheless, having accurate forecasts about the hurricane’s attributes does not guarantee a successful disaster relief effort. For instance, notwithstanding the unprecedented number of relief commodities prepositioned in preparation for hurricane *Katrina*, the Federal Emergency Management Agency (FEMA) was heavily criticized for failing to prepare enough relief commodities (Menzel, 2006), after the NHC provided accurate predictions about 48 hours before the hurricane hit (Knabb et al., 2005). Indeed, logistics planning plays a central role in disaster relief efforts since it directly deals with the procurement, allocation, and distribution of critical supplies and services to the affected areas (Ergun et al., 2010; Murray-Tuite and Wolshon, 2013; Pérez-Rodríguez and Holguín-Veras, 2016), such as first-aid kits, food, water, etc. The aforementioned assertions and the dynamic evolution of the hurricanes’ attributes, emphasize the need for a proactive, yet flexible and reliable logistics plan.

We consider a typical hurricane disaster relief logistics planning (HDRLP) problem where a decision maker (e.g., emergency management officials in FEMA) procures and prepositions relief commodities in anticipation of an impending hurricane. Based on the hurricane forecast, the decision maker can identify a study region that is at risk of being struck by the hurricane. Within the study region, there is a set of demand points (DPs) that represent the locations from which demand for relief commodities is expected. These relief commodities can be procured from a main distribution center (MDC), and the procured items can be prepositioned at a set of supply points (SPs). Once the hurricane makes landfall, the prepositioned relief commodities can be shipped from the SPs to the DPs to serve the demand. We assume that the demand level at each DP is based on two of the hurricane’s attributes at the time of its landfall: intensity and location. If these attributes and the time of landfall were known in advance, this problem could be modeled as a deterministic multi-period network flow problem defined over a logistics network. However, this is not the case due to the dynamic evolution of the hurricane’s attributes, which we model as a stochastic process characterized by a Markov chain (MC).

To obtain a dynamic decision policy that is adaptive to the evolving hurricane’s attributes, we develop a multi-stage stochastic programming (MSP) model for the HDRLP problem. MSP is a class of mathematical programming models for sequential decision-making under uncertainty, with the goal of constructing an optimal policy over a planning horizon consisting of multiple decision stages. With an optimal policy at hand, at each decision stage, the decision maker observes the current state of the system and implements an optimal decision prescribed by the policy. Standard MSP models assume a fixed number of stages in the planning horizon. In the HDRLP problem, however, the last stage in the planning horizon, which is referred to as the *terminal stage*, is signified by the hurricane’s landfall. In reality, this terminal stage is a random event that depends on several factors, such as the hurricane’s forward speed, intensity, and trajectory. For instance, a rapid surge

in the hurricane’s forward speed could cause the hurricane to arrive much earlier than anticipated. We show that this challenge can be addressed by an MSP model where the underlying stochastic process follows an MC with absorbing states, which correspond to the hurricane’s landfall locations. Our proposed model is inspired by the work of Guigues (2021), who developed an MSP model with a random number of stages for the case where the underlying stochastic process is assumed to be stage-wise independent. Instead, we consider a stage-wise dependent stochastic process (the MC), which characterizes the evolution of the hurricane’s attributes.

We compare the optimal policy obtained from the MSP model with several alternative solution approaches based on two-stage stochastic programming (2SSP) models that are well-studied in the literature. These alternative approaches produce approximate decision policies for the HDRLP problem, which include a static policy, a rolling-horizon policy, a wait-and-see policy, and a decision-tree-based policy. To that end, we conduct extensive numerical experiments that lead to key insights into the value of MSP in HDRLP. In sum, the contributions of our paper are as follows:

- First, to the best of our knowledge, our paper is the first that apply a fully adaptive MSP model to the HDRLP problem under demand uncertainty characterized by the hurricane’s evolving attributes. In addition to the MSP model, we also study several alternative solution approaches based on two-stage stochastic programming models for the HDRLP problem, including a novel decision-tree-based approach that achieves partial adaptability.
- Second, we develop a novel extension to MSP modeling for situations where the underlying stochastic process is modeled as an MC and the number of stages in the planning horizon is random, which addresses the unique challenge arising in the HDRLP problem with a random terminal stage due to the hurricane’s landfall.
- Finally, through an extensive numerical experiment based on both synthetic instances and case study instances inspired by Hurricane Florence (in 2018), we provide key managerial insights into the value of the adaptive decision policy obtained from the MSP model compared to those derived from alternative approaches based on two-stage stochastic programming in the HDRLP.

The remainder of this paper is organized as follows. In Section 2, we review the literature on applying optimization models to disaster relief logistics problems. In Section 3, we describe the problem setting. In Section 4, we describe the two stochastic programming models for the HDRLP problem, namely the 2SSP model and the MSP model. In Section 5, we discuss how these stochastic programming models can be used to constructed to design decision policies for the HDRLP problem. We discuss our numerical experiment results in Section 6 and conclude with final remarks in Section 7.

2. Literature Review

There is a growing stream of literature on applying optimization models to address different aspects of disaster relief logistics planning. A significant part of this literature is related to the last-mile delivery problems (Balci et al., 2008), which focus on the scheduling and delivery of relief commodities to the affected individuals. Other aspects of the disaster relief effort have also been studied, including the social cost of unmet demand (Holguín-Veras et al., 2013; Pérez-Rodríguez and Holguín-Veras, 2016;

Rawls and Turnquist, 2010; Yi and Özdamar, 2007), evacuation and rescue missions (Murray-Tuite and Wolshon, 2013), disruptions in the power and supply chain systems (Li et al., 2017; Yang et al., 2022), etc. Our work focuses mainly on the interplay between the dynamic evolution of the hurricane attributes and the logistics operational decisions made under uncertainty. In what follows, we give a brief overview of the literature most relevant to these two aspects.

2.1. Hurricane Evolution

As mentioned in Section 1, the NHC provides data support to aid with disaster preparation and response. When a tropical cyclone is likely to form in the Atlantic or eastern North Pacific Ocean, the NHC issues advisory notices, at least every six hours, starting approximately five days before the anticipated time of landfall. Given any of the forecast advisories, the forecast cone indicates a set of vulnerable locations from which demand for relief commodities is expected. These locations are thereby considered as the DPs. The demand quantities expected from these locations depend on the intensity level of the hurricane, as classified according to the Saffir-Simpson hurricane wind scale (Schott et al., 2012). Figure E.29 in the Appendix shows an example of the forecast cone produced by the NHC for Hurricane Florence (in 2018). It is important to note that these forecast advisories may be subject to significant forecast errors, especially the early ones. Analyzing and constructing statistical models that capture the multi-layered uncertainty about the hurricane evolution can be deployed in providing proactive decision policies (Yang et al., 2022). However, the vast majority of the literature that incorporates the dynamic uncertainty about the hurricane evolution into a decision model, focuses on hurricane evacuation (Davidson et al., 2020; Regnier, 2008; Regnier and Harr, 2006; Regnier, 2020; Yang et al., 2017), which is typically a one-time, irreversible decision. Instead, we focus on a fully adaptive disaster relief supply allocation policy, concurrently with the dynamic evolution of the hurricane attributes.

2.2. Solving HDRLP by Optimization under Uncertainty

A significant part of the literature focuses on decision-making under uncertainty at the strategic level using two-stage stochastic optimization models and robust optimization models. In most of the 2SSP models (e.g., Alem et al., 2016; Duran et al., 2011; Morrice et al., 2016; Paul and Zhang, 2019; Rawls and Turnquist, 2010; Salmerón and Apte, 2010; Velasquez et al., 2020), the first-stage decisions are typically defined as preparatory actions before the hurricane makes landfall, e.g., deciding on the locations of facilities for stockpiling relief commodities, delivery fleet assignment, and inventory levels. In the second stage, the decisions are typically defined to model recourse actions, after observing the realized hurricane attributes, e.g., paying penalties for unmet demand or salvaging unused items. One main concern with these 2SSP models is that it does not adequately address the dynamic nature of the evolving hurricane attributes. On the one hand, prepositioning the relief commodities too early in advance might become ineffective in situations where the hurricane’s realized path diverges from its predicted trajectory. On the other hand, postponing the prepositioning decisions until the forecast about the hurricane’s attributes becomes sufficiently accurate, might be costly or even infeasible due to the limited time that remains to accomplish these tasks. This dilemma emphasizes the need for a policy that can be adapted to the dynamic evolution of the hurricane, especially when

the landfall time of the hurricane is uncertain. A recent survey by Sabbaghtorkan et al. (2020) has identified the question of “when to start prepositioning” as one of the major research gaps for the HDRLP problem.

In the literature, there have been earlier attempts to address this issue. For example, Lodree Jr and Taskin (2009); Pacheco and Batta (2016) view the process of making the pre-disaster planning decisions as having two layers: (i) when to start the prepositioning process; and (ii) how many relief commodities should be prepositioned at different DPs once the start time of prepositioning has been determined. From this perspective, the goal of the decision maker is to identify the “best” period that strikes a balance between maximizing the level of confidence in the predicted hurricane attributes (at the termination time) and minimizing the logistics cost. Lodree Jr and Taskin (2009) approach this problem as an optimal stopping problem, and Pacheco and Batta (2016) use a combination of decision theory and stochastic programming techniques. Similarly, we show that the timing of implementing a prepositioning decision that is adaptive to the hurricane’s dynamic evolution can be optimized using a decision-tree model (see Section 5). However, the resulting policy is not as competitive as the fully adaptive decision policy offered by the MSP model due to its lack of adaptability once a prepositioning plan is implemented.

An alternative way to achieve an adaptive decision policy is to apply the so-called rolling-horizon approach where a look-ahead model is constructed and solved in every period in the planning horizon, but only the decisions pertaining to the current period are implemented. This heuristic approach naturally gives rise to an adaptive decision policy as the look-ahead model is constructed based on the most up-to-date forecast information in each period (Alden and Smith, 1992; Chand et al., 2002; Sethi and Sorger, 1991). In the context of disaster relief logistics, 2SSP models have been a popular choice for the look-ahead model in the rolling-horizon approach (Chang et al., 2022; Pacheco and Batta, 2016; Rivera-Royero et al., 2016; Yang et al., 2022). In our numerical study, we show that the rolling-horizon policy is not as competitive as the fully adaptive decision policy offered by the MSP model since the dynamic programming equations are not enforced.

Another stream of literature considers an adjustable robust optimization where a specific form of linear decision rules is imposed to derive adaptive decision policies for disaster relief logistics planning under uncertainty (e.g., Ben-Tal et al., 2011). Although this approach is appealing when the probability distribution for the underlying uncertain data is difficult to accurately identify, the imposed decision structure may undermine the resulting solution quality despite of the computational advantage that it brings. The MSP model proposed in our paper does not impose any decision rule structures and thus offers full-scale adaptability in the decision policy.

3. Problem Description and Formulation

3.1. Problem description

We consider an HDRLP problem that is modeled as a *multi-period network flow* problem defined on a directed network $G = (V, A)$. The set of nodes $V = \{0\} \cup I \cup J$ consists of the MDC $\{0\}$, the SPs I , and the DPs J . The planning horizon is given by $\{1, 2, \dots, T\}$, where time T corresponds to the terminal stage of the problem. At each period $t \in \{1, 2, \dots, T - 1\}$, the decision-maker has the

opportunity to procure relief commodities from the MDC and preposition them at different SPs. Relief commodities are generic items that represent any kind of demand that may arise at the DPs. Prepositioned items at an SP $i \in I$ can also be rerouted to any other SP $i' \in I$. At the terminal stage, the demand is realized for each DP $j \in J$, and to satisfy the demand, the prepositioned relief commodities can be shipped from the set $I(j)$ of eligible SPs for this DP j . The objective of the HDRLP problem is to minimize the expected total cost, which consists of the procurement and transportation costs over the planning horizon, the penalty cost for unsatisfied demand, and the salvage value for unused relief commodities.

We make the following assumptions to simplify the problem formulation.

Assumption 1. *All shipments made at the start of a period t will arrive at their destinations by the end of period t , or equivalently, the start of period $t + 1$.*

Assumption 2. *All SPs are assumed to be open starting from period $t = 1$.*

Assumption 3. *The demand for relief commodities is only incurred at the hurricane's landfall time T .*

Assumption 1 is made to avoid having to model the flow over a time-expanded graph that complicates the underlying optimization model significantly. Assumptions 2 and 3 are made to avoid introducing discrete decision variables (to model when and where the SPs should be activated/deactivated) and a nonlinear term in the objective (to capture the deprivation cost for delayed demand satisfaction). These would result in a multi-stage stochastic *mixed-integer nonlinear* programming problem, which is computationally challenging to solve. These simplifications allow us to focus on designing prepositioning policies that adapt to the evolving hurricane characteristics.

3.2. Deterministic Formulation

We start by presenting a deterministic version of the HDRLP problem for the following reasons. First, it allows us to introduce the logistics model in a setting that is easier to understand and makes it easy to transition to the different stochastic optimization models. Second, the solution to the deterministic version can be viewed from the perspective of a clairvoyant who has perfect information about the evolution of the hurricane and the demand arising from it. This serves as a (lower-bound) benchmark for all the presented approaches in our numerical experiments. The notations for the parameters and decision variables used in the deterministic formulation are summarized in Table 1.

The deterministic formulation of the HDRLP problem is defined as follows.

Parameter	Description
$c_{ii',t}^b$	unit shipping cost from SP/MDC $i \in \{0\} \cup I$ to SP $i' \in I$, at $t = 1, 2, \dots, T-1$.
$c_{i,t}^h$	unit holding cost at SP $i \in I$, at $t = 1, 2, \dots, T-1$.
h_t	unit procurement cost from the MDC at t , at $t = 1, 2, \dots, T-1$.
c_{ij}^a	unit transportation cost from SP $i \in I$ to DP $j \in J$.
$I(j)$	subset of SPs eligible to send relief items to DP $j \in J$.
p	unit penalty cost for unsatisfied demand.
q	unit salvage value for unused relief items.
d_j	demand for relief items in DP $j \in J$.
$\bar{x}_i$	inventory capacity of SP $i \in I$.
$x_{i,0}$	initial inventory level of SP $i \in I$.
$\bar{f}_{ii'}$	flow capacity from SP/MDC $i \in \{0\} \cup I$ to SP $i' \in I$.
$\bar{y}_{ij}$	flow capacity from SP $i \in I(j)$ to DP $j \in J$.
Decision variable	
$x_t = (x_{i,t})_{i \in I}$	inventory at SP $i \in I$, at the end of period $t = 1, 2, \dots, T-1$.
$f_t = (f_{ii',t})_{i \in \{0\} \cup I, i' \in I}$	items shipped from node $i \in \{0\} \cup I$ to $i' \in I$, at $t = 1, 2, \dots, T-1$.
$y = (y_{ij})_{i \in I, j \in J}$	items shipped from SP $i \in I$ to DP $j \in J$.
$\underline{e} = (\underline{e}_j)_{j \in J}$	unmet demand at DP $j \in J$.
$\bar{e} = (\bar{e}_i)_{i \in I}$	unused inventory prepositioned at SP $i \in I$.

Table 1: Description of the parameters and decision variables in the HDRLP problem.

$$\begin{aligned}
\min \quad & \sum_{t=1}^{T-1} \left(h_t \sum_{i \in I} f_{0i,t} + \sum_{i \in I \cup \{0\}} \sum_{i' \in I} c_{ii',t}^b f_{ii',t} + \sum_{i \in I} c_{i,t}^h x_{i,t} \right) + \sum_{j \in J} \sum_{i \in I(j)} c_{ij}^a y_{ij} + p \sum_{j \in J} \underline{e}_j + q \sum_{i \in I} \bar{e}_i \\
\text{s.t.} \quad & x_{i,t} = x_{i,t-1} + \sum_{i' \in I \cup \{0\}, i' \neq i} f_{i'i,t} - \sum_{i' \in I, i' \neq i} f_{ii',t}, \quad \forall i \in I, t = 1, 2, \dots, T-1, \quad (1a) \\
& x_{i,t} \leq \bar{x}_i, \quad \forall i \in I, t = 1, 2, \dots, T-1, \quad (1b) \\
& \sum_{i' \in I, i' \neq i} f_{ii',t} \leq x_{i,t-1}, \quad \forall i \in I, t = 1, 2, \dots, T-1, \quad (1c) \\
& f_{ii',t} \leq \bar{f}_{ii'}, \quad \forall i \in \{0\} \cup I, i' \in I, t = 1, 2, \dots, T-1, \quad (1d) \\
& y_{ij} \leq \bar{y}_{ij}, \quad \forall j \in J, i \in I(j) \quad (1e) \\
& \sum_{j \in J: i \in I(j)} y_{ij} + \bar{e}_i = x_{i,T}, \quad \forall i \in I, \quad (1f) \\
& \sum_{i \in I(j)} y_{ij} + \underline{e}_j = d_j, \quad \forall j \in J, \quad (1g) \\
& x_t, f_t \geq 0 \quad \forall t = 1, 2, \dots, T-1 \quad (1h) \\
& y, \bar{e}, \underline{e} \geq 0 \quad (1i)
\end{aligned}$$

The objective function of the deterministic formulation (1) minimizes the procurement cost, transportation cost from the MDC to SPs and between the SPs, the holding cost at the SPs, the transportation cost from the SPs to the DPs, the penalty cost for unsatisfied demand, and the salvage value for unused relief commodities. Constraint (1a) represents the flow balance constraint at each SP $i \in I$. Constraint (1b) ensures that the inventory capacity at each SP $i \in I$ is not

exceeded. Constraint (1c) ensures that the total amount of items shipped from SP $i \in I$ cannot exceed the inventory level at time $t - 1$. Constraints (1d) and (1e) represent the flow capacity constraint from the MDC/SPs to the SPs and from the SPs to the DPs, respectively. Constraint (1f) and (1g) quantify the amount of unused inventory prepositioned at each SP $i \in I$ and the amount of unsatisfied demand at each DP $j \in J$, respectively. Constraint (1h) and (1i) are non-negativity constraints.

3.3. Estimating the demand and modelling the hurricane evolution

Estimating the demand for relief commodities is one of the most challenging and understudied topics in humanitarian logistics research (Çelik et al., 2012; Holguín-Veras et al., 2012; Holguín-Veras and Jaller, 2012). Studies have shown that the demand for relief commodities depends on several factors including the hurricane’s intensity (Baker, 1991) and landfall location (Kirac and Milburn, 2018; Morrice et al., 2016), which can change dynamically over time. Let $\xi_t^{\text{int}} \in \Xi^{\text{int}}$ denote the hurricane intensity and $\xi_t^{\text{loc}} \in \Xi^{\text{loc}}$ denote the hurricane location at time t . We measure the hurricane intensity using the Saffir-Simpson hurricane wind scale (SSHWS) (Schott et al., 2012). The SSHWS is used by the NHC to estimate the potential damage and flooding from a hurricane. The SSHWS classifies hurricanes into five categories based on the maximum sustained wind speed with Category 1 being the least severe and Category 5 being the most severe (National Hurricane Center and Central Pacific Hurricane Center, 2022), and we define $\Xi^{\text{int}} := \{0, 1, 2, 3, 4, 5\}$, where states 1 to 5 correspond to these five categories and state 0 corresponds to the case when the wind speed is less than Category 1. Furthermore, the hurricane’s location is specified using the latitude and longitude coordinates of its center (Kirac and Milburn, 2018). We discretize the continuous space of the coordinates of the hurricane’s center ξ_t^{loc} into a finite set of possible locations Ξ^{loc} as illustrated in Figure E.27 in Appendix E.

