

Markov Chain-based Policies for Multi-stage Stochastic Integer Linear Programming with an Application to Disaster Relief Logistics

Margarita P. Castro

Department of Industrial and Systems Engineering, Pontificia Universidad Católica de Chile, Santiago 7820436, Chile
margarita.castro@ing.puc.cl

Merve Bodur

Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario M5S 3GH, Canada,
bodur@mie.utoronto.ca

Yongjia Song

Department of Industrial Engineering, Clemson University, Clemson, South Carolina 29631, yongjis@clemson.edu

We introduce an aggregation framework to address multi-stage stochastic programs with mixed-integer state variables and continuous local variables (MSILPs). Our aggregation framework imposes additional structure to the integer state variables by leveraging the information of the underlying stochastic process, which is modeled as a Markov chain (MC). We demonstrate that the aggregated MSILP can be solved exactly via a branch-and-cut algorithm integrated with a variant of stochastic dual dynamic programming. To improve tractability, we propose to use this approach to obtain dual bounds. Moreover, we apply two-stage linear decision rule (2SLDR) approximations, in particular a new MC-based variant that we propose, to obtain high-quality decision policies with significantly reduced computational effort. We test the proposed methodologies in an MSILP model for hurricane disaster relief logistics planning. Our empirical evaluation compares the effectiveness of the various proposed approaches and analyzes the trade-offs between policy flexibility, solution quality, and computational effort. Specifically, the 2SLDR approximation yields provable high-quality solutions for our test instances supported by the proposed bounding procedure. We also extract valuable managerial insights from the solution behaviors exhibited by the underlying decision policies.

Key words: Multi-stage stochastic programming, Markov chain, linear decision rules, stochastic dual dynamic programming, disaster relief logistics

1. Introduction

Multi-stage stochastic integer linear programming (MSILP) problems form an important class of optimization models for sequential decision-making under uncertainty. These problems consider decisions at each stage of the sequential process, which consists of local variables—that only participate in a single stage locally—and state variables—that link multiple stages together. In this paper, we consider MSILP problems with continuous local variables and mixed-integer state variables, which arise in many applications such as hydro-power scheduling (Hjelmeland et al. 2018), unit commitment (Zou et al. 2018), and disaster relief logistics planning (see Section 4).

Despite the vast applicability of MSILP, these problems are computationally prohibitive to solve due to the non-convexity of the feasibility set caused by the integer state variables. Existing solution methods utilize the stage-wise decomposition framework, such as the stochastic dual dynamic programming (SDDP) algorithm, with convexification procedures for the expected cost-to-go functions, which require certain limiting assumptions such as pure binary state variables (Zou et al. 2019) and Lipschitz continuous cost-to-go functions (Ahmed et al. 2020).

We propose a different perspective to handle such MSILP problems without attempting to convexify the expected cost-to-go functions that are parametrized by a mixture of continuous and integer state variables. Instead, we construct a *partially extended reformulation* by transforming the original problem into one that has integer variables only in the first stage. By doing so, we can employ decomposition algorithms, such as the branch-and-cut (B&C) algorithm, to decompose the problem into a master problem that corresponds to a first-stage mixed integer linear program, and a subproblem defined for the remaining stages, which is a multi-stage stochastic linear program that can be handled by stage-wise decomposition algorithms such as the SDDP algorithm. However, this reformulation may lead to exponentially many first-stage integer variables, thus can be computationally challenging to solve.

To alleviate the main challenge of the partially extended reformulation, we present an aggregation framework that imposes certain structures to the reformulation by aggregating the integer state variables. This framework could potentially leverage the structure of the underlying stochastic process to impose a restriction on the problem at hand that leads to an informative decision policy. In particular, we consider the case where the underlying stochastic process is modeled as a Markov chain (MC) and present several aggregation schemes based on the MC state information that lead to a wide range of aggregated MSILP approximations.

We demonstrate that the aggregated MSILP models can be solved exactly by a B&C framework integrated with a variant of the SDDP algorithm. The framework decomposes the problem into a master problem and a set of subproblems, where the master problem is concerned with the first-stage decisions including all the integer state variables, and the subproblems deal with the remaining stages, which involve continuous state and local variables. While this framework guarantees an optimal solution for the aggregated MSILP, solving a multi-stage stochastic program (via SDDP) at each incumbent solution encountered in the branch-and-bound process can be computationally prohibitive. We therefore settle with a relaxation bound from this approach (by relaxing the SDDP termination criteria), and pursue a more tractable approach that further approximates the aggregated MSILP as a two-stage stochastic program using two-stage linear decision rules (2SLDRs) (Bodur and Luedtke 2022). We present 2SLDR alternatives from the literature

and propose a new variant that leverages the MC structure of the underlying stochastic process. The resulting problem can be solved via Benders decomposition or any of its enhanced versions.

We perform a case study for the proposed aggregation framework and solution methods on a hurricane disaster relief logistic planning problem (HDR), where the stochastic process is given by an MC that models the evolution of the hurricane. We show how to create several aggregated MSILPs and 2SLDR approximations for this application. Our numerical experiments compare these alternatives, find suitable transformations and 2SLDR variants to obtain high-quality solutions with provable bounds, and present managerial insights about the resulting policies.

The remainder of the paper is organized as follows. Section 2 presents