The hurricane’s intensity and location depend on several random factors such as the hurricane’s forward speed, trajectory, atmospheric conditions and geographical features (Brettschneider, 2008; Bosart et al., 2000). We follow Pacheco and Batta (2016); Regnier and Harr (2006); Regnier (2008); Zhang et al. (2022) and model the evolution of the hurricane’s attributes $\xi_t = (\xi_t^{\text{int}}, \xi_t^{\text{loc}})$ as a stochastic process $\{\xi_t\}_{t=1}^T$ that follows a Markov chain.

Assumption 4. *The stochastic process $\{\xi_t\}_{t=1}^T$ follows a finite discrete-time MC where $\mathbb{P}(\xi_t \mid \xi_1, \xi_2, \dots, \xi_{t-1}) = \mathbb{P}(\xi_t \mid \xi_{t-1}), \forall t = 2, 3, \dots, T$.*

We let $\Xi = \Xi^{\text{int}} \times \Xi^{\text{loc}}$ be the set of states associated with this MC, and denote the one-step transition probability matrix of this MC as $\mathbf{P}$. $\mathbf{P}$ is a $|\Xi|$ -by- $|\Xi|$ matrix where each row/column index corresponds to an MC state $\xi \in \Xi$. For notational convenience, we refer to the row/column indices and the corresponding MC states interchangeably. Table 2 summarizes the notation for the MC model used to characterize the evolution of the hurricane’s attributes.

Given an initial state ξ_1 at $t = 1$, the terminal stage T is signified by the hurricane reaching a landfall location, which can be a random variable depending on the MC model. The hurricane’s landfall location is a location where the hurricane’s center may cross the coastline as shown in Figure E.27 in Appendix E. Given the set of landfall locations $L \subset \Xi^{\text{loc}}$, we partition the set $\Xi = \Xi^{\text{int}} \times \Xi^{\text{loc}}$ into a set of absorbing states $\Xi^{\text{absor}} = \{(\xi_t^{\text{int}}, \xi_t^{\text{loc}}) \in \Xi \mid \xi_t^{\text{loc}} \in L\}$, where

Notation	Description
ξ_t^{int}	hurricane intensity at time t measured by the SSHWS.
ξ_t^{loc}	hurricane location at time t measured by latitude and longitude coordinates.
Ξ^{int}	set of all possible hurricane intensity states.
Ξ^{loc}	set of all possible hurricane location states.
Ξ	set of all MC states.
L	set of all possible hurricane landfall locations.
$\mathbf{P}$	one-step transition probability matrix of the MC.
$\mathbf{P}^{(n)}$	n -step transition probability matrix of the MC.
Ξ^{absor}	set of all possible absorbing states.
Ξ^{trans}	set of all possible transient states.
T_{max}	maximum number of periods until the system reaches an absorbing state, starting from the initial MC state ξ_1 .

Table 2: Description of the notation used to model the hurricane evolution.

$\mathbf{P}_{ii} = 1, \forall i \in \Xi^{\text{absor}}$, and a set of transient states $\Xi^{\text{trans}} = \Xi \setminus \Xi^{\text{absor}}$. Therefore, given an initial state ξ_1 at $t = 1$, the maximum possible number of periods T_{max} for the MC to reach any absorbing state can be computed as:

$$T_{\text{max}} := \max \left\{ n \in \mathbb{Z}_+ \mid \exists \xi \in \Xi^{\text{trans}} \text{ such that } \mathbf{P}_{\xi_1, \xi}^{(n)} > 0 \right\} + 1, \quad (2)$$

where $\mathbf{P}^{(n)}$ denotes the n -step transition probability matrix for the MC. We note that T_{max} is deterministically derived given an initial state ξ_1 , while the terminal stage T is a random variable. Figure 1 shows a symbolic representation of how reaching an absorbing state models the random number of periods.




Figure 1: A symbolic representation of how to model the random number of periods until the system reaches an absorbing state. The example shows a transient state ξ_1^1 at $t = 1$, two transient states $\{\xi_t^1, \xi_t^2\}$ and one absorbing state ξ_t^3 for each $t = 2, 3, \dots, T_{\text{max}} - 1$, and a single absorbing state $\xi_{T_{\text{max}}}^3$ for $t = T_{\text{max}}$. An arrow from ξ_t^i to ξ_{t+1}^j indicates that the system can transition from state ξ_t^i to state ξ_{t+1}^j with a positive probability. The arrows from ξ_t^i to itself indicate that the system remains in the absorbing state and the terminal period T is reached. We see that T is a random variable with $2, 3, \dots, T_{\text{max}}$ being its possible realizations in this example.

We are now ready to make the following assumption on the demand for relief commodities, which is based on the hurricane’s intensity and landfall location.

Assumption 5. *The demand for relief commodities at a demand point $DP\ j \in J$ is given by function $d_j(\xi_T)$ that depends on the hurricane’s state ξ_T at its landfall time $t = T$, i.e., $\xi_T \in \Xi^{absor}$.*

4. Stochastic Programming Models for Multi-period HDRLP

In this section, we present two stochastic programming models for solving the multi-period HDRLP described in Section 3. We first describe a 2SSP model (Section 4.1), which is a static model that can only adapt once after all uncertain information is revealed, and then describe an MSP model (Section 4.2), which is a fully adaptive model that can adapt to new information as it becomes available.

To understand the difference between a “two-stage” and a “multi-stage” model, it is crucial to distinguish between the concepts of *stages* and *periods*. We use the term *stage* to refer to a point in time in which a decision can be made or adapted, and we use the term *period* to refer to a point in time in which a new realization of the uncertain parameters is revealed. In that sense, *stages* concern decision-making, while *periods* concern the availability of new information. More importantly, in a single stage, one can make decisions that span multiple periods, but one can only observe new information about the current period in each period. For a two-stage stochastic programming model for a multi-period problem specifically, to obtain an implementable policy, one would treat the state variables over the entire horizon as the first-stage decision variables, while the recourse decisions on each possible sample path are defined as second-stage decision variables. Once the first-stage decision is made, one would observe the realization of the random variables in each period sequentially and make the corresponding recourse decisions, although equivalently, this can be seen as observing the realizations of all period at once and make the corresponding recourse decisions in each period. This equivalence is based on the fact that the state variables, which are decided in the first stage, will not adapt to the sequential realization of the random variables in each period. We elaborate this further in the following subsection.

4.1. Two-Stage Stochastic Programming (2SSP) model

A 2SSP model has two stages: a first stage to plan for all future periods and a second stage to adapt after all uncertain information is revealed. In the context of the HDRLP, the first-stage problem decides the prepositioning of relief items for all periods in the planning horizon, and the second-stage problem decides the allocation of relief items to meet the demand after the hurricane makes landfall, pay the penalty for demand shortages, and salvage the unused items. To that end, the *first-stage* problem includes the flow variables $\mathbf{f} = (f_1, f_2, \dots, f_{T_{\max}-1})$ and inventory variables $\mathbf{x} = (x_1, x_2, \dots, x_{T_{\max}-1})$ over the maximum possible planning horizon, and the *second-stage* problem includes the variables $y, \bar{e}, \underline{e}$ that are used to allocate the relief items to meet the demand after the hurricane makes landfall.

Given an initial state $\xi_1 \in \Xi^{\text{trans}}$, we can enumerate all possible sample paths $\Xi_{[T]}$ in which the

stochastic process starts in state ξ_1 and ends in an absorbing state $\xi_T \in \Xi^{\text{absor}}$, i.e.,

$$\Xi_{[T]} := \{ \xi_{[T]} = (\xi_1, \xi_2, \dots, \xi_T) \mid \xi_t \in \Xi^{\text{trans}}, \forall t = 1, 2, \dots, T-1 \text{ and } \xi_T \in \Xi^{\text{absor}} \}.$$

Each sample path $\xi_{[T]} \in \Xi_{[T]}$ is associated with a probability $p_{\xi_{[T]}}$ that can be calculated from the one-step transition probability matrix $\mathbf{P}$. The *first-stage* problem is then defined as:

$$\min_{\mathbf{x}, \mathbf{f}} \sum_{\xi_{[T]} \in \Xi_{[T]}} p_{\xi_{[T]}} \cdot Q^{\xi_{[T]}}(\mathbf{x}, \mathbf{f}) \quad (3a)$$

$$\text{s.t. (1a), (1b), (1c), (1d) and (1h) with } T \text{ replaced by } T_{\max}, \quad (3b)$$

and the *second-stage* problem $Q^{\xi_{[T]}}(\mathbf{x}, \mathbf{f})$ is defined for each scenario $\xi_{[T]} \in \Xi_{[T]}$ as:

$$\min_{y, \bar{e}, \underline{e}} \sum_{t=1}^{T-1} \left(h_t \sum_{i \in I} f_{0i,t} + \sum_{i \in I_0} \sum_{i' \in I} c_{ii',t}^b f_{ii',t} + \sum_{i \in I} c_{i,t}^h x_{i,t} \right) + \sum_{j \in J} \sum_{i \in I(j)} c_{ij}^a y_{ij} + p \sum_{j \in J} \underline{e}_j + q \sum_{i \in I} \bar{e}_i \quad (4a)$$

$$\text{s.t. (1e), (1f), (1g) and (1i) with } d_j \text{ replaced by } d_j(\xi_T). \quad (4b)$$

The 2SSP model defined by eqs. (3) and (4) can be solved efficiently using standard techniques, such as Benders decomposition (Birge and Louveaux, 2011), but has the drawback of making static decisions for all future periods at $t = 1$, before observing any updated information about the hurricane, that can only be adapted, once, after the hurricane makes landfall. It is not difficult to see that this static approach to decision-making is likely to lead to suboptimal decisions that compromise the quality of the resulting policy.

4.2. Multi-stage Stochastic Programming (MSP) model

MSP models that have the same number of stages as periods are fully adaptive as they can adapt to new information in every period as it becomes available. A generic *nested* formulation for an MSP model is defined as follows.

$$\min_{a_1 \in \chi_1(a_0, \xi_1)} z_1(a_1, \xi_1) + \mathbb{E}_{|\xi_{[1]}} \left[\min_{a_2 \in \chi_2(a_1, \xi_2)} z_2(a_2, \xi_2) + \mathbb{E}_{|\xi_{[2]}} \left[\dots + \mathbb{E}_{|\xi_{[T-1]}} \left[\min_{a_T \in \chi_T(a_{T-1}, \xi_T)} z_T(a_T, \xi_T) \right] \right] \right], \quad (5)$$

where ξ_1 and a_0 are given as input, and $\xi_{[t]} := (\xi_1, \xi_2, \dots, \xi_t)$ denotes the history of the stochastic process up to time t . The goal of the decision-maker in MSP models is to find an optimal decision policy that gives an implementable action $a_t \in \chi_t(a_{t-1}, \xi_t)$ at every period $t = 1, 2, \dots, T$. The vector a_t can be split into two types of decision variables: state variables $u_t \in \mathbb{R}^{n_t}$ and local (control) variables $v_t \in \mathbb{R}^{m_t}$, i.e., $a_t := (u_t, v_t)$. The set of state variables u_t is what links different stages together, and the set of local variables v_t participates exclusively in each period t . When $t < T$, the decision-maker aims to find an optimal action a_t which minimizes the immediate cost given by $z_t(a_t, \xi_t)$ plus the expected future cost given by the nested expectation $\mathbb{E}_{|\xi_t}[\cdot]$. When $t = T$, there is no future cost to be incurred, and the goal is to minimize the immediate cost $z_T(a_T, \xi_T)$ only.

In the context of the HDRLP with random landfall time, we replace T with $T_{\max}$ and define the decision vector a_t as

$$a_t := \begin{cases} (x_t, f_t) & \text{if } \xi_t \in \Xi^{\text{trans}}, \\ (y, \bar{e}, \underline{e}) & \text{if } \xi_t \in \Xi^{\text{absor}}, \end{cases} \quad (6)$$

where the state variable $u_t := x_t$ and the local variable $v_t := f_t$ for transient states $\xi_t \in \Xi^{\text{trans}}$, and $v_t := (y, \bar{e}, \underline{e})$ for absorbing states $\xi_t \in \Xi^{\text{absor}}$. Furthermore, we define the immediate cost function $z_t(a_t, \xi_t)$ as

$$z_t(a_t, \xi_t) := \begin{cases} h_t \sum_{i \in I} f_{0i,t} + \sum_{i \in I_0} \sum_{i' \in I} c_{ii',t}^b f_{ii',t} + \sum_{i \in I} c_{i,t}^h x_{i,t} & \text{if } \xi_t \in \Xi^{\text{trans}}, \\ \sum_{j \in J} \sum_{i \in I(j)} c_{ij}^a y_{ij} + p \sum_{j \in J} \underline{e}_j + q \sum_{i \in I} \bar{e}_i, & \text{if } \xi_t \in \Xi^{\text{absor}}, \end{cases} \quad (7)$$

and the feasible set $\chi_t(a_{t-1}, \xi_t)$ as

$$\chi_t(a_{t-1}, \xi_t) := \begin{cases} \text{Constraints eqs. (1a) to (1d) and (1h) associated with period } t & \text{if } \xi_t \in \Xi^{\text{trans}}, \\ (1e), (1f), (1g), \text{ and } (1i) \text{ with } d_j \text{ replaced by } d_j(\xi_t) & \text{if } \xi_t \in \Xi^{\text{absor}}. \end{cases} \quad (8)$$

Dynamic programming formulation: Under Assumption 4, the conditional expectation $\mathbb{E}_{|\xi_t}[\cdot]$ in (5) can be replaced by $\mathbb{E}_{|\xi_t}[\cdot]$ thanks to the Markovian property, and the MSP model can be solved using the following dynamic programming formulation (Bellman, 1966):

$$Q_t(a_{t-1}, \xi_t) := \min_{a_t} \left\{ z_t(a_t, \xi_t) + \mathfrak{Q}_{t+1}^{\xi_t}(a_t) \mid a_t \in \chi_t(a_{t-1}, \xi_t) \right\}, \quad \forall t = 1, 2, \dots, T_{\max}, \quad (9)$$

where $\mathfrak{Q}_{t+1}^{\xi_t}(a_t)$ is referred to as the expected *cost-to-go* function and is given by

$$\mathfrak{Q}_{t+1}^{\xi_t}(a_t) = \begin{cases} \mathbb{E}[Q_{t+1}(a_t, \xi_{t+1}) \mid \xi_t] & \text{if } \xi_t \in \Xi^{\text{trans}} \\ 0 & \text{if } \xi_t \in \Xi^{\text{absor}}. \end{cases} \quad (10)$$

Formulation (9) decomposes the nested structure of formulation (5) into $T_{\max}$ stages by recursively defining the expected cost-to-go function $\mathfrak{Q}_{t+1}^{\xi_t}(a_t)$ for each MC state ξ_t in period t . We note that the expected cost-to-go function $\mathfrak{Q}_{t+1}^{\xi_t}(a_t)$ models the random number of periods in the problem by treating any time that an absorbing state ξ_t is reached as the final stage of the problem, and it takes a constant value 0 in this case as no future cost will be incurred.

Nested Benders decomposition algorithm: The dynamic programming formulation (9) can be solved using the nested Benders' decomposition algorithm (Birge and Louveaux, 2011), which we modify to accommodate the random number of periods in the problem. The basic idea of the algorithm is to iteratively approximate the expected cost-to-go functions $\mathfrak{Q}_{t+1}^{\xi_t}(\cdot)$ by a collection of cutting-plane approximations that are obtained by iteratively executing the following two main steps.

1. *Forward* step: Simulate a trajectory (sample path) of MC states $\xi_{[T]}$ by going forward in time starting from $t = 1$ until time T such that $\xi_T \in \Xi^{\text{absor}}$. Along this trajectory, solve each non-terminal stage- t problem based on the current approximation of the expected cost-to-go function $\check{\mathbf{Q}}_{t+1}^{\xi_t}(\cdot)$, and carry the values of the state variables in the corresponding optimal solution $\check{a}_t$ (also known as trial points) over to the next stage.
2. *Backward* step: Go backward in the simulated trajectory $\xi_{[T]}$, starting from $t = T$ where $\xi_T \in \Xi^{\text{absor}}$ until $t = 1$, and place valid inequalities for the epigraph of $\mathbf{Q}_t^{\xi_{t-1}}(\cdot)$, which are constructed by the approximate subgradient of $\mathbf{Q}_t^{\xi_{t-1}}(\cdot)$ evaluated at the trial points collected. Specifically, given a sequence of trial points $(\check{a}_1, \check{a}_2, \dots, \check{a}_{T-1})$, for each stage $t = T, T-1, \dots, 2$, a linear underestimator for $\mathbf{Q}_t^{\xi_{t-1}}(\cdot)$ is given by $q_t(a_{t-1}) := \beta_t^\top a_{t-1} + \pi_t$, where

$$\beta_t \in \mathbb{E}[\partial \check{Q}_t(\check{a}_{t-1}, \xi_t) \mid \xi_{t-1}], \quad (11)$$

and

$$\pi_t := \begin{cases} \mathbb{E}[b_t^{\xi_t \top} \check{\lambda}_t] & \text{for } t = T \\ \mathbb{E}\left[b_t^{\xi_t \top} \check{\lambda}_t + \sum_{c \in C_{t+1}^{\xi_t}} \pi_{t+1,c} \check{\rho}_{t,c}\right] & \text{for } t = T-1, \dots, 2, \end{cases} \quad (12)$$

where $b_t^{\xi_t \top}$ is the right-hand side of the constraints in $\chi_t(\check{a}_{t-1}, \xi_t)$, $C_{t+1}^{\xi_t}$ is the collection of cutting-plane approximation so far to $\mathbf{Q}_{t+1}^{\xi_t}(\cdot)$, $\check{\rho}_{t,c}$ values are the optimal dual vectors associated with cutting-plane approximation for the epigraph of $\check{\mathbf{Q}}_{t+1}^{\xi_t}(\cdot)$, given by $\beta_{t+1,c}^\top a_t + \pi_{t+1,c} \leq \theta_{t+1}^{\xi_t}$, $\forall c \in C_{t+1}^{\xi_t}$, and $\check{\lambda}_t$ values are the optimal dual vectors associated with constraints defined in $\chi_t(\check{a}_{t-1}, \xi_t)$.

Algorithm 1 Nested Benders decomposition algorithm for an MSP model where the underlying stochastic process is an MC with transient states Ξ^{trans} and absorbing states Ξ^{absor} .

Input: $a_0, \xi_1, T_{\max}, \mathbf{P}$ and initial lower bounds $\underline{\theta}$ for the approximate expected cost-to-go functions.

STEP 0: (*Initialization*).

- For each $\xi_t \in \Xi^{\text{trans}}$ that is reachable from ξ_1 : set $C_{t+1}^{\xi_t} = \emptyset$, introduce variable $\theta_{t+1}^{\xi_t}$, and add constraint $\theta_{t+1}^{\xi_t} \geq \underline{\theta}$, for $t = 1, 2, \dots, T_{\max} - 1$.
- For each $\xi_t \in \Xi^{\text{absor}}$ that is reachable from ξ_1 : set $\mathbf{Q}_{t+1}^{\xi_t}(\cdot) = 0$, for $t = 2, 3, \dots, T_{\max}$.
- Initialize the iteration count $n = 0$, $t = 1$, $\xi_t^n = \xi_1$, and $\check{a}_{t-1}^n = a_0$.

STEP 1: (*Forward Pass*).

- Solve the stage- t problem $Q_t^n(\check{a}_{t-1}^n, \xi_t^n)$ to obtain a_t^* . Set $\check{a}_t^n = a_t^*$ and let $t \leftarrow t + 1$.
- If $\xi_t^n \in \Xi^{\text{trans}}$, sample a new state ξ_t^n given ξ_{t-1}^n according to the one-step transition probability matrix $\mathbf{P}$. Otherwise, go to **STEP 2**.

STEP 2: (*Termination check*).

- If the termination criterion is met, *STOP*; otherwise, go to **STEP 3**.

STEP 3: (*Backward Pass*).

- Starting from the terminal period T of the trajectory that was sampled in the forward pass, for every $\xi_t \in \Xi^{\text{trans}}$, solve the stage- t problem $Q_t^n(\check{a}_{t-1}^n, \xi_t)$.
- Generate a cutting-plane approximation to $\mathbf{Q}_t^{\xi_{t-1}}(\cdot)$ following the description in (11) and (12), add the cut to the collection of cuts $C_t^{\xi_{t-1}}$, and let $t \leftarrow t - 1$.
- If $t = 1$, let $n \leftarrow n + 1$, set $\check{a}_{t-1}^n = a_0$, $\xi_t^n = \xi_1$ and go to **STEP 1**.

Output: $\check{\mathbf{Q}}_{t+1}^{\xi_t}(\cdot)$ for every $\xi_t \in \Xi^{\text{trans}}$, $t = 1, 2, \dots, T_{\max} - 1$ that is reachable from ξ_1 .

Due to the polyhedral structure of the epigraph of $\mathfrak{Q}_{t+1}^{\xi_t}(\cdot)$, this iterative procedure builds a piecewise linear lower approximation to $\mathfrak{Q}_{t+1}^{\xi_t}(\cdot)$ that converges in a finite number of iterations with probability one (Shapiro, 2011). A pseudocode description of this modified nested Benders decomposition algorithm is provided in Algorithm 1. Note that the HDRLP problem has *relatively complete recourse* (Shapiro et al., 2009), which eliminates the need for feasibility cuts. Therefore, only the Benders optimality cuts are considered in Algorithm 1.

5. Stochastic Programming-based Policies

The stochastic multi-period HDRLP described in Section 3 is a sequential decision-making problem under uncertainty. Sequential decision-making problems are characterized by a repeating pattern of being in some situation, generally known as a state, and making decisions based on that situation. Each time we make a decision, we incur a cost, and every decision we make affects not only the immediate cost but also the next situation and, through that, all future costs. The goal is to find a policy that tells us what to do in each situation to optimize some performance criterion (e.g., minimize the expected cumulative cost).

There are many ways to define a policy using the stochastic programming models described in Section 4. The policies we discuss in this section differ based on their level of adaptability to new information. We consider three types of policies: (i) *static* policies; (ii) *partially adaptive* policies; and (iii) *fully adaptive* policies.

5.1. Static policies

A static policy can be obtained by solving the 2SSP model defined by eqs. (3) and (4) at time $t = 1$. Solving the 2SSP model provides an optimal first-stage decision $\tilde{\mathbf{f}} = (\tilde{f}_1, \tilde{f}_2, \dots, \tilde{f}_{T_{\max}-1})$ and $\tilde{\mathbf{x}} = (\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_{T_{\max}-1})$. Given a sample path $\xi_{[T]}$, decisions $\tilde{f}_t$ and $\tilde{x}_t$ are then implemented for every period $t = 1, 2, \dots, T - 1$ where $\xi_t \in \Xi^{\text{trans}}$, until reaching a period $t = T$ where $\xi_T \in \Xi^{\text{absor}}$ (i.e., the hurricane makes landfall). At the landfall time $t = T$, (4) is solved to obtain the total cost of applying the static policy on sample path $\xi_{[T]}$. In our numerical study in Section 6, we refer to this as the static 2SSP (s-2SSP) policy.

5.2. Partially adaptive policies

A significant drawback of the static 2SSP policy is its lack of adaptability to additional information that is expected to be collected during the execution of the static logistics plan, which is committed at $t = 1$. Such an early commitment will likely suffer from having to mitigate a wide range of hurricane scenarios due to lack of information. One natural idea to address this issue is to “wait” and only activate a (static) logistics plan until sufficient information (e.g., about the hurricane’s states at its landfall) is obtained. Once the activation decision is made, e.g., at time t , a static logistics plan is implemented thereafter by solving a 2SSP model (eqs. (3) and (4)) for the periods $t, t + 1, \dots, T_{\max}$ with the initial condition $x_{t-1} = x_0$ and the initial MC state set to be the MC state ξ_t at that time. This *wait-and-activate* type policy is partially adaptive: it is only adaptive

until a static logistics plan is activated (and implemented), once the plan is committed, it cannot adapt to how the remainder of the sample-path unfolds. The wait-and-activate decision structure is consistent with the practical operational constraints associated with these large-scale logistics operations, which have an “all or nothing” feature: once activated, a (static) logistics plan must be committed until the end of the process.

The activation decision $g(\xi_t) \in \{0, 1\}$ at each MC state ξ_t is made according to an activation policy, where $g(\xi_t) = 0$ and 1 indicate the wait and activate decisions, respectively. We consider two types of activation policies: a simple heuristic activation policy and an optimal adaptive activation policy obtained from a decision-tree approach. The key difference between the two is that the heuristic activation policy does not incorporate the expected cost of postponing the activation. Indeed, the delay in executing a logistics plan may have the costly consequence of insufficiently fulfilling the demand due to a reduced timeline for prepositioning, and this negative consequence will be exacerbated if there exists a surge in the logistics cost close to the hurricane’s landfall. In contrast, an optimal adaptive activation policy aims to strike a balance between the value of information gained through waiting and the expected cost of postponing the execution of a logistics plan.

A heuristic activation policy: A simple heuristic activation policy can be defined as follows: Given an MC state ξ_t at time t , a static logistics plan is constructed and activated at time t if the expected remaining time to landfall $\bar{T}(\xi_t)$ is less than a user-defined threshold $\tau > 0$. That is, we set $g(\xi_t) = 1$ if $\bar{T}(\xi_t) \leq \tau$ and $g(\xi_t) = 0$ otherwise. The expected time to landfall $\bar{T}(\xi_t)$ can be calculated using the transition probability matrix $\mathbf{P}$ of the MC. In our numerical study in Section 6, we refer to this as the wait-and-see (WS) policy.

Optimal adaptive activation policy: We propose a decision-tree framework for finding an optimal adaptive activation policy, i.e., an optimal activation decision $g(\xi_t)$ for each $\xi_t \in \Xi$. Specifically, for each transient state $\xi_t \in \Xi^{\text{trans}}$ at time t , we use the 2SSP model (eqs. (3) and (4)) to compute the expected cost $\mathcal{Z}(\xi_t)$ associated with the static logistics plan if it is activated at ξ_t , where the initial condition $x_{t-1} = x_0$. On the other hand, the expected cost of waiting is given by:

$$\sum_{\xi_{t+1} \in \Xi} \mathbf{P}_{\xi_t \xi_{t+1}} \mathcal{Z}^*(\xi_{t+1}) \quad (13)$$

where

$$\mathcal{Z}^*(\xi_t) = \begin{cases} \min \left\{ \mathcal{Z}(\xi_t), \sum_{\xi_{t+1} \in \Xi} \mathbf{P}_{\xi_t \xi_{t+1}} \mathcal{Z}^*(\xi_{t+1}) \right\}, & \text{if } \xi_t \in \Xi^{\text{trans}} \\ \mathcal{Z}(\xi_t), & \text{if } \xi_t \in \Xi^{\text{absor}}. \end{cases} \quad (14)$$

The $\mathcal{Z}(\xi_t)$ value associated with an absorbing state $\xi_t \in \Xi^{\text{absor}}$ is calculated by setting $x_{t-1} = x_0$ and solving the following deterministic optimization problem:

$$\min_{y, \underline{e}, \underline{e}} \sum_{j \in J} \sum_{i \in I(j)} c_{ij}^a y_{ij} + p \sum_{j \in J} \underline{e}_j + q \sum_{i \in I} \bar{e}_i \quad (15a)$$

$$\text{s.t. (1e), (1f), (1i) and (1g) with } d_j \text{ replaced by } d_j(\xi_t). \quad (15b)$$

The recursive formulation in (13) and (14) can be solved using a backward induction algorithm, starting from $t = T_{\max} - 1$ and moving backward until $t = 1$, as shown in Algorithm 2. In our numerical study in Section 6, we refer to this as the decision-tree policy.

Algorithm 2 Optimal adaptive activation condition algorithm.

- 1: *Input:* $x_0, \xi_1, T_{\max}$ and $\mathbf{P}$.
 - 2: For all $\xi_t \in \Xi^{\text{absor}}$: solve (15) to obtain $\mathcal{Z}(\xi_t)$, set $\mathcal{Z}^*(\xi_t) = \mathcal{Z}(\xi_t)$ and $g(\xi_t) = 1$.
 - 3: For all $\xi_t \in \Xi^{\text{trans}}$ that is reachable from ξ_1 : solve the 2SSP model (eqs. (3) and (4)) with initial condition $x_{t-1} = x_0$ to obtain $\mathcal{Z}(\xi_t)$. Set $t = T_{\max} - 1$.
 - 4: **while** $t > 0$ **do**
 - 5: Let Ξ_t be the set of MC states that can be reachable from ξ_1 in exactly $t - 1$ steps.
 - 6: **for** $\xi_t \in \Xi_t \cap \Xi^{\text{trans}}$ **do**
 - 7: **if** $\mathcal{Z}(\xi_t) < \sum_{\xi_{t+1} \in \Xi} \mathbf{P}_{\xi_t \xi_{t+1}} \mathcal{Z}^*(\xi_{t+1})$ **then**
 - 8: Set $\mathcal{Z}^*(\xi_t) = \mathcal{Z}(\xi_t)$ and $g(\xi_t) = 1$
 - 9: **else**
 - 10: Set $\mathcal{Z}^*(\xi_t) = \sum_{\xi_{t+1} \in \Xi} \mathbf{P}_{\xi_t \xi_{t+1}} \mathcal{Z}^*(\xi_{t+1})$ and $g(\xi_t) = 0$
 - 11: **end if**
 - 12: **end for**
 - 13: $t \leftarrow t - 1$
 - 14: **end while**
 - 15: *Output:* The optimal value $\mathcal{Z}^*(\xi_t)$ and activation decision $g(\xi_t)$ for each $\xi_t \in \Xi$ reachable from ξ_1 .
-

Implementing the partially adaptive policy: Given an activation policy $g(\xi_t)$ and a sample path $\xi_{[T]}$, the partially adaptive policy can be implemented by setting $x_t = 0$ for all periods t until reaching a period t' where $g(\xi_{t'}) = 1$.

- If $\xi_{t'} \in \Xi^{\text{trans}}$, the 2SSP model (eqs. (3) and (4)) is solved for the periods $t', t' + 1, \dots, T_{\max}$ with the initial condition $x_{t'-1} = x_0$ and initial MC state $\xi_{t'}$ as input, which gives an optimal first-stage solution $\tilde{\mathbf{f}} = (\tilde{f}_{t'}, \tilde{f}_{t'+1}, \dots, \tilde{f}_{T_{\max}-1})$ and $\tilde{\mathbf{x}} = (\tilde{x}_{t'}, \tilde{x}_{t'+1}, \dots, \tilde{x}_{T_{\max}-1})$. A static policy is then implemented on the remainder of the sample path $\xi_{[T]}$ by setting $x_t = \tilde{x}_t$ and $f_t = \tilde{f}_t$ for all periods $t = t', t' + 1, \dots, T - 1$. At the final period $t = T$, implying that $\xi_T \in \Xi^{\text{absor}}$, the following second-stage problem is solved to collect the overall cost of the wait-and-activate policy on sample path $\xi_{[T]}$:

$$\mathcal{C}(\xi_{[T]}) = \min_{y, \bar{e}, \underline{e}} \sum_{t=t'}^{T-1} \left(h_t \sum_{i \in I} \tilde{f}_{0i,t} + \sum_{i \in I_0} \sum_{i' \in I} c_{ii',t}^b \tilde{f}_{ii',t} + \sum_{i \in I} c_{i,t}^h x_{i,t} \right) + \sum_{j \in J} \sum_{i \in I(j)} c_{ij}^a y_{ij} + p \sum_{j \in J} \underline{e}_j + q \sum_{i \in I} \bar{e}_i$$

s.t. (1e), (1f), (1i) and (1g) with d_j replaced by $d_j(\xi_T)$.

- If $\xi_{t'} \in \Xi^{\text{absor}}$, implying that $t' = T$, solve (15) by setting $x_{T-1} = x_0$, and set $\mathcal{C}(\xi_{[T]})$ to the corresponding optimal objective value.

5.3. Fully adaptive policies

A fully adaptive policy can be constructed in an online fashion by using the rolling-horizon procedure or in an offline fashion by solving the MSP model defined in (9).

Rolling-horizon-based policy The rolling-horizon procedure is a fully adaptive policy that constructs a new plan at each period based on the most up-to-date information at that point but only commits to the decision of the current period. In other words, this policy is obtained in an online fashion as it is implemented on a sample path $\xi_{[T]}$, by solving a new optimization problem typically referred to as the look-ahead model, which captures the remaining planning horizon at each period along sample path $\xi_{[T]}$. In our experiments, we use the 2SSP model (eqs. (3) and (4)) as the look-ahead model, and the rolling-horizon procedure is executed as follows: For every period $t = 1, 2, \dots, T$ and given (x_{t-1}, ξ_t) , solve the 2SSP model (eqs. (3) and (4)) to obtain $(\check{f}_t, \check{x}_t)$, set $x_t = \check{x}_t$, go to the next period, observe the new ξ_t and repeat the process. Algorithm 3 provides a pseudocode description of the rolling-horizon procedure. In our numerical study in Section 6, we refer to this as the rolling-horizon 2SSP (rh-2SSP) policy.

Algorithm 3 The rolling horizon procedure.

```

1: Input: Given  $x_0$  and a sample path  $\xi_{[T]} = (\xi_1, \xi_2, \dots, \xi_T)$ .
2: Initialize the cumulative cost of the sample path as  $\mathcal{C}(\xi_{[T]}) = 0$ .
3: for  $t \in \{1, 2, \dots, T\}$  do
4:   if  $t < T$  (implying that  $\xi_t \in \Xi^{\text{trans}}$ ) then
5:     Define and 2SSP model with  $T_{\max} - t + 1$  periods and  $(x_{t-1}, \xi_t)$  as the initial condition.
6:     Solve the model to obtain a new  $(\check{x}_t, \check{x}_{t+1}, \dots, \check{x}_{T_{\max}-1})$  and  $(\check{f}_t, \check{f}_{t+1}, \dots, \check{f}_{T_{\max}-1})$ .
7:     Set  $x_t = \check{x}_t$  and let  $\mathcal{C}(\xi_{[T]}) \leftarrow \mathcal{C}(\xi_{[T]}) + h_t \sum_{i \in I} \check{f}_{0i,t} + \sum_{i \in I_0} \sum_{i' \in I} c_{ii',t}^b \check{f}_{ii',t} + \sum_{i \in I} c_{i,t}^h \check{x}_{i,t}$ .
8:   else
9:     Let  $\mathcal{C}(\xi_{[T]}) \leftarrow \mathcal{C}(\xi_{[T]}) + z(x_{T-1}, \xi_T)$  where  $z(x_{T-1}, \xi_T)$  is the optimal objective value of model (15).
10:  end if
11: end for
12: Output: The cumulative cost  $\mathcal{C}(\xi_{[T]})$  of the sample path  $\xi_{[T]}$ .
```

MSP-based policy A fully adaptive policy is obtained upon solving the MSP model (5) or its dynamic programming formulation (9) using the nested Benders decomposition algorithm described in Algorithm 1 through the cutting plane approximations of the expected cost-to-go functions $\check{\mathcal{Q}}_{t+1}^{\xi_t}(\cdot)$ for each $\xi_t \in \Xi$ and $t = 1, 2, \dots, T_{\max} - 1$. Given a sample path $\xi_{[T]} = (\xi_1, \xi_2, \dots, \xi_T)$, the MSP-based policy can be implemented as follows: for every period $t = 1, 2, \dots, T$, we set

$$\check{a}_t := \operatorname{argmin}_{a_t} \left\{ z_t(a_t, \xi_t) + \check{\mathcal{Q}}_{t+1}^{\xi_t}(a_t) \mid a_t \in \chi_t(\check{a}_{t-1}, \xi_t) \right\} \quad (16)$$

where a_t is defined in (6), $z_t(a_t, \xi_t)$ is defined in (7) and $\chi_t(\check{a}_{t-1}, \xi_t)$ is defined in (8). The cumulative cost $\mathcal{C}(\xi_{[T]})$ of the sample path $\xi_{[T]}$ is then calculated by: $\mathcal{C}(\xi_{[T]}) = \sum_{t=1}^T z_t(\check{a}_t, \xi_t)$. In our numerical study in Section 6, we refer to this as the MSP policy.

Remark 1. *The level of adaptability in the policies we discussed in this section is determined based on three features: (1) the stochastic programming model used to design the policy (i.e., 2SSP or MSP); (2) whether the policy is constructed in an offline or online (i.e., using a rolling-horizon procedure) setting, and (3) whether or not one adopts a wait-and-activate decision structure. Every combination of these three features can result in a policy, but not each of the resulting policies is meaningful. For instance, combining an MSP model with either an online setting or a wait-and-activate structure*

is redundant in our setting, as the MSP-based policy has full adaptability and optimality guarantee through the dynamic programming equations (9). On the other hand, the 2SSP model (eqs. (3) and (4)) can benefit significantly from the adaptability of the online setting through the rolling-horizon procedure or the partial adaptability through the wait-and-activate decision structure. However, their performances are expected to be inferior to the MSP policy in general: rh-2SSP does not follow the dynamic programming equations (9) and thus does not have any optimality guarantee, while the decision-tree approach provides a wait-and-activate policy that only has partial adaptability. We will demonstrate the performances of these policies in our numerical study next.

6. Numerical Results

In this section, we present and analyze the computational performances of the following approaches and the associated decision policies presented in Section 5: (i) the fully-adaptive MSP model (MSP), (ii) the static 2SSP model (s-2SSP), (iii) the rolling-horizon 2SSP approach (rh-2SSP), (iv) an optimal activation policy for the wait-and-activate type policy obtained by the decision-tree approach (Decision-tree), and (v) a simple heuristic activation policy for the wait-and-activate type policy based on the expected remaining time to landfall (WS). We present results using two different sets of instances: a set of synthetic instances from the literature (Galindo and Batta, 2013; Pacheco and Batta, 2016) and a case study using realistic data based on Hurricane Florence (in 2018) (Bhattarai and Song, 2023). We evaluate each approach in these instances under deterministic and random landfall time settings and other configuration parameters. We use the synthetic instances to provide a comprehensive comparison of solution quality and computational efficiency among different approaches. We then present the results of the case study instances where sensitivity analyses and solution-level comparisons are performed to provide managerial insights for different decision policies.

6.1. Problem data

In this section, we give an overview of how we generate the configuration of the disaster relief logistics network and how we specify the problem parameters in our numerical experiments. We provide a high-level description of the test instances below, and the detailed problem data generation for both the synthetic and case study instances can be found in Appendix B.

6.1.1. Synthetic instances

For the set of synthetic instances, we randomly generate HDRLP instances with a different number of DPs $|J| \in \{10, 20, 30\}$ and SPs $|I| \in \{3, 6, 9\}$ under the deterministic landfall time case and random landfall time case. Many of the problem parameters for this set of instances follow those given in the literature, including the MC for the deterministic landfall time case (see Galindo and Batta (2013) and Pacheco and Batta (2016)). The random landfall time case is constructed by assuming that the hurricane’s location state ξ_t^{loc} corresponds to a pair of coordinates $(\xi_{t,x}^{\text{loc}}, \xi_{t,y}^{\text{loc}})$, each of which follows an independent MC. For the x -coordinate $\xi_{t,x}^{\text{loc}}$, we use the same MC as the one used for the

deterministic landfall time case, and we specify the MC for the y -coordinate $\xi_{t,y}^{\text{loc}}$ by assuming that the hurricane’s forward speed on the y -axis follows a certain probability distribution. See Appendix B.1 for more details.

6.1.2. Case study instances

Our case study is constructed based on realistic data from Hurricane Florence (in 2018) for a study region that contains coastal South Carolina. We use the realistic logistics network data described in SCEMD (2023). Specifically, we approximate the coastline by a straight line segment and consider the potential landfall location of the hurricane to be anywhere within an extension of the coastline by 200 miles from either of its endpoints. We consider a set of points of distribution (PoDs) as the DPs, a set of regional staging areas (RSAs) as the SPs, and consider the MDC to be a FEMA warehouse located near Atlanta, GA. The details about these locations, including their coordinates and inventory capacities, are provided in Table B.11, Table B.12, and Table B.13 in Appendix B.2. The detailed specification for the costs of logistics operations, including procurement, transportation, and inventory costs, can also be found in Appendix B.2.

The full description of how we model the hurricane’s evolution as an MC can be found in Bhattarai and Song (2023), and we provide a high-level description here for the sake of completeness. We consider each period to be 12 hours long, corresponding to the interval between consecutive forecast lead times in the forecast error database published by the National Hurricane Center (National Hurricane Center, 2012). We set the start time $t = 1$ to be H-120 (i.e., 120 hours before the actual hurricane landfall time) for Hurricane Florence (in 2018). The hurricane’s track and intensity (point) forecast is given by the NHC’s official best forecast product at H-120. To incorporate the forecast uncertainty, we fit a time series model on the historical forecast error data (on both track and intensity) from the last 20 years, and generate sample paths of the forecast errors from the time series model. These sample paths of forecast errors, together with the point forecast, provide sample paths for potential trajectories of hurricane’s locations and intensities over time. We create two versions of MCs, corresponding to a deterministic landfall time and a random landfall time, respectively. The MC with a deterministic landfall time is constructed with the assumption that the hurricane’s forward speed towards the orthogonal direction of the coastline (known as the along-track direction) is constant, and the track forecast error is reflected in the cross-track direction (known as the cross-track error) only. For this deterministic landfall time case, we have a planning horizon of $T = 10$ periods (landfall occurs at $t = 10$). The MC with a random landfall time is constructed without such assumption, and both the along-track error and cross-track error are considered. As a result of the MC discretization (see Figure E.28 in Appendix E), we have a planning horizon of $T_{\max} = 12$. All the forecast errors (along-track error, cross-track error, and intensity error) are assumed to be independent in our numerical experiments, which is supported by a statistical analysis using realistic forecast error data from the NHC database (National Hurricane Center, 2022). To construct an MC model, we apply an optimization-based discretization approach to find the MC states and the associated transition probability matrix that fit the best with these sample paths (see, e.g., Löhndorf and Shapiro (2019) and the references therein). Figure E.28 in Appendix E provides an illustration of the study region, the sample paths of hurricane trajectories, and the MC discretization. We note that the performances of a given decision policy on the two MCs, i.e., the one

with deterministic landfall time and the one with random landfall time, cannot be directly compared with each other. The two MCs are constructed via different levels of discretization on sample paths generated under different assumptions on the forecast errors. The choice of the underlying stochastic processes (e.g, MCs) to approximate the true process (for obtaining a better policy) is out of the scope of this paper. It falls into the topic of the predict/estimate-and-optimize framework, and readers can refer to, e.g., Sadana et al. (2024) for more details.

6.1.3. Logistics cost structure

We apply a logistics cost structure that is motivated by the disruptions to the day-to-day operations of the logistics network, caused by the imminence of the hurricane. In particular, the days before the hurricane’s predicted landfall are typically associated with damage-inducing weather conditions, which may disrupt the logistics network. These disruptions can ultimately lead to a surge in the prices of relief commodities and their logistics costs (Beatty et al., 2021). Specifically, we focus on the cost structure based on a safety threshold: the unit logistics costs associated with state ξ_t stay as the baseline values when the expected time until landfall $\bar{T}(\xi_t)$ is larger than a pre-defined safety threshold τ_{safe} , and increase to $\nu \geq 1$ times the baseline values otherwise. For example, let ω be a baseline transportation cost parameter (e.g., unit fuel cost) and let ℓ_i and $\ell_{i'}$ be the locations of $i \in \{0\} \cup I$ and $i' \in I$, respectively. Given a state $\xi_t \in \Xi$ at time t , we set the unit cost transporting relief commodities from node $i \in \{0\} \cup I$ to node $i' \in I$ as:

$$c_{ii',t}^b(\xi_t) = \begin{cases} \omega \times \text{dist}(\ell_i, \ell_{i'}) & \text{if } \bar{T}(\xi_t) > \tau_{\text{safe}} \\ \omega \times \text{dist}(\ell_i, \ell_{i'}) \times \nu & \text{otherwise,} \end{cases}$$

where $\nu \geq 1$ is the cost surge parameter and $\text{dist}(\cdot, \cdot)$ is a distance function that measures the distance between the two locations. We note that letting $\nu \rightarrow \infty$ (or a large enough value) corresponds to the situation where the logistics operations are forbidden when the expected time until landfall falls within the safety threshold, following a standard safety protocol in pre-disaster logistics operation (Morrice et al., 2016). In our numerical experiments, we use a baseline safety threshold parameter $\tau_{\text{safe}} = 2$, i.e., two periods (24 hours), and conduct a sensitivity analysis on this parameter. For the implementation of the WS policy on this cost structure, we set the activation threshold $\tau = \tau_{\text{safe}} + 1$, i.e., the activation decision $g(\xi_t) = 1$ when $\bar{T}(\xi_t) \leq \tau = \tau_{\text{safe}} + 1$. This means that for each sample path $\xi_{[T]}$, to follow this WS policy, one would wait (to activate a static plan) until the last period before the cost surge is expected to take place. We also consider an alternative cost structure where the cost increases linearly over time, regardless of the corresponding states, and the results of the instances with this alternative cost structure can be found in Appendix C.1.

6.1.4. Demand estimation

We define the demand function $d_j(\xi_t)$ following Assumption 3 and Assumption 5 as:

$$d_j(\xi_t) = P_j \times \frac{\xi_t^{\text{int}^2}}{(|\Xi^{\text{int}}| - 1)^2} \times \max \left\{ \left(1 - \frac{||\xi_t^{\text{loc}} - \ell_j||}{D_{\text{max}}} \right), 0 \right\} \times \mathbb{1}_{\{\xi_t^{\text{loc}} \in L\}}, \quad (17)$$

where ℓ_j denotes the location of DP j , P_j is the maximum possible demand from DP j , $D_{\max}$ is the maximum distance from a landfall location at which a DP is affected by the hurricane, and the indicator function $\mathbb{1}_{\{\xi_t^{\text{loc}} \in L\}}$ ensures that the demand is only incurred at the hurricane’s landfall. Equation (17) implies that the demand for relief commodities $d_j(\xi_t)$ is positively correlated to the hurricane intensity and negatively correlated to the distance between the hurricane’s landfall location and the DP’s location, which aligns with the literature (Baker, 1991; Kirac and Milburn, 2018; Morrice et al., 2016). We emphasize that (17) is a simplification that is used for illustrative purposes. A detailed study of the relationship between the hurricane’s characteristics and the demand for relief commodities is beyond the scope of this paper.

All the data sets and codes developed for this paper can be found in a GitHub repository (link to be provided upon the acceptance of this manuscript).

6.2. Main results

We next present the numerical experiment results based on a set of synthetic instances as well as real-world instances inspired by Hurricane Florence (in 2018). To evaluate the performances of these different approaches, we create a set of out-of-sample scenarios according to the MC, starting from a given initial state ξ_1 , with a sample size of $N = 1000$. We compute the sample mean $\hat{z} := \sum_{n=1}^N z(\hat{\xi}^n)/N$ and the sample standard deviation $\hat{\sigma} := \sqrt{\sum_{n=1}^N (z(\hat{\xi}^n) - \hat{z})^2/(N-1)}$, where $\hat{\xi}^n := (\hat{\xi}_2^n, \hat{\xi}_3^n, \dots, \hat{\xi}_T^n)$ is the realized trajectory of the stochastic process corresponding to sample path n and $z(\hat{\xi}^n)$ is the total objective value of the corresponding decision policy accumulated in sample-path n . We compute a 95% confidence interval (CI) for the performance of the corresponding policy on the set of out-of-sample scenarios, which is given by $\hat{z}_{\pm} := \hat{z} \pm 1.96\hat{\sigma}/\sqrt{N}$. We benchmark the performance of each approach by computing the relative gap to the sample mean performance of the perfect-information (PI) solution obtained by solving the deterministic formulation (1) for each sample path $\hat{\xi}^n$. Additional implementation details can be found in Appendix D.

6.2.1. Numerical results on the synthetic instances

In Table 3 and Table 4, we report the numerical results for the PI solution, MSP, s-2SSP, rh-2SSP, WS, and Decision-tree approaches on the synthetic instances for HDRLP with a different number of DPs $|J| \in \{10, 20, 30\}$ and SPs $|I| \in \{3, 6, 9\}$ with deterministic and random landfall time, respectively, using the safety threshold-based logistics cost structure under different cost surge parameter $\nu \in \{1, 2, 5, 10\}$. In particular, $\nu = 1$ corresponds to the special case of no cost surge. Column “time” shows the computational time (in seconds) for training each policy. We do not report the time for training policies for rh-2SSP since it is obtained in an online fashion during the out-of-sample evaluation. We do not report the time for WS for a similar reason as it pertains to solving a (single) static 2SSP model for each sample-path during the out-of-sample evaluation. Column “gap” shows the relative gap of the out-of-sample performance of each approach to that of the PI solution, and column “CI” shows the confidence interval for the performance of the PI solution in the out-of-sample evaluation, illustrating the variability in different sample paths. Additionally, we provide the lower bound (LB) and the statistical upper bound (UB) of MSP, illustrating the solution quality of the trained fully-adaptive policy by the MSP. Figure 2 and Figure 3 show the

performances of different policies, *averaged* across different configurations of the logistics network instances (in terms of different $|I|$ and $|J|$) for each value for the cost surge parameter ν , with deterministic and random landfall time, respectively.

ν	$ I $	$ J $	PI		MSP				Decision-tree		WS	rh-2SSP	s-2SSP	
			$\hat{z}$	CI	LB	UB	time	gap	time	gap	gap	gap	time	gap
1	3	10	710.8	93.0	1693.2	1646.4±147.5	20.1	131%	2.8	131%	350%	131%	0.0	753%
		20	1421.6	179.0	3175.3	3061.8±265.8	22.6	115%	4.6	115%	331%	115%	0.0	696%
		30	2123.2	267.0	4859.6	4696.2±407.4	24.3	121%	6.2	121%	337%	121%	0.1	707%
	6	10	1093.4	286.9	1953.1	1965.2±299.4	28.1	79%	4.2	79%	216%	79%	0.0	462%
		20	1645.7	294.4	3315.3	3202.7±337.8	31.1	94%	6.7	94%	281%	94%	0.1	603%
		30	2542.8	477.8	5205.7	5017.4±548.5	38.1	97%	9.2	97%	276%	97%	0.1	581%
	9	10	908.8	202.0	1807.3	1801.0±224.0	29.6	98%	5.6	98%	267%	98%	0.0	579%
		20	1458.4	194.4	3249.1	3058.3±267.0	39.7	109%	8.8	109%	325%	109%	0.1	700%
		30	2243.7	321.6	5021.9	4788.5±437.7	48.7	113%	12.0	113%	320%	113%	0.1	672%
2	3	10	828.2	109.4	2422.1	2344.3±197.5	25.3	183%	3.0	236%	317%	251%	0.0	695%
		20	1653.4	209.0	4494.8	4345.6±393.1	31.0	162%	5.2	220%	303%	257%	0.0	654%
		30	2451.9	308.6	6852.3	6598.5±598.5	36.2	169%	6.5	219%	312%	249%	0.1	656%
	6	10	1202.7	300.3	2636.7	2642.1±329.3	29.9	119%	4.2	151%	210%	161%	0.0	450%
		20	1862.3	322.1	4589.6	4494.6±447.2	44.5	141%	7.1	189%	268%	221%	0.1	573%
		30	2850.8	514.0	7114.9	6855.9±695.8	43.0	140%	9.7	182%	262%	209%	0.1	554%
	9	10	1029.0	218.0	2516.9	2486.8±259.6	41.4	141%	5.9	182%	252%	194%	0.0	543%
		20	1690.6	228.9	4559.9	4393.2±404.5	63.6	159%	9.9	215%	302%	249%	0.1	642%
		30	2572.8	367.8	6995.0	6702.0±624.8	48.0	160%	13.0	208%	298%	238%	0.1	625%
5	3	10	1032.2	140.0	3123.7	3000.7±277.5	28.1	190%	3.2	259%	259%	218%	0.0	558%
		20	2047.9	261.7	5851.2	5736.2±614.2	32.6	180%	5.5	249%	249%	220%	0.1	529%
		30	2985.0	377.1	8859.4	8960.5±1137.4	36.2	200%	7.1	257%	257%	234%	0.1	540%
	6	10	1384.6	327.8	3292.3	3223.6±385.4	36.4	132%	4.3	182%	182%	153%	0.0	389%
		20	2213.2	374.7	5888.9	5791.7±636.5	35.3	161%	7.8	225%	225%	200%	0.1	481%
		30	3323.4	578.3	9021.7	9059.7±1176.0	44.2	172%	10.2	226%	226%	203%	0.1	475%
	9	10	1242.2	252.7	3222.4	3133.6±330.7	64.7	152%	6.3	208%	208%	176%	0.1	447%
		20	2085.0	296.2	5915.1	5749.1±612.0	70.0	175%	11.0	245%	245%	217%	0.1	518%
		30	3102.2	453.8	8983.2	9021.0±1150.4	54.7	190%	14.7	247%	247%	223%	0.1	517%
10	3	10	1372.1	191.7	3811.9	3806.6±609.4	25.2	177%	3.4	194%	194%	184%	0.0	419%
		20	2705.5	350.5	7305.6	7464.7±1225.8	34.5	175%	5.8	189%	189%	183%	0.1	401%
		30	3873.4	492.3	10781.1	11110.3±1869.1	38.5	186%	7.7	199%	199%	192%	0.1	417%
	6	10	1687.7	375.5	3936.1	3971.3±665.9	24.5	135%	4.8	149%	149%	140%	0.0	317%
		20	2797.7	465.0	7252.0	7436.5±1238.5	39.8	165%	8.2	178%	178%	173%	0.1	379%
		30	4110.5	689.9	10810.4	11146.4±1897.1	62.7	171%	11.0	182%	182%	176%	0.1	383%
	9	10	1597.3	313.6	3909.8	3935.0±641.2	54.9	146%	6.7	161%	161%	151%	0.1	343%
		20	2742.0	410.3	7339.4	7451.7±1222.2	68.1	171%	11.9	186%	186%	179%	0.1	391%
		30	3983.2	600.8	10849.9	11168.1±1876.2	45.1	180%	16.3	192%	192%	185%	0.2	400%

Table 3: Performances of the PI solution and decision policies given by MSP, Decision-tree, WS, rh-2SSP, and s-2SSP for synthetic instances, deterministic landfall time case, and the safety threshold-based cost structure.

From Tables 3 and 4 and Figures 2 and 3, we observe the general pattern that MSP yields the best performance, followed by rh-2SSP, Decision-tree, WS, and finally s-2SSP. This pattern can be observed across all different cost surge values ν . We note that the performance of different policies is closely related to the adaptability of the associated policies: higher adaptability leads to better performance, due to the capability of leveraging the updated information during the stochastic process modeled by the underlying MC. We also see the general pattern that policies with higher adaptability are computationally more expensive to solve. In terms of computational time, with our computational enhancements to the nested Benders decomposition (see Table D.18 and D.19 in Appendix D), we managed to obtain high-quality policies associated with the MSP model (measured by the gap between LB and UB) within a reasonable computational time (mostly under 10 minutes). We note that the computational time associated with Decision-tree and s-2SSP policies is significantly longer for the random landfall time case than the deterministic case. This

ν	$ I $	$ J $	PI		MSP				Decision-tree		WS	rh-2SSP	s-2SSP	
			$\hat{z}$	CI	LB	UB	time	gap	time	gap	gap	gap	time	gap
1	3	10	575.3	65.2	1755.0	1745.0±182.7	138.7	203%	55.4	395%	551%	238%	0.7	666%
		20	1154.7	124.5	3324.1	3332.0±293.7	238.1	188%	68.0	375%	489%	222%	0.5	648%
		30	1718.5	185.0	5109.5	5123.1±437.5	216.7	198%	76.1	371%	457%	234%	0.5	680%
	6	10	639.3	129.1	1855.8	1797.3±210.8	222.9	181%	121.1	346%	483%	210%	1.1	601%
		20	1176.2	153.0	3370.0	3334.6±302.4	179.3	183%	148.0	365%	477%	217%	1.2	647%
		30	1782.0	253.7	5241.1	5152.1±461.6	277.5	189%	159.5	353%	436%	224%	1.3	664%
	9	10	602.0	95.0	1798.9	1767.2±193.7	447.7	193%	229.0	373%	521%	226%	2.1	641%
		20	1153.1	128.3	3348.9	3330.0±293.8	219.4	188%	278.1	376%	492%	223%	2.3	657%
		30	1732.4	201.9	5173.1	5143.2±449.4	287.8	196%	296.4	368%	453%	233%	2.4	679%
2	3	10	654.8	74.9	2268.8	2215.3±185.4	178.4	238%	94.8	462%	507%	418%	0.7	622%
		20	1313.2	142.1	4253.6	4216.1±365.1	240.8	221%	126.9	412%	472%	385%	0.8	655%
		30	1940.8	209.1	6510.5	6713.2±626.8	334.8	245%	138.6	406%	476%	410%	0.8	703%
	6	10	710.1	136.9	2344.3	2237.3±205.3	258.8	215%	233.7	412%	454%	384%	1.7	570%
		20	1315.1	169.3	4271.6	4196.4±364.7	219.4	219%	302.8	410%	467%	387%	2.0	659%
		30	1976.9	274.5	6591.0	6675.9±629.3	254.2	237%	330.3	394%	461%	401%	1.9	694%
	9	10	678.6	104.6	2300.5	2210.9±187.3	295.1	225%	397.6	440%	482%	407%	1.9	601%
		20	1299.5	147.4	4269.5	4214.8±367.5	467.6	224%	548.6	417%	477%	403%	4.0	669%
		30	1936.2	227.5	6557.1	6704.8±627.1	750.2	246%	574.8	405%	476%	422%	3.9	712%
5	3	10	816.8	95.7	2920.9	2870.1±277.6	178.1	251%	96.8	389%	409%	362%	0.8	504%
		20	1630.5	177.5	5504.5	5390.0±503.4	235.2	230%	129.1	362%	391%	329%	0.9	492%
		30	2369.7	254.5	8396.7	8413.5±919.3	207.2	255%	151.9	380%	407%	362%	0.9	512%
	6	10	846.3	155.4	2954.3	2825.4±281.2	181.2	233%	237.9	367%	384%	340%	1.7	478%
		20	1574.0	203.5	5453.8	5374.3±551.5	227.0	241%	326.6	387%	403%	340%	2.1	510%
		30	2322.6	315.4	8384.0	8322.6±906.8	261.0	258%	348.8	385%	412%	366%	2.4	522%
	9	10	831.6	127.6	2945.9	2852.4±285.1	289.1	243%	445.2	377%	396%	350%	1.9	490%
		20	1578.6	188.2	5490.8	5353.7±504.4	433.1	239%	620.4	389%	405%	342%	4.5	510%
		30	2303.7	278.9	8396.1	8315.1±892.7	459.9	260%	641.3	392%	419%	375%	4.9	528%
10	3	10	1086.7	130.9	3459.1	3635.3±578.2	131.6	234%	105.3	310%	323%	288%	0.8	379%
		20	2159.2	237.2	6680.7	6925.1±1019.4	187.0	220%	134.2	322%	310%	267%	0.9	373%
		30	3084.6	331.0	9961.1	9976.9±1318.9	195.0	223%	150.0	336%	330%	292%	0.9	395%
	6	10	1073.3	187.9	3446.0	3562.7±582.0	194.1	231%	248.7	313%	318%	282%	1.4	378%
		20	2005.3	262.2	6551.6	6685.8±980.5	211.7	233%	345.2	346%	333%	285%	2.5	401%
		30	2898.2	386.4	9805.9	9839.4±1373.3	152.4	239%	368.6	356%	349%	306%	2.2	420%
	9	10	1086.5	168.0	3464.3	3596.3±579.0	260.4	230%	463.6	305%	317%	281%	1.9	374%
		20	2043.5	257.8	6610.8	6753.9±1002.1	329.2	230%	650.2	341%	328%	283%	4.4	393%
		30	2915.3	366.6	9866.8	9828.1±1336.6	322.2	237%	690.2	356%	349%	308%	5.1	417%

Table 4: Performances of the PI solution and decision policies given by MSP, Decision-tree, WS, rh-2SSP, and s-2SSP for synthetic instances, random landfall time case, and the safety threshold-based cost structure.




Figure 2: Performances of the PI solution and decision policies given by MSP, Decision-tree, WS, rh-2SSP, and s-2SSP for synthetic instances, deterministic landfall time case, and the safety threshold-based cost structure.




Figure 3: Performances of the PI solution and decision policies given by MSP, Decision-tree, WS, rh-2SSP, and s-2SSP for synthetic instances, random landfall time case, and the safety threshold-based cost structure.

is because for the random landfall time case, the expected remaining time until landfall at each time t depends on the MC state ξ_t at that time, and thus, the logistics cost parameters are MC state-dependent. Therefore, to exactly solve the 2SSP model (3), we need to consider all possible sample paths, unlike the deterministic landfall time case where the logistics cost parameters are only time-dependent, and we only need to focus on the absorbing states that are reachable from ξ_1 and the corresponding probability of absorption. In our implementation, we resort to a standard sample average approximation approach by sampling a large number of sample paths, which are often referred to as the “in-sample” scenarios, when solving the static two-stage SPs (used in s-2SSP, rh-2SSP, Decision-tree, and WS). We also attribute the slightly inferior performances of decision-tree compared to WS in a few instances in Table 4 to the sampling errors introduced by these in-sample scenarios.

Despite the general trend, we see that the comparisons in the performances of different policies vary significantly with different cost parameters ν . For example, for the deterministic landfall time case (Table 3 and Figure 2), when $\nu = 1$ (no cost surge occurs), we see almost identical performances between MSP, Decision-tree, and rh-2SSP, and an inferior performance by WS, followed by s-2SSP. In this case, policies that exhibit adaptability in their decisions (including MSP, Decision-tree, and rh-2SSP) have successfully exploited the value of information gained by waiting until the uncertainty in the hurricane’s landfall state is sufficiently resolved and then implementing their logistics operations. On the contrary, the lack of adaptability has caused the performances of WS and s-2SSP to be inferior: s-2SSP commits to a static logistics plan at $t = 1$ and does not adapt to new information completely, while WS waits until the last period before the expected remaining time to landfall falls below the safety threshold $\tau_{\text{safe}} = 2$, but it may be advantageous to wait longer in this case as the cost surge parameter $\nu = 1$. On the other hand, when $\nu = 10$, we see that all policies perform similarly (despite a slight advantage by MSP) except s-2SSP. This is due to the significant cost surge when the expected remaining time to landfall falls below τ_{safe} , making it too expensive to perform an adaptive plan compared to committing (and implementing) a static plan right before the cost surge occurs. As a result, it is reasonable to commit to a static logistics plan and implement it before the cost surge occurs, explaining the competitive performances of the

partial adaptive policies of Decision-tree and WS in this situation. Between these extreme cases, the superiority of MSP becomes clear: it strikes a balance between the benefit of waiting to adapt to new information and the cost surge incurred from waiting. The performance comparison between rh-2SSP and Decision-tree has mixed results in different cases.

Comparing the deterministic and random landfall time cases, we see that although the general trends discussed above are consistent in both cases, the superiority of the fully adaptive policies (MSP and rh-2SSP) becomes more clear when the landfall time is random, while the performance of the partially adaptive policies (Decision-tree and WS) is affected by the randomness in the landfall time. We also observe that the advantage of MSP compared to rh-2SSP is significant not only when $\nu = 2$ or 5, as in the deterministic landfall time case, but across all different values of ν . Additionally, by comparing the computational performance results shown in Table 4 and 3, we see that the random landfall time case is more computationally demanding for solving both the static 2SSP model and the MSP model.

6.2.2. Numerical results on the case study instances

In Tables 5 and Table 6, we report the numerical results for the PI solution, MSP, s-2SSP, rh-2SSP, WS, and Decision-tree approaches on the realistic instances of HDRLP inspired by Hurricane Florence (in 2018) with deterministic and random landfall time, respectively, using the safety threshold-based logistics cost structure with different cost surge parameters ν . In particular, we include an additional parameter value $\nu = \infty$ (we use a sufficiently large number, 10^6 in our implementation), which corresponds to the special case that no logistics operation is allowed when the expected time until landfall falls below the safety threshold τ_{safe} .

ν	MSP				s-2SSP		rh-2SSP	WS	Decision-tree	
	LB	UB	time	gap	time	gap	gap	gap	time	gap
1	2381917.4	2224109.6±128755.9	172.8	148%	0.1	313%	155%	256%	13.9	209%
2	2821117.5	2702341.0±158119.7	134.8	182%	0.1	314%	199%	370%	12.6	260%
5	3124434.3	2985251.8±154233.0	318.2	213%	0.2	315%	216%	625%	13.1	262%
10	3208616.2	3119646.2±195906.3	182.0	228%	0.1	316%	231%	889%	14.5	264%
∞	3279854.3	3100418.8±183713.7	99.7	223%	0.1	313%	232%	1125%	11.7	262%

Table 5: Performances of the decision policies given by MSP, s-2SSP, rh-2SSP, decision-tree, and WS for the case study instances with deterministic landfall time.

ν	MSP				s-2SSP		rh-2SSP	WS	Decision-tree	
	LB	UB	time	gap	time	gap	gap	gap	time	gap
1	2526503.1	2446379.8±214354.7	76.7	216%	4.7	329%	248%	457%	1812.9	302%
2	2829311.5	2765838.6±195737.9	66.5	239%	4.7	322%	247%	540%	1798.1	326%
5	3009203.0	3025749.2±259595.5	87.1	272%	4.7	323%	294%	820%	1774.6	320%
10	3081331.1	3083036.9±268750.5	81.4	281%	4.6	325%	305%	1049%	1787.6	321%
∞	3124531.8	3078117.7±263259.7	61.8	277%	4.7	322%	310%	1230%	1860.7	318%

Table 6: Performances of the decision policies given by MSP, s-2SSP, rh-2SSP, decision-tree, and WS for the case study instances with random landfall time.

We observe from Table 5 and Table 6 that the general pattern observed from the performances of

different approaches on the synthetic instances is the same, except that the performances of s-2SSP and rh-2SSP are much closer to MSP. In addition, we observe that the performance of WS gets worse when the value of the cost surge factor ν increases. This is attributed to the limited flow capacity in the logistics network, making it insufficient to only rely on the prepositioning conducted in the period when the static logistics plan is activated, which results in logistics operations at higher logistics costs (to fulfill the demand) due to the cost surge. Policies with adaptability either in the logistics operation itself, i.e., MSP and rh-2SSP, or in the activation policy for the wait-and-activate type policies, i.e., Decision-tree, do not significantly suffer from this issue. The s-2SSP policy, on the other hand, does not suffer from the limited flow capacity either, thanks to its “early commitment” feature, allowing it to build an inventory of prepositioned commodities over time. To better understand the impact of the flow capacity on the behaviors of different policies, we conduct a sensitivity analysis on the flow capacity level next.

Sensitivity analysis on the flow capacity level. We perform sensitivity analysis by considering flow capacities that are quarter, half, double, and quadruple of the baseline value (0.25, 0.5, 2, and 4, respectively). We set $\nu = \infty$ for this sensitivity analysis to focus on the impact of flow capacity levels on the behaviors of different policies. Tables 7 and 8 present the performances of different policies for the deterministic landfall time case and random landfall time case, respectively. When the flow capacity level is set to the highest (4), the WS policy is optimal as the flow capacity is sufficiently large so that all the logistics operations can be performed at the very last period, after which no logistics operation is allowed (as ν is set to ∞). The same holds true for Decision-tree and MSP. (We note that the performance of MSP is slightly inferior to WS and Decision-tree, by less than 0.6%, and this is attributed to the practical termination criterion used in the training process, see details in Appendix D, which leads to a slightly suboptimal policy.) On the other hand, when the flow capacity level is set to the lowest (0.25), WS underperforms significantly due to insufficient prepositioned commodities caused by its heuristic activation policy based on a fixed threshold parameter. In this case, we also see that the value of adaptability diminishes when comparing the MSP and s-2SSP policies, since the limited flow capacity restricts the adaptability to the new information. We next characterize the behaviors of different policies by illustrating the solution-level structures exhibited by MSP, s-2SSP, and Decision-tree.

flow-capacity	PI	MSP	s-2SSP	WS	Decision-tree
	$\hat{z}$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$
0.25	5185.1K	(6342.9K, 22%)	(6405.1K, 23%)	(23739.1K, 357%)	(6405.1K, 23%)
0.5	1170.4K	(3455.7K, 195%)	(4174.8K, 256%)	(18739.6K, 1501%)	(3842.1K, 228%)
1	957.7K	(3100.4K, 223%)	(3962.3K, 313%)	(11735.8K, 1125%)	(3467.3K, 262%)
2	946.4K	(2960.0K, 212%)	(3858.7K, 307%)	(5004.9K, 428%)	(3183.5K, 236%)
4	943.4K	(2922.9K, 209%)	(3806.9K, 303%)	(2906.0K, 208%)	(2906.0K, 208%)

Table 7: Performances of the PI solution and decision policies given by MSP, s-2SSP, decision-tree, and WS for the case study instances with deterministic landfall time and safety threshold-based cost structure under different flow capacity levels.

Figure 4, 5 and 6 illustrate the solution-level behavior of MSP, Decision-tree, and s-2SSP for the deterministic landfall time case, respectively, under different flow capacity levels. Similar illustrations for the random landfall time case are provided in Appendix C.2. In all three figures, we plot the average amount of commodity procurement over time (averaged over the set of sample paths in the

flow-capacity	PI	MSP	s-2SSP	WS	Decision-tree
	$\hat{z}$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$
0.25	3007.7K	(4832.4K, 60%)	(4897.5K, 62%)	(20638.8K, 586%)	(4897.4K, 62%)
0.5	838.1K	(3322.9K, 296%)	(3629.1K, 333%)	(16412.7K, 1858%)	(3650.6K, 335%)
1	815.3K	(3078.1K, 277%)	(3443.0K, 322%)	(10848.6K, 1230%)	(3413.0K, 318%)
2	805.2K	(2950.5K, 266%)	(3357.1K, 316%)	(5087.4K, 531%)	(3115.6K, 286%)
4	802.1K	(2893.7K, 260%)	(3314.1K, 313%)	(2936.8K, 266%)	(2893.9K, 260%)

Table 8: Performances of the PI solution and decision policies given by MSP, s-2SSP, decision-tree, and WS for the case study instances with random landfall time and safety threshold-based cost structure under different flow capacity levels.




Figure 4: Behavior of MSP solution for the case study instances with deterministic landfall time under varying flow capacity levels with the safety threshold-based cost structure.




Figure 5: Behavior of Decision-tree solution for the case study instances with deterministic landfall time under varying flow capacity levels.




Figure 6: Behavior of s-2SSP solution for the case study instances with deterministic landfall time under varying flow capacity levels with the safety threshold-based flow restrictions.

out-of-sample test). Additionally, in Figure 4 we provide, for each period, the percentage of sample paths where a positive amount of procurement occurs in the MSP solution (%Procurement). In Figure 5, we use a bar chart to illustrate, for each period, the percentage of sample paths that a static logistic plan is activated by the Decision-tree policy at that period (%Go). Note that we do not provide this bar chart in Figure 6 for s-2SSP as the corresponding static plan is committed at $t = 1$ for all sample paths.

We can see from Figure 5 a clear shift of pattern in the activation timing decision by the Decision-tree policy from “early commitment” to “wait-and-activate” as the flow capacity level increases: when the flow capacity is at the lowest, it is crucial to activate at $t = 1$ to ensure a sufficient prepositioned inventory level; when the flow capacity is at the highest, it is sufficient to wait until the last period where the logistics operation is allowed (and perform all the logistics operation only in that period); when the flow capacity is between the two extreme cases, we see that the timing of activation varies based on the realization of uncertainty in each different sample path, resulting in a distribution over multiple periods. For MSP, we see a similar pattern in the periods of procurement from Figure 4: when the flow capacity level is low, the procurement is performed at its full capacity starting from the beginning of the planning horizon; when the flow capacity level increases, the procurement activity is spread across several consecutive periods closer to the landfall time, leveraging the adaptability to new information gained through waiting. For s-2SSP, we see that when the flow capacity level is higher, in which cases it is unnecessary to make an early commitment, the amount of procurement given by the s-2SSP solution appears to be larger than MSP and Decision-tree solutions, leading to its inferior performance in these cases. This is attributed to the lack of information at the time when the static plan is constructed at $t = 1$, where a larger amount of procurement helps mitigate a wide range of scenarios, and the procurement amount in the solutions given by the adaptive policies can be adapted according to the revealed information along the sample path.

We provide a heat map to further highlight the adaptability associated with MSP in Figure 7, where we plot the average amount of inventory over time at each intensity state (taken from all sample paths from the out-of-sample experiment that visit the corresponding state at the corresponding time) for each different flow capacity levels. The areas masked by white correspond to the pairs of

intensity state and time that are not reachable from the initial MC state ξ_1 . We see from Figure 7 that at each period, especially the ones closer to the terminal stage, states associated with a higher intensity level correspond to a higher (average) inventory level. This illustrates the adaptive behavior of the prepositioning operation from the MSP policy to address the demand uncertainty at landfall, which is more likely to be high at higher intensity states, and vice versa. On the other hand, when the flow capacity level is low (0.25), we do not observe this adaptability due to the limited flow capacity, which coincides with the observation in Figure 4, explaining the similar performance between MSP and s-2SSP in this case.




Figure 7: Behavior of MSP solution for the case study instances with deterministic landfall time under varying flow capacity levels with the safety threshold-based flow restrictions.




Figure 8: Behavior of Decision-tree solution for the case study instances with deterministic landfall time under varying flow capacity levels.

In Figure 8, we illustrate the optimal activation decisions $g(\xi_t)$ for all MC states ξ_t reachable from ξ_1 , obtained by the Decision-tree approach. Figure 8 can also be viewed as a look-up table that prescribes the “wait vs. activate” decision for each (intensity, time) pair. We note that such a look-up table can be extended to prescribe the “wait vs. activate” for each (intensity, location, time) combination. However, for illustrative purposes, we focus on the proportion of states that correspond

to each (intensity, time) pair that have an “activate” decision. From Figure 8, we clearly see the practicality of the solution provided by the Decision-tree approach in guiding decision-makers on when to activate a static plan constructed at that time, if no activation has been made yet. When the flow capacity level is 0.25, we see that the activation is made right at the initial state, making it identical to the solution provided by s-2SSP – again, this is due to the very limited flow capacity in this case. When the flow capacity level is 4, we see that the activation behavior is almost identical to WS, which is optimal due to the high flow capacity. When the flow capacity is 1, we see that the proportion of activation increases over time for each state, and increases as the intensity gets higher for each period, which demonstrates the trade-off between the value of timing in prepositioning commodities under a limited flow capacity and the value of information gained by waiting.

τ_{safe}	PI	MSP	s-2SSP	WS	Decision-tree
	$\hat{z}$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$
1	926.8K	(2696.3K, 190%)	(3834.3K, 313%)	(11649.2K, 1156%)	(3069.4K, 231%)
2	957.7K	(3100.4K, 223%)	(3962.3K, 313%)	(11735.8K, 1125%)	(3467.3K, 262%)
3	988.6K	(3466.9K, 250%)	(4090.2K, 313%)	(11771.4K, 1090%)	(3629.6K, 267%)
4	1019.5K	(3667.5K, 259%)	(4218.2K, 313%)	(11815.0K, 1058%)	(3723.8K, 265%)

Table 9: Performances of the PI solution and decision policies given by MSP, s-2SSP, decision-tree, and WS for the case study instances with deterministic landfall time and safety threshold-based cost structure under various safety threshold values τ_{safe} .

In Appendix C.2, we present the sensitivity analysis for the safety threshold τ_{safe} and logistics cost parameters for the case study instance.

Practical insights from the sensitivity analyses. Practically speaking, the flow capacity level depends on the logistics system’s transportation capability, such as the size and efficiency of the fleets of vehicles. Through the sensitivity analysis on the flow capacity levels, we observe that when the flow capacity level is low, the value of adaptability is relatively low as well, and it might be the most effective to increase the flow capacity level to reduce the overall cost. When the flow capacity level is sufficiently high, policies with partial adaptability (wait-and-activate) yield similar performance as the fully adaptive policies, in which case the operational adaptability is less critical. In all other cases, i.e., with an intermediate flow capacity level, we see from our results on the case study instance that greater benefit comes from improving the adaptability than increasing the transportation capability. This observation offers valuable insights to decision-makers on making long-term investments in transportation capability, such as increasing the fleet size, versus improving the adaptability and flexibility in the logistics operation, under different circumstances. Similarly, through the sensitivity analysis on the safety threshold (shown in the Appendix), we see a higher value of adaptability with a smaller safety threshold value. This suggests that the benefit of additional operational adaptability will be amplified with additional safety measures (to decrease the safety threshold).

Our numerical experiment and sensitivity analysis results provide useful information on the value of the adaptive decision policy obtained via the MSP model in various situations. However, we recognize that implementing a fully adaptive decision policy in a practical sense usually involves higher costs than a static decision policy like s-2SSP or a partially adaptive policy like the wait-and-activate approach given by Decision-tree, due to their higher requirements on data communication and/or the coordination between different logistics units. In this regard, our results provide valuable

guidance to the decision-maker on whether or not it is worthwhile to make investments to support adaptive logistics decision-making in various situations.

7. Conclusion

In this paper, we considered the problem of prepositioning relief commodities prior to impending hurricane landfall, referred to as the HDRLP problem. We assume that the evolution of the hurricane attributes is modeled as a discrete-time MC model, and the demand for relief commodities at DPs can be derived from the hurricane’s intensity and location at the time of its landfall. The solution methodology that we have presented in this work is capable of handling the case when the number of stages in the planning horizon is random due to the uncertainty about the hurricane’s landfall time. To solve the HDRLP problem, we have proposed a novel fully-adaptive MSP model that gives an optimal decision policy for the decision maker to make adaptive logistics operational decisions in every period of the planning horizon. We have also considered alternative decision policies: a static policy given by a 2SSP model, a rolling-horizon online decision policy that employs the 2SSP model as the look-ahead model in each period of the planning horizon, and partially adaptive policies including a decision-tree based policy, and a wait-and-see policy. We have conducted a set of numerical experiments to validate the proposed approaches on both synthetic instances and realistic instances inspired by Hurricane Florence (in 2018). Our numerical results show that in general, the performances of different policies clearly depend on the levels of adaptability inherited in different approaches. Furthermore, our sensitivity analyses demonstrate how the performances of the different policies depend on several key parameters in the model, including the flow capacity level, the cost surge parameter, and the safety threshold value. These results offer valuable insights to practitioners on how to choose appropriate decision policies under different operational constraints in practice.

We have identified several key limitations in our work that motivate future research directions accordingly. First, several simplifications made in this paper can be lifted to incorporate a richer set of real-world aspects associated with the hurricane’s attributes, the hazard/damage that the hurricane might cause to study regions of different geographical and socioeconomic characteristics, and the demand profiles for relief items resulting from the damages. Second, it would be valuable to incorporate the decision maker’s adaptive decisions regarding the selection of the SPs and the timings of their activation and deactivation, as these facility location decisions typically make up a significant part of the overall logistics cost. Third, as is true in many real-world applications, the probability distributions of the random variables involved in the problem are almost impossible to know exactly, and it would be of interest to study risk-averse MSP models and data-driven distributionally robust MSP models for the HDRLP problem. In addition, the proposed models and approaches assume that a fixed MC model is used to characterize the evolution of the hurricane over the planning horizon, e.g., in the case study, the MC model is constructed by overlaying sampled forecast errors on a point forecast made at $t = 1$. As the point forecasts themselves are updated over time, the MC models need to be updated accordingly. It will be interesting to investigate how the proposed approaches can be extended to address this challenge. Finally, it will be interesting to incorporate some qualitative aspects of the HDRLP problem such as the trade-off between the value and the executability of the adaptive decision policies in their practical implementations.

Acknowledgement The authors acknowledge partial support by the National Science Foundation [Grant CMMI 2045744]. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Alden, J. M., Smith, R. L., 1992. Rolling horizon procedures in nonhomogeneous markov decision processes. *Operations Research* 40 (3-supplement-2), S183–S194.
- Alem, D., Clark, A., Moreno, A., 2016. Stochastic network models for logistics planning in disaster relief. *European Journal of Operational Research* 255 (1), 187–206.
- Baker, E. J., 1991. Hurricane evacuation behavior. *International Journal of Mass Emergencies & Disasters* 9 (2), 287–310.
- Balcik, B., Beamon, B. M., Smilowitz, K., 2008. Last mile distribution in humanitarian relief. *Journal of intelligent transportation systems* 12 (2), 51–63.
- Beatty, T. K., Lade, G. E., Shimshack, J., 2021. Hurricanes and gasoline price gouging. *Journal of the Association of Environmental and Resource Economists* 8 (2), 347–374.
- Bellman, R., 1966. Dynamic programming. *Science* 153 (3731), 34–37.
- Ben-Tal, A., Do Chung, B., Mandala, S. R., Yao, T., 2011. Robust optimization for emergency logistics planning: Risk mitigation in humanitarian relief supply chains. *Transportation research part B: methodological* 45 (8), 1177–1189.
- Bhattarai, S., Song, Y., 2023. Integrated hurricane relief logistics and evacuation planning under forecast uncertainty: A case study for hurricane florence. In: K. Babski-Reeves, B. Eksioglu, D. H. (Ed.), *Proceedings of the IISE Annual Conference & Expo 2023*.
- Birge, J. R., Louveaux, F., 2011. *Introduction to stochastic programming*. Springer Science & Business Media.
- Bosart, L. F., Bracken, W. E., Molinari, J., Velden, C. S., Black, P. G., 2000. Environmental influences on the rapid intensification of hurricane opal (1995) over the gulf of mexico. *Monthly Weather Review* 128 (2), 322–352.
- Brettschneider, B., 2008. Climatological hurricane landfall probability for the united states. *Journal of Applied Meteorology and Climatology* 47 (2), 704–716.
- Çelik, M., Ergun, Ö., Johnson, B., Keskinocak, P., Lorca, Á., Pekgün, P., Swann, J., 2012. Humanitarian logistics. In: *New directions in informatics, optimization, logistics, and production*. INFORMS, pp. 18–49.
- Chand, S., Hsu, V. N., Sethi, S., 2002. Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing & Service Operations Management* 4 (1), 25–43.

- Chang, Y., Song, Y., Eksioglu, B., 2022. A stochastic look-ahead approach for hurricane relief logistics operations planning under uncertainty. *Annals of Operations Research* 319, 1231–1263.
- Davidson, R. A., Nozick, L. K., Wachtendorf, T., Blanton, B., Colle, B., Kolar, R. L., DeYoung, S., Dresback, K. M., Yi, W., Yang, K., et al., 2020. An integrated scenario ensemble-based framework for hurricane evacuation modeling: Part 1—decision support system. *Risk analysis* 40 (1), 97–116.
- Duran, S., Gutierrez, M. A., Keskinocak, P., 2011. Pre-positioning of emergency items for care international. *Interfaces* 41 (3), 223–237.
- Ergun, O., Karakus, G., Keskinocak, P., Swann, J., Villarreal, M., 2010. Operations research to improve disaster supply chain management. *Encyclopedia of operations research and management science* 6, 3802–3810.
- Galindo, G., Batta, R., 2013. Prepositioning of supplies in preparation for a hurricane under potential destruction of prepositioned supplies. *Socio-Economic Planning Sciences* 47 (1), 20–37.
- Guigues, V., 2021. Multistage stochastic programs with a random number of stages: dynamic programming equations, solution methods, and application to portfolio selection. *Optimization Methods and Software* 36 (1), 211–236.
- Holguín-Veras, J., Jaller, M., 2012. Immediate resource requirements after hurricane katrina. *Natural Hazards Review* 13 (2), 117–131.
- Holguín-Veras, J., Jaller, M., Van Wassenhove, L. N., Pérez, N., Wachtendorf, T., 2012. On the unique features of post-disaster humanitarian logistics. *Journal of Operations Management* 30 (7-8), 494–506.
- Holguín-Veras, J., Pérez, N., Jaller, M., Van Wassenhove, L. N., Aros-Vera, F., 2013. On the appropriate objective function for post-disaster humanitarian logistics models. *Journal of Operations Management* 31 (5), 262–280.
- Horry County Emergency Management Division, 2024. Comprehensive emergency management plan. Last accessed: June 05, 2024.
URL <https://www.horrycountysc.gov/media/hwmpsvyk/2-2-logistics-plan-2024.pdf>
- Kirac, E., Milburn, A. B., 2018. A general framework for assessing the value of social data for disaster response logistics planning. *European Journal of Operational Research* 269 (2), 486–500.
- Knabb, R. D., Rhome, J. R., Brown, D. P., 2005. Tropical cyclone report: Hurricane katrina, 23-30 august 2005. National Hurricane Center.
- Li, X., Batta, R., Kwon, C., 2017. Effective and equitable supply of gasoline to impacted areas in the aftermath of a natural disaster. *Socio-Economic Planning Sciences* 57, 25–34.
- Lodree Jr, E. J., Taskin, S., 2009. Supply chain planning for hurricane response with wind speed information updates. *Computers & Operations Research* 36 (1), 2–15.
- Löhndorf, N., Shapiro, A., 2019. Modeling time-dependent randomness in stochastic dual dynamic programming. *European Journal of Operational Research* 273 (2), 650–661.

- Menzel, D. C., 2006. The katrina aftermath: A failure of federalism or leadership? *Public Administration Review* 66 (6), 808–812.
- Morrice, D. J., Cronin, P., Tanrisever, F., Butler, J. C., 2016. Supporting hurricane inventory management decisions with consumer demand estimates. *Journal of Operations Management* 45, 86–100.
- Murray-Tuite, P., Wolshon, B., 2013. Evacuation transportation modeling: An overview of research, development, and practice. *Transportation Research Part C: Emerging Technologies* 27, 25–45.
- National Hurricane Center, 2012. Description of forecast error file format. https://www.nhc.noaa.gov/verification/pdfs/Error_Tabulation_File_Format.pdf, last accessed: May 1st, 2024.
- National Hurricane Center, 2022. National hurricane center data archive. <https://www.nhc.noaa.gov/data/>, last accessed: January 15, 2022.
- National Hurricane Center and Central Pacific Hurricane Center, 2022. Saffir-simpson hurricane scale. <https://www.nhc.noaa.gov/aboutsshws.php>, last accessed: June 04, 2022.
- NOAA NCEI, 2022. Noaa national centers for environmental information (ncei) u.s. billion-dollar weather and climate disasters. <https://www.ncei.noaa.gov/access/billions/>, last accessed: November 24, 2022.
- Pacheco, G. G., Batta, R., 2016. Forecast-driven model for prepositioning supplies in preparation for a foreseen hurricane. *Journal of the Operational Research Society* 67 (1), 98–113.
- Paul, J. A., Zhang, M., 2019. Supply location and transportation planning for hurricanes: A two-stage stochastic programming framework. *European Journal of Operational Research* 274 (1), 108–125.
- Pérez-Rodríguez, N., Holguín-Veras, J., 2016. Inventory-allocation distribution models for postdisaster humanitarian logistics with explicit consideration of deprivation costs. *Transportation Science* 50 (4), 1261–1285.
- Rawls, C. G., Turnquist, M. A., 2010. Pre-positioning of emergency supplies for disaster response. *Transportation research part B: Methodological* 44 (4), 521–534.
- Regnier, E., 2008. Public evacuation decisions and hurricane track uncertainty. *Management Science* 54 (1), 16–28.
- Regnier, E., Harr, P. A., 2006. A dynamic decision model applied to hurricane landfall. *Weather and forecasting* 21 (5), 764–780.
- Regnier, E. D., 2020. What is six hours worth? the impact of lead time on tropical-storm preparation decisions. *Decision Analysis* 17 (1), 9–23.
- Rivera-Royero, D., Galindo, G., Yie-Pinedo, R., 2016. A dynamic model for disaster response considering prioritized demand points. *Socio-economic planning sciences* 55, 59–75.
- Sabbaghtorkan, M., Batta, R., He, Q., 2020. Prepositioning of assets and supplies in disaster operations management: Review and research gap identification. *European Journal of Operational Research* 284 (1), 1–19.

- Sadana, U., Chenreddy, A., Delage, E., Forel, A., Frejinger, E., Vidal, T., 2024. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*.
- Salmerón, J., Apte, A., 2010. Stochastic optimization for natural disaster asset prepositioning. *Production and operations management* 19 (5), 561–574.
- SCEMD, 2023. Attachment a logistics plan to the south carolina emergency operations plan. Last accessed: May 1st, 2024.
- Schott, T., Landsea, C., Hafele, G., Lorens, J., Taylor, A., Thurm, H., Ward, B., Willis, M., Zaleski, W., 2012. The saffir-simpson hurricane wind scale. NOAA/National Weather Service [Internet], 1–4.
- Sethi, S., Sorger, G., 1991. A theory of rolling horizon decision making. *Annals of Operations Research* 29 (1), 387–415.
- Shapiro, A., 2011. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research* 209 (1), 63–72.
- Shapiro, A., Dentcheva, D., Ruszczyński, A., 2009. *Lectures on stochastic programming: modeling and theory*. SIAM.
- The National Hurricane Center and Central Pacific Hurricane Center, 2022. Graphics archive. <https://www.nhc.noaa.gov/archive/2018/>, last accessed: June 14, 2023.
- Velasquez, G. A., Mayorga, M. E., Özaltın, O. Y., 2020. Prepositioning disaster relief supplies using robust optimization. *IIE Transactions* 52 (10), 1122–1140.
- Yang, H., Duque, D., Morton, D. P., 2022. Optimizing diesel fuel supply chain operations to mitigate power outages for hurricane relief. *IIE Transactions* 54 (10), 936–949.
- Yang, K., Davidson, R. A., Nozick, L. K., Blanton, B., Colle, B., 2017. Scenario-based hazard trees for depicting resolution of hurricane uncertainty over time. *Natural Hazards Review* 18 (3), 04017001.
- Yi, W., Özdamar, L., 2007. A dynamic logistics coordination model for evacuation and support in disaster response activities. *European journal of operational research* 179 (3), 1177–1193.
- Zhang, Y., Richter, A. R., Shanthikumar, J. G., Shen, Z.-J. M., 2022. Dynamic inventory relocation in disaster relief. *Production and Operations Management* 31 (3), 1052–1070.

Appendix A. Acronyms

Acronym	Definition
2SSP	Two-stage stochastic programming
CI	Confidence interval
Decision-tree	A decision-tree based solution approach for optimizing the activation policy for wait-and-activate type policies
DP	Demand point
FEMA	Federal Emergency Management Agency
HDRLP	Hurricane disaster relief logistics planning
MC	Markov chain
MDC	Main distribution center
MSP	Multi-stage stochastic programming
NHC	National Hurricane Center
PI	Perfect-information
rh-2SSP	Rolling-horizon two-stage stochastic programming approach
SCEMD	South Carolina Emergency Management Division
SSHWS	Saffir-Simpson hurricane wind scale
SP	Supply point
s-2SSP	Static policy by a two-stage stochastic programming model
WS	Wait-and-see policy

Table A.10: Definitions and notations.

Appendix B. Data Generation for Test Instances

Appendix B.1. Synthetic instance generation

In this section, we first provide details on how the network instances are generated randomly, discuss how the stochastic process of the evolution of the hurricane’s attributes is modeled, and then show how the hurricane’s attributes are mapped to the demand for relief commodities at different DPs for the set of synthetic instances used in our numerical experiments. Many of the problem parameters for this set of synthetic instances follow those given in the relevant literature (see Galindo and Batta (2013) and Pacheco and Batta (2016)).

We consider a study region similar to the one depicted in Figure B.9. This figure shows three example sample paths of $\{\xi_1^{\text{loc}}, \xi_2^{\text{loc}}, \dots\}$, where each location state ξ_t^{loc} is defined as a pair of coordinates $(\xi_{t,x}^{\text{loc}}, \xi_{t,y}^{\text{loc}})$ that represent the hurricane’s location at time t . The three sample paths depicted in Figure B.9 are sampled from our MC model, which leads to three different landfall times (we omit the intensity states on these sample paths in the plot):

- $(50, -600) \rightarrow (150, -500) \rightarrow (150, -400) \rightarrow (250, -300) \rightarrow (350, -200) \rightarrow (450, -100) \rightarrow (550, 0)$; Landfall time is $T = 7$.
- $(50, -600) \rightarrow (150, -500) \rightarrow (150, -400) \rightarrow (250, -300) \rightarrow (350, -200) \rightarrow (450, 0)$; Landfall time is $T = 6$.
- $(50, -600) \rightarrow (150, -500) \rightarrow (150, -400) \rightarrow (250, -300) \rightarrow (350, 0)$; Landfall time is $T = 5$.

Let $\ell_i = (\ell_{i,x}, \ell_{i,y})$ and $\ell_j = (\ell_{j,x}, \ell_{j,y})$ be the coordinates of an SP i and a DP j , respectively. In all instances, we independently sample the coordinates of the SPs and DPs randomly from a set of uniform distributions: $\ell_{i,x} \sim U(0, 700)$, $\ell_{i,y} \sim U(100, 200)$, $\ell_{j,x} \sim U(0, 700)$ and $\ell_{j,y} \sim U(0, 100)$. We assume that the MDC and SPs are outside of the girded area, meaning that they are not at risk of




Figure B.9: An example hurricane trajectory towards the study region. The units in both axes are in miles. The solid line indicates the realized hurricane trajectory and the dashed lines indicate examples of alternative trajectories that could have been realized.

damage by the hurricane. In all instances, we set the coordinates of the MDC to be $(350, 450)$.

In addition to the safety threshold based cost structure presented in the main paper, which is MC-state dependent, we also consider a simple alternative cost structure that is time-dependent only: the unit logistics costs are assumed to increase linearly over time by a scaling factor ν , regardless of the corresponding MC states. In particular, we set the unit cost $c_{ii',t}^b$ to as:

$$c_{ii',t}^b = \omega(1 + \nu(t - 1))\|\ell_i - \ell_{i'}\|_2.$$

Similarly, the unit cost of transporting relief commodities from SP $i \in I$ to DP $j \in J$ is defined as:

$$c_{ij}^a = \omega(1 + \nu(T - 1))\|\ell_i - \ell_j\|_2, \quad (\text{B.1})$$

where the base transportation cost $\omega = 0.0038$ (e.g., fuel cost per mile) and T is the landfall time. Moreover, we define the remaining parameters h_t, c_t^h, p and q as a multiple of a base cost $\beta = 5$ such that $c_{i,t}^h = 0.2\beta, \forall i \in I, t = 1, 2, \dots, T_{\max}$, $h_t = \beta(1 + \nu(t - 1)), \forall t = 1, 2, \dots, T_{\max}$, $p = 50\beta$, and $q = -0.2\beta$.

The capacities of the SPs are generated randomly according to a uniform distribution $U(0.05\bar{d} \cdot |J|/|I|, 0.5\bar{d} \cdot |J|/|I|)$, where $\bar{d}$ is the total maximum possible demand for relief commodities at any DP $j \in J$. In our test instances, we set $\bar{d} = 400$. The rational behind this setup is that the capacity of each SP $i \in I$ is proportional to the maximum possible demand that can occur at a given DP, increasing in the total number of DPs in the study region, and decreasing in the total number of SPs in the study region that could supply different DPs with the relief commodities.

Modeling the evolution of the hurricane's attributes We define the state space Ξ of the MC model as $\Xi = \Xi^{\text{int}} \times \Xi^{\text{loc}}$. We assume that the stochastic processes associated with the hurricane's intensity and the hurricane's location (x - and y -coordinate) are independent. In our test instances, we consider $\Xi^{\text{int}} = \{0, 1, \dots, 5\}$, where 0 corresponds to the state where the hurricane's intensity is not high enough to generate any demand, and states 1 to 5 correspond to the five categories pertaining to the Saffir-Simpson hurricane wind scale. The one-step transition probability matrix for the MC associated with the intensity level is given by:

$$\mathbf{P}_{\xi^{\text{int}}} = \begin{bmatrix} 1.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.11 & 0.83 & 0.06 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.15 & 0.60 & 0.25 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.04 & 0.68 & 0.28 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.18 & 0.79 & 0.03 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.50 & 0.50 \end{bmatrix} \quad (\text{B.2})$$

We assume that the state space for the hurricane's landfall location $\Xi^{\text{loc}} = \Xi_x^{\text{loc}} \times \Xi_y^{\text{loc}}$, where $\Xi_x^{\text{loc}} = \{50, 150, \dots, 650\}$ and $\Xi_y^{\text{loc}} = \{-600, -500, \dots, 0\}$. The one-step transition probability matrix for the MC associated with the hurricane's landfall location along the x -coordinate is given by:

$$\mathbf{P}_{\xi_x^{\text{loc}}} = \begin{bmatrix} 0.004 & 0.300 & 0.395 & 0.198 & 0.049 & 0.038 & 0.016 \\ 0.150 & 0.202 & 0.249 & 0.222 & 0.117 & 0.033 & 0.027 \\ 0.198 & 0.249 & 0.029 & 0.168 & 0.206 & 0.099 & 0.051 \\ 0.099 & 0.222 & 0.169 & 0.012 & 0.150 & 0.198 & 0.150 \\ 0.025 & 0.117 & 0.206 & 0.15 & 0.004 & 0.150 & 0.348 \\ 0.019 & 0.033 & 0.098 & 0.198 & 0.150 & 0.004 & 0.498 \\ 0.008 & 0.019 & 0.025 & 0.098 & 0.198 & 0.150 & 0.502 \end{bmatrix} \quad (\text{B.3})$$

We note that $\mathbf{P}_{\xi_x^{\text{loc}}}$ and $\mathbf{P}_{\xi^{\text{int}}}$ are taken directly from Pacheco and Batta (2016). For the random landfall time case, the one-step transition probability matrix for the MC associated with the hurricane's landfall location along the y -coordinate is given by:

$$\mathbf{P}_{\xi_y^{\text{loc}}} = \begin{bmatrix} 0.0 & 0.6 & 0.3 & 0.1 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.6 & 0.3 & 0.1 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.6 & 0.3 & 0.1 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.6 & 0.3 & 0.1 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.6 & 0.4 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 1.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 1.0 \end{bmatrix}, \quad (\text{B.4})$$

which can be interpreted as having a random forward speed along the y -coordinate, so that the hurricane may move one or multiple steps closer towards the landfall at each period.

For the deterministic landfall time case, the one-step transition probability matrix for the MC

associated with the hurricane’s landfall location along the y -coordinate is given by:

$$\mathbf{P}_{\xi_y^{\text{loc}}} = \begin{bmatrix} 0.0 & 1.0 & 0.0 & 0.10 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 1.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 1.0 \end{bmatrix}, \quad (\text{B.5})$$

which can be interpreted as moving one step closer towards the landfall at each period in a deterministic fashion.

Consequently, the stochastic process associated with the hurricane’s evolution is characterized by an MC whose transition probability matrix $\mathbf{P} := \mathbf{P}_{\xi^{\text{int}}} \times \mathbf{P}_{\xi_x^{\text{loc}}} \times \mathbf{P}_{\xi_y^{\text{loc}}}$. We note that state 0 in Ξ_y^{loc} is an absorbing state corresponding to the landfall event. Accordingly, we have the following sets of transient states and absorbing states in S :

- Set of transient states $\Xi^{\text{trans}} := \{\xi_t = (\xi_t^{\text{int}}, \xi_{t,x}^{\text{loc}}, \xi_{t,y}^{\text{loc}}) \mid \xi_{t,y}^{\text{loc}} < 0\}$, and
- Set of absorbing states $\Xi^{\text{absor}} := \{\xi_t = (\xi_t^{\text{int}}, \xi_{t,x}^{\text{loc}}, \xi_{t,y}^{\text{loc}}) \mid \xi_{t,y}^{\text{loc}} = 0\}$.

For our experiments, we assume that the initial state is $\xi_1 = (1, 150, -600)$. We define the demand function $d_j(\xi_t)$ following equation (17) in the main paper, where $D_{\max} = 300$ and $P_j = 400$, $\forall j \in J$.

Appendix B.2. Case study instance generation

In this section, we briefly describe the input data used to construct the logistics network for the case study. Table B.11, B.12, and B.13 provide the data for the MDC and SPs, DPs, and the set of eligible SPs to serve DPs, respectively, which are collected from SCEMD (2023) as well as the authors’ personal communication with officials at the SCEMD.

Type	Facility	Address	Capacity
MDC	FEMA Warehouse	3780 Southside Industrial Pkwy, Atlanta, GA 30354	∞
SP	Florence Airport	2100 Terminal Dr, Florence, SC 29501	230000
SP	North AAF	270 Slab Landing Rd, North, SC 29112	480000
SP	Charleston Convention Center	5001 Coliseum Dr, North Charleston, SC 29418	120000
SP	I-26E (Harleyville)	495 Judge St, Harleyville, SC 29448	96000
SP	I-95S (Hendersonville)	2410 State Rd S-15-28, Walterboro, SC 29488	150000
SP	Farmers Market	3483 Charleston Hwy, West Columbia, SC 29172	120000
SP	SCEMD Warehouse	100 Tidewater Dr, Winnsboro, SC 29180	50000

Table B.11: Detailed information about the MDC and SPs used in the case study instance.

The capacity for each SP depends on the lot size of each facility, which is estimated using the maximum number of trucks that the facility can support, the maximum number of pallets of commodity items that each truck can carry, and the number of demand units that each pallets of commodity provide. All of these numbers and estimations were obtained from SCEMD (2023) as well as the authors’ personal communications with logistics managers at the South Carolina Emergency Management Division (SCEMD). Following this idea, we specialize the flow capacity constraint to

Facility	County	Address	Demand
Andrews Elementary School	Georgetown	13072 County Line Road, Andrews, SC 29510	12681
Loris High School	Horry	301 Loris Lion Road, Loris, SC 29569	70206
Cane Bay High School	Berkeley	1624 State Rd. Summerville, SC 29483	45972
Stall High School	Charleston	3625 Ashley Phosphate Road, North Charleston, SC 29418	80635
Fort Dorchester High School	Dorchester	8500 Patriot Boulevard, N Charleston, SC 29420	31457
Colleton County High School	Colleton	150 Cougar Nation Dr. Walterboro, SC 29488	7682
Ridgeland School Complex	Jasper	250 Jaguar Trail, Ridgeland, SC 29936	6020

Table B.12: Detailed information about DPs used in the case study instance.

Type	Facility	Counties Served
SP	Florence Airport	Georgetown, Horry, Berkeley, Charleston, Dorchester
SP	North AAF	Berkeley, Charleston, Dorchester, Colleton, Jasper
SP	Charleston Convention Center	Colleton, Jasper
SP	I-26E (Harleyville)	Berkeley, Charleston, Dorchester, Colleton, Jasper
SP	I-95S (Hendersonville)	Berkeley, Charleston, Dorchester, Colleton, Jasper
SP	Farmers Market	Berkeley, Charleston, Dorchester
SP	SCEMD Warehouse	All

Table B.13: Detailed information about the eligible SPs to serve each DP for the case study instance.

limit the total amount of flow value in a period to be bounded above by a given threshold $\bar{f}$, i.e.,

$$\sum_{i \in \{0\} \cup I, i' \in I} f_{ii',t} \leq \bar{f}, \quad \forall t = 1, 2, \dots, T.$$

In our experiments we set the baseline value to be $\bar{f} = 34000$, which is estimated by assuming a fleet size of 20 trucks for conducting the logistics operation per period, along with the aforementioned estimation on the maximum number of pallets of commodity items that each truck can carry, and the number of demand units that each pallet of commodities provide. We conduct a sensitivity on the flow capacity $\bar{f}$ to demonstrate its impact to the solution behaviors for different policies.

In the case study, we define the demand function $d_j(\xi_t)$ following equation (17) in the main paper. We set P_j to be 20% of the population of the corresponding county served by DP $j \in J$, which is obtained from the US Census 2020 data. Additionally, we set $D_{\max} = 300$. The remaining parameters are set in the same way as the synthetic instances, and since we do not find reliable sources to estimate these parameters, we perform sensitivity analyses on these parameters. We use the safety-threshold based cost structure for the case study instances, and given a state $\xi_t \in \Xi$ at time t , the unit cost transporting relief commodities from node $i \in \{0\} \cup I$ to node $i' \in I$ is defined as:

$$c_{ii',t}^b(\xi_t) = \begin{cases} \omega \times \text{dist}(\ell_i, \ell_{i'}) & \text{if } \bar{T}(\xi_t) > \tau_{\text{safe}} \\ \omega \times \text{dist}(\ell_i, \ell_{i'}) \times \nu & \text{otherwise,} \end{cases}$$

where the base transportation cost $\omega = 0.038$, $\nu \geq 1$ is the cost surge parameter and $\text{dist}(\cdot, \cdot)$ is a distance function that measures the distance between the two locations, which we obtained from Google Map API.

Similarly, the unit cost of transporting relief commodities from SP $i \in I$ to DP $j \in J$ is defined as:

$$c_{ij}^a = \omega \times \|\ell_i - \ell_j\|_2$$

The unit procurement cost h_t is set to be the base cost $\beta = 22.4$ at all periods $t = 1, 2, \dots, T_{\max}$. The base cost is estimated based on the market price for commodities obtained in a typical disaster relief kit, including an 8-pound ice bag, three liters of water, and two military-grade MREs (Horry County Emergency Management Division, 2024). The unit inventory cost c_t^h is set to be a constant: $c_{i,t}^h = 0.05 \times \beta, \forall i \in I, t = 1, 2, \dots, T_{\max}$. The remaining parameters are set in a similar way as the synthetic instances. We conduct a sensitivity analysis on these parameters to investigate how they affect the solution behavior.

Appendix C. Additional Numerical Results

Appendix C.1. Additional numerical results on the synthetic instances

ν	$ I $	$ J $	PI		MSP				Decision-tree		rh-2SSP	s-2SSP	
			$\hat{z}$	CI	LB	UB	time	gap	time	gap	gap	time	gap
0	3	10	710.8	93.0	1693.2	1646.4±147.5	20.1	131%	2.8	131%	131%	0.0	753%
		20	1421.6	179.0	3175.3	3061.8±265.8	22.6	115%	4.6	115%	115%	0.0	696%
		30	2123.2	267.0	4859.6	4696.2±407.4	24.3	121%	6.2	121%	121%	0.1	707%
	6	10	1093.4	286.9	1953.1	1965.2±299.4	28.1	79%	4.2	79%	79%	0.0	462%
		20	1645.7	294.4	3315.3	3202.7±337.8	31.1	94%	6.7	94%	94%	0.1	603%
		30	2542.8	477.8	5205.7	5017.4±548.5	38.1	97%	9.2	97%	97%	0.1	581%
	9	10	908.8	202.0	1807.3	1801.0±224.0	29.6	98%	5.6	98%	98%	0.0	579%
		20	1458.4	194.4	3249.1	3058.3±267.0	39.7	109%	8.8	109%	109%	0.1	700%
		30	2243.7	321.6	5021.9	4788.5±437.7	48.7	113%	12.0	113%	113%	0.1	672%
0.5	3	10	1038.2	139.3	4047.6	3923.3±342.2	40.3	277%	3.3	343%	464%	0.0	597%
		20	2065.0	262.6	7590.1	7391.7±614.1	42.9	257%	5.5	324%	488%	0.0	573%
		30	3031.4	382.3	11504.0	11183.2±988.9	50.6	268%	7.0	333%	506%	0.1	582%
	6	10	1396.8	325.7	4192.8	4118.4±431.4	43.5	194%	4.5	239%	333%	0.0	418%
		20	2244.2	373.3	7629.6	7434.9±647.9	58.6	231%	8.0	290%	430%	0.1	522%
		30	3388.8	580.8	11642.1	11282.4±1069.7	53.2	232%	10.8	293%	418%	0.1	516%
	9	10	1244.5	248.7	4130.3	4032.6±387.1	73.3	224%	6.3	273%	381%	0.1	482%
		20	2102.1	291.9	7662.6	7422.9±632.2	92.8	253%	11.1	318%	471%	0.1	563%
		30	3149.7	452.0	11639.2	11310.1±1075.8	75.8	259%	15.1	322%	475%	0.1	560%
5	3	10	2874.1	420.6	8734.3	8876.4±1148.1	25.7	208%	4.5	208%	207%	0.0	208%
		20	5616.1	745.6	16672.6	17138.5±1928.9	26.3	205%	7.7	207%	204%	0.1	207%
		30	7829.0	1008.8	24343.3	25016.9±2765.2	19.2	219%	11.1	223%	219%	0.1	223%
	6	10	3033.8	596.1	8522.5	8648.0±1160.1	19.0	185%	7.0	184%	184%	0.0	184%
		20	5400.5	875.4	16214.5	16657.5±1943.1	25.8	208%	13.5	210%	207%	0.1	210%
		30	7639.1	1212.0	23723.9	24397.3±2781.7	23.9	219%	19.0	223%	219%	0.1	223%
	9	10	3162.4	593.3	8671.5	8683.3±1191.9	25.1	174%	9.2	176%	175%	0.1	176%
		20	5649.4	913.9	16463.3	16734.2±1997.2	54.0	196%	17.1	200%	196%	0.1	200%
		30	7907.0	1259.5	24019.3	24684.5±2846.4	30.2	212%	25.5	215%	212%	0.2	215%

Table C.14: Performances of the PI solution and decision policies given by MSP, Decision-tree, rh-2SSP, and s-2SSP for synthetic instances, deterministic landfall time case, and the time-dependent cost structure.

In this section, we present performance summaries for the PI solution and different decision policies for the synthetic instances under the time-dependent cost structure described in Appendix B.1. From Table C.14 and C.15 and Figure C.10 and C.11, we see similar patterns in the performance of different approaches as the safety threshold based cost structure presented in the main paper: MSP yields the best performance, followed by rh-2SSP, Decision-tree, and finally s-2SSP. In terms of computational performance, again we see that the random landfall time case is computationally more demanding than the deterministic landfall time case.

ν	$ I $	$ J $	PI		MSP				Decision-tree		rh-2SSP	s-2SSP	
			$\hat{z}$	CI	LB	UB	time	gap	time	gap	gap	time	gap
0	3	10	575.3	65.2	1755.0	1745.0±182.7	138.7	203%	55.4	395%	238%	0.7	666%
		20	1154.7	124.5	3324.1	3332.0±293.7	238.1	188%	68.0	375%	222%	0.5	648%
		30	1718.5	185.0	5109.5	5123.1±437.5	216.7	198%	76.1	371%	234%	0.5	680%
	6	10	639.3	129.1	1855.8	1797.3±210.8	222.9	181%	121.1	346%	210%	1.1	601%
		20	1176.2	153.0	3370.0	3334.6±302.4	179.3	183%	148.0	365%	217%	1.2	647%
		30	1782.0	253.7	5241.1	5152.1±461.6	277.5	189%	159.5	353%	224%	1.3	664%
	9	10	602.0	95.0	1798.9	1767.2±193.7	447.7	193%	229.0	373%	226%	2.1	641%
		20	1153.1	128.3	3348.9	3330.0±293.8	219.4	188%	278.1	376%	223%	2.3	657%
		30	1732.4	201.9	5173.1	5143.2±449.4	287.8	196%	296.4	368%	233%	2.4	679%
0.5	3	10	758.2	89.2	2987.0	3040.4±256.2	203.9	301%	62.5	542%	488%	0.7	544%
		20	1516.3	167.4	5599.8	5561.5±437.4	215.1	266%	82.2	514%	443%	0.9	517%
		30	2226.1	243.2	8526.1	8639.2±713.7	178.8	288%	88.4	525%	487%	0.9	552%
	6	10	804.3	149.7	3038.7	3039.3±271.0	198.5	277%	145.0	500%	452%	2.2	507%
		20	1497.4	193.9	5594.6	5520.1±429.9	287.2	268%	185.6	521%	449%	2.6	525%
		30	2232.5	304.7	8553.7	8578.7±710.8	221.0	284%	196.5	522%	485%	2.9	551%
	9	10	779.8	119.7	3015.0	3036.3±259.1	313.6	289%	268.9	520%	472%	3.3	528%
		20	1490.4	175.2	5613.2	5551.0±436.2	430.6	272%	334.6	523%	455%	5.0	529%
		30	2201.3	263.7	8563.1	8630.1±717.6	499.1	292%	364.0	532%	495%	4.9	561%
5	3	10	1810.9	239.5	6219.5	6011.9±740.5	176.5	231%	76.8	227%	228%	0.9	227%
		20	3570.0	421.5	11909.7	11486.0±1286.6	121.6	221%	95.7	219%	223%	1.1	219%
		30	4996.1	567.3	17626.3	17208.1±2214.0	183.5	244%	114.1	246%	248%	1.1	246%
	6	10	1698.4	295.0	6048.4	5790.9±731.5	320.8	240%	191.4	238%	239%	1.9	238%
		20	3189.1	448.0	11549.5	11085.3±1277.5	570.5	247%	268.8	245%	250%	2.7	245%
		30	4484.0	612.4	17179.5	16539.9±2166.1	230.4	268%	295.1	274%	275%	2.6	274%
	9	10	1787.0	297.3	6118.9	5826.1±740.4	195.1	226%	352.7	225%	225%	2.9	225%
		20	3318.8	474.3	11632.7	11168.2±1293.1	261.8	236%	482.4	234%	237%	5.1	234%
		30	4599.9	637.8	17279.5	16531.8±2165.8	252.8	259%	529.7	266%	268%	4.7	266%

Table C.15: Performances of the PI solution and decision policies given by MSP, Decision-tree, rh-2SSP, and s-2SSP for synthetic instances, random landfall time case, and the time-dependent cost structure.




Figure C.10: Performances of the PI solution and decision policies given by MSP, Decision-tree, rh-2SSP, and s-2SSP for synthetic instances, deterministic landfall time case, and the time-dependent cost structure.




Figure C.11: Performances of the PI solution and decision policies given by MSP, Decision-tree, rh-2SSP, and s-2SSP for synthetic instances, random landfall time case, and the time-dependent cost structure.

Additionally, we observe that the differences between different approaches decrease as ν gets larger. Intuitively, this is because when the value of ν is relatively small, the cost of adaptability is small and for adaptive policies (MSP, Decision-tree, and rh-2SSP), it encourages a wait-and-see behavior of delaying the prepositioning of relief commodities until the uncertainty in the hurricane's attributes is significantly reduced. On the other hand, the static policy s-2SSP is unable to leverage this cost structure, leading to its poor performance in this case. As ν increases, the cost of waiting outweighs the benefit of adaptability, and it encourages the behavior of early commitment from solutions obtained by different approaches.

Comparing between deterministic and random landfall time cases, we see that the behaviors of different policies are similar in general, except that: (i) when $\nu = 0$, MSP and rh-2SSP show significant advantage over Decision-tree on the random landfall time case, while all three approaches exhibit identical performances in the deterministic landfall time case; (ii) rh-2SSP outperforms Decision-tree in the random landfall time case in contrast to the deterministic landfall time case. This underscores the value of adaptability especially in the random landfall time case.

Appendix C.2. Additional numerical results on the case study instances

In this section, we present additional numerical results on the case study instances: (i) solution-level illustration for the flow capacity sensitivity analysis on case study instances with random landfall time; (ii) sensitivity analysis on the safety threshold parameter; (iii) sensitivity analyses on the performance of MSP solution under various cost factors, including the inventory cost, unmet demand penalty cost, and transportation (flow) cost.

Solution-level illustration for the flow capacity sensitivity analysis on case study instances with random landfall time. We provide solution-level illustrations in Figure C.12 to Figure C.15 for the sensitivity analysis on the flow capacity level, based on instances with random landfall time, similar to those presented for the deterministic landfall time case in the main paper. We observe that the solution behaviors mostly follow the patterns identified for the deterministic

landfall time case, except the slight increase on the number of periods where procurement operations are conducted (%Procurement) from MSP and the variability on when the plan commitment is made (%Go) from Decision-tree, which is expected due to the variability in the landfall time and the associated demand realization.




Figure C.12: Behavior of MSP solution for the case study instances with random landfall time under varying flow capacity levels.




Figure C.13: Behavior of Decision-tree solution for the case study instances with random landfall time under varying flow capacity levels.




Figure C.14: Behavior of the s-2SSP solution for the case study instances with random landfall time under varying flow capacity levels.




Figure C.15: Behavior of MSP solution for the case study instances with random landfall time under varying flow capacity levels with the safety threshold based flow restrictions.

Sensitivity analysis for the safety threshold τ_{safe} . We conduct a sensitivity analysis on the safety threshold value τ_{safe} and investigate its impact on the performances of different policies. Table C.16 and Table C.17 present performance summary results for the deterministic landfall time case and random landfall time case, respectively. Figure C.16, C.17, and C.18 illustrate the solution-level behaviors of different policies for the deterministic landfall time case, and Figure C.19 and C.20 illustrate the solution-level behaviors of the MSP policy for the random landfall time case.

From Table C.16, we see that as the safety threshold value τ_{safe} increases from 1 to 4, the performance of MSP and Decision-tree is more negatively affected than that of s-2SSP. Considering that we set $\nu = \infty$ in this sensitivity analysis, larger values of τ_{safe} force the logistics operation to be activated earlier, which in turn reduce the value of adaptability, as there is less time to adapt to the new information. However, since s-2SSP is a static policy that commits to a plan at the beginning of the planning horizon, it is less affected by the safety threshold value. The performance of WS is not significantly affected by the increased τ_{safe} either, since the bottleneck of its performance lies in

τ_{safe}	PI	MSP	s-2SSP	WS	Decision-tree
	$\hat{z}$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$
1	926.8K	(2696.3K, 190%)	(3834.3K, 313%)	(11649.2K, 1156%)	(3069.4K, 231%)
2	957.7K	(3100.4K, 223%)	(3962.3K, 313%)	(11735.8K, 1125%)	(3467.3K, 262%)
3	988.6K	(3466.9K, 250%)	(4090.2K, 313%)	(11771.4K, 1090%)	(3629.6K, 267%)
4	1019.5K	(3667.5K, 259%)	(4218.2K, 313%)	(11815.0K, 1058%)	(3723.8K, 265%)

Table C.16: Performances of the PI solution and decision policies given by MSP, s-2SSP, decision-tree, and WS for the case study instances with deterministic landfall time and safety threshold-based cost structure under various safety threshold values τ_{safe} .




Figure C.16: Behavior of MSP solution for the case study instances with deterministic landfall time under various safety threshold values τ_{safe} .

the limited flow capacity, as we discussed earlier. As we observe from Figure C.16, the procurement activity shifts to earlier periods due to the increase of the safety threshold value. As a result, the (average) total amount of procurement increases due to the lack of information during these earlier periods, leading to worse performance.




Figure C.17: Behavior of MSP solution for the case study instances with deterministic landfall time under various safety threshold values.




Figure C.18: Behavior of Decision-tree solution for the case study instances with deterministic landfall time under various safety threshold values.

In Figure C.17 we provide a heat map on the average amount of inventory over time at each intensity state for each different safety threshold value. From Figure C.17, we can see the impact of the safety threshold on the inventory (prepositioning) behavior associated with the MSP policy. As the safety threshold increases: (i) prepositioning is conducted at earlier stages, especially for higher intensity states; (ii) the prepositioning activities are concentrated at fewer periods as the overall timeline for prepositioning is shortened with a larger safety threshold. Similarly, in Figure C.18 we demonstrate the activation timing behavior corresponding to the Decision-tree approach by illustrating the “wait vs. activate” decision for each (intensity, time) pair under different safety threshold values. We see that the impact of adaptability diminishes as the safety threshold values get larger. In particular, higher safety threshold values make the activation timing more homogeneous across different intensity levels, except for the first few periods. This explains the observation made in Table C.16 that the performance gap between Decision-tree and s-2SSP gets smaller as the safety

threshold increases.

τ_{safe}	PI	MSP	s-2SSP	WS	Decision-tree
	$\hat{z}$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$	$(\hat{z}, \text{gap})$
1	788.5K	(2568.5K, 225%)	(3323.9K, 321%)	(10793.3K, 1268%)	(3095.6K, 292%)
2	815.3K	(3078.1K, 277%)	(3443.0K, 322%)	(10848.6K, 1230%)	(3413.0K, 318%)
3	842.0K	(3293.9K, 291%)	(3562.2K, 323%)	(10903.2K, 1194%)	(3573.9K, 324%)
4	868.7K	(3494.9K, 302%)	(3681.3K, 323%)	(10932.3K, 1158%)	(3704.7K, 326%)

Table C.17: Performances of the PI solution and decision policies given by MSP, s-2SSP, decision-tree, and WS for the case study instances with random landfall time under various safety threshold values τ_{safe} .




Figure C.19: Behavior of MSP solution for the case study instances with random landfall time under various safety threshold values τ_{safe} .

From Table C.17, Figure C.19 and Figure C.20, we observe that the behaviors of different policies in the random landfall time case are similar to the deterministic landfall time case.

Sensitivity analyses on various cost factors. We now present some sensitivity analysis results for the MSP solutions in terms of the average amount of inventory, salvage, and demand shortage across the set of out-of-sample scenarios under various cost factors. Specifically, Figure C.21 to Figure C.26 illustrate the sensitivity analysis results under various cost factors on the penalty cost for demand shortage, inventory cost, and transportation cost, respectively, for the deterministic and random landfall time cases.

From Figure C.21 and C.22, we see a clear trade-off in the inventory level to balance between the amount of salvage and the demand shortage as the penalty cost factor increases from 0.25 to 4. As the penalty parameter increases, a higher inventory level is maintain prior to the landfall stage to avoid the high penalty cost at the expense of an increased amount of salvage. From Figure C.23 and C.24, we see that the solution is not sensitive to the inventory cost factor, including the special case when the inventory cost is zero. From Figure C.25 and C.26, we see that the solution is not sensitive to the transportation cost factor either, with a slight decrease in both inventory level and




Figure C.20: Behavior of MSP solution for the case study instances with random landfall time under various safety threshold values τ_{safe} , with the safety threshold based flow restrictions.

salvage amount and a slight increase in the amount of unmet demand, as the transportation cost increases.




Figure C.21: Average amount of prepositioned commodities, salvaged commodities, and unmet demand under various penalty cost factors for the case study instance with deterministic landfall time.




Figure C.22: Average amount of prepositioned commodities, salvaged commodities, and unmet demand under various penalty cost factors for the case study instance with random landfall time.




Figure C.23: Average amount of prepositioned commodities, salvaged commodities, and unmet demand under various inventory cost factors for the case study instance with deterministic landfall time.




Figure C.24: Average amount of prepositioned commodities, salvaged commodities, and unmet demand under various inventory cost factors for the case study instance with random landfall time.




Figure C.25: Average amount of prepositioned commodities, salvaged commodities, and unmet demand under various transportation cost factors for the case study instance with deterministic landfall time.




Figure C.26: Average amount of prepositioned commodities, salvaged commodities, and unmet demand under various transportation cost factors for the case study instance with random landfall time.

Appendix D. Implementation Details

In this section, we describe our implementation details on the solution approaches for the stochastic optimization models, including the MSP and 2SSP models, along with the termination criteria used for solving these models. In all of our experiments, we use programming language Python (version 3.8), with commercial solver *Gurobi* (version 11). All of the numerical experiments are conducted on an iMac desktop computer with an Apple M1 chip, 8GB memory, and 8 cores.

Implementation details on the nested Benders algorithm for solving the MSP model

We use the nested Benders algorithm (see Algorithm 1) to obtain optimal offline decision policies associated with the MSP model. In principle, during the training process of the nested Benders algorithm, every iteration of Algorithm 1 can be carried over a set of multiple sample paths. That is, in the forward step, one can collect trial points over multiple sample paths and, analogously, generate cuts along each of these sample paths during the backward pass. However, in all of our implementation, we observed that using one sample path per forward/backward step works the best. Furthermore, to accelerate the computations, we implement a *cut-sharing* strategy that is typically helpful for the application of nested Benders decomposition on problems with structured stage-wise dependence such as the MC structure. Indeed, the inequalities for approximating the expected cost-to-go functions at each stage can be generated by recombining the dual multipliers associated with the subproblems at the subsequent stage using different weights, according to the MC transition probability matrix. This allows us to generate valid inequalities to the approximation of the expected cost-to-go functions associated with MC states even if they are not encountered during the forward pass. Table D.18 and D.19 show the computational savings obtained by the cut-sharing strategy: to obtain solutions of similar quality (as indicated by LB and UB), much less computational time and number of iterations (forward and backward passes) are expended with the cut-sharing strategy. We also observe that the computational time saving is much more significant in the instances with random landfall time than the ones with deterministic landfall time.

We use the following termination criteria for the nested Benders decomposition algorithm: (i) *the maximum number of iterations* is set to be 10^5 ; (ii) *the time limit* is set to be three hours; and (iii) a *stability test* which keeps track of the progress of the lower bound for the optimal objective value: if this lower bound does not progress by more than $\epsilon = 10^{-4}$ (in the relative term) for more than 200 consecutive iterations, we terminate the algorithm.

Implementation details for solving the 2SSP model To solve the 2SSP model, we employ a standard implementation of the L-shaped method (Birge and Louveaux, 2011). The basic idea of the L-shaped method is to successively build an outer approximation of the second-stage value function $Q^{\xi_{[T]}}(\mathbf{x}, \mathbf{f})$ in (3) for each scenario (sample-path) $\xi_{[T]}$. To do this, we can generate cuts in the spirit of the one generated at the terminal stage, in the backward step of Algorithm 1. When the cost parameters are not state-dependent (e.g., they can just be time-dependent), the set of scenarios to be considered here can be enumerated based on the initial MC state ξ_1 , which corresponds to the absorbing states that are reachable from ξ_1 . This enumeration can be done efficiently offline through a tree-search algorithm such as the breadth-first search on the scenario tree. The same procedure is used in every roll of the r-2SSP approach and for implementing the decision-tree approach. When the cost parameters are state-dependent, such as the cost structure based on a safety threshold in

ν	$ I $	$ J $	MSP without cut-sharing				MSP with cut-sharing			
			LB	UB	time	# of iterations	LB	UB	time	# of iterations
0	3	10	1687.8	1643.9±148.1	93.2	5434	1702.6	1644.9±148.3	25.8	508
		20	3167.7	3060.0±268.7	104.3	4708	3187.2	3059.9±268.7	17.9	331
		30	4841.4	4735.4±421.8	107.7	3900	4891.4	4735.6±421.8	33.7	540
	6	10	1943.9	1964.5±299.8	147.5	6395	1965.8	1964.5±299.8	21.1	296
		20	3234.2	3202.2±340.3	88.8	3058	3329.1	3202.1±340.3	28.2	359
		30	5102.9	5056.9±559.7	123.4	3479	5232.4	5056.8±559.7	27.4	326
	9	10	1802.8	1800.3±224.7	149.3	5083	1819.4	1800.3±224.7	27.0	285
		20	3208.0	3056.2±270.0	188.3	5193	3263.4	3056.3±270.0	33.7	329
		30	4977.4	4829.8±451.8	147.9	3466	5048.8	4824.1±451.6	35.8	321
0.5	3	10	3956.1	3851.3±342.0	88.9	4591	3975.5	3859.5±345.6	55.9	913
		20	7413.9	7259.2±621.4	177.7	6903	7436.2	7240.6±619.4	38.6	632
		30	11253.2	11009.4±999.2	176.1	5876	11297.3	11036.0±1009.1	45.5	670
	6	10	4104.6	4053.9±433.3	132.2	4852	4122.5	4050.0±433.4	52.5	643
		20	7428.6	7287.1±652.5	167.1	4915	7479.7	7301.6±656.1	50.0	572
		30	11339.8	11131.5±1101.8	151.2	3811	11452.5	11113.3±1090.8	77.0	778
	9	10	4024.6	3979.9±392.1	181.0	5121	4059.7	3965.0±389.3	90.4	796
		20	7462.0	7275.0±635.0	309.0	6893	7512.3	7266.6±634.6	86.0	720
		30	11387.2	11125.3±1081.6	331.8	6420	11448.6	11126.8±1081.8	81.7	655
5	3	10	8320.1	8598.0±1050.0	78.1	3924	8473.8	8641.6±1087.8	43.1	711
		20	16090.8	16698.9±1907.5	160.7	6326	16182.9	16670.3±1902.4	27.9	467
		30	23357.3	24318.6±2795.2	205.1	6515	23593.7	24289.0±2778.0	20.3	320
	6	10	8210.8	8384.0±1095.3	174.7	6467	8264.0	8410.7±1100.2	68.7	758
		20	15645.5	16164.6±1919.4	346.3	9580	15733.4	16203.5±1916.9	43.4	480
		30	22770.4	23676.1±2783.9	230.7	5603	22980.3	23675.3±2792.6	33.8	368
	9	10	8247.5	8552.7±1160.5	148.4	4365	8413.4	8460.5±1136.6	46.3	427
		20	15758.8	16325.3±1990.8	221.6	5036	15978.2	16263.1±1968.3	60.2	502
		30	23116.7	24019.2±2880.2	368.6	6889	23274.7	23957.9±2856.7	44.8	366

Table D.18: Performances of the nested Benders decomposition algorithm on the synthetic instances with deterministic landfall time for solving the MSP model with and without the cut-sharing strategy.

ν	$ I $	$ J $	MSP without cut-sharing				MSP with cut-sharing			
			LB	UB	time	# of iterations	LB	UB	time	# of iterations
0	3	10	1729.8	1767.9±189.9	524.9	12120	1755.0	1745.0±182.7	138.7	916
		20	3290.6	3615.2±643.6	809.4	14411	3324.1	3332.0±293.7	238.1	1343
		30	5044.8	5605.7±1029.1	881.1	12801	5109.5	5123.1±437.5	216.7	1174
	6	10	1826.2	1796.9±210.9	809.5	13893	1855.8	1797.3±210.8	222.9	1012
		20	3328.2	3362.0±309.9	948.8	12547	3370.0	3334.6±302.4	179.3	779
		30	5115.3	5151.6±460.7	870.7	9787	5241.1	5152.1±461.6	277.5	1085
	9	10	1779.5	1766.1±193.6	988.3	13363	1798.9	1767.2±193.7	447.7	1436
		20	3308.1	3331.3±294.2	1225.2	13218	3348.9	3330.0±293.8	219.4	735
		30	5116.2	5139.5±449.8	1474.3	13401	5173.1	5143.2±449.4	287.8	895
0.5	3	10	2966.1	3034.8±254.9	626.2	13954	2987.0	3040.4±256.2	203.9	1269
		20	5566.2	5573.3±440.8	996.8	17083	5599.8	5561.5±437.4	215.1	1238
		30	8457.0	8863.6±889.2	1004.9	14412	8526.1	8639.2±713.7	178.8	994
	6	10	3006.8	3031.0±269.0	800.0	13120	3038.7	3039.3±271.0	198.5	885
		20	5547.7	5518.6±431.7	984.0	12842	5594.6	5520.1±429.9	287.2	1175
		30	8507.5	8578.8±709.7	1336.0	14368	8553.7	8578.7±710.8	221.0	878
	9	10	2984.7	3033.1±257.7	1039.6	13319	3015.0	3036.3±259.1	313.6	1071
		20	5573.4	5552.6±436.1	1588.8	15869	5613.2	5551.0±436.2	430.6	1296
		30	8491.6	8643.7±714.8	1474.1	12706	8563.1	8630.1±717.6	499.1	1395
5	3	10	6112.4	6066.3±749.7	593.3	13220	6219.5	6011.9±740.5	176.5	1117
		20	11710.1	11355.0±1266.5	953.3	16327	11909.7	11486.0±1286.6	121.6	759
		30	17507.7	16917.0±2141.7	1105.9	15981	17626.3	17208.1±2214.0	183.5	1005
	6	10	5914.4	5730.8±722.6	844.1	13686	6048.4	5790.9±731.5	320.8	1364
		20	11194.9	10964.6±1239.2	786.5	10311	11549.5	11085.3±1277.5	570.5	1985
		30	16877.4	16457.8±2178.2	1236.5	13461	17179.5	16539.9±2166.1	230.4	907
	9	10	6027.1	5884.5±751.1	1126.9	14891	6118.9	5826.1±740.4	195.1	704
		20	11529.7	11101.1±1276.8	1564.5	16436	11632.7	11168.2±1293.1	261.8	853
		30	17070.1	16537.0±2153.4	1859.9	16379	17279.5	16531.8±2165.8	252.8	770

Table D.19: Performances of the nested Benders decomposition algorithm on the synthetic instances with random landfall time for solving the MSP model with and without the cut-sharing strategy.

the random landfall time case, the number of sample paths can be exponentially large in general. We resort to sample average approximation approach and use a sample size of 100 for solving the 2SSP models for the synthetic instances, and use a sample size of 500 for the case study instances. We terminate the L-shaped method when the relative gap between the upper and the lower bounds falls below a stopping threshold $\epsilon = 10^{-4}$.

Appendix E. Supplementary figures




Figure E.27: Example of SPs and DP in South Carolina coastline (Left) and landfall locations shown in red (right).




Figure E.28: An illustration of the study region, the sample paths of hurricane trajectories, and the MC states (only location states are shown) for the case study. The black central line represents the point forecast provided at H-120 of Hurricane Florence (in 2018), and the red lines represent different trajectory sample paths.




Figure E.29: The NHC track forecast cone for Hurricane Florence (in 2018) (The National Hurricane Center and Central Pacific Hurricane Center, 2022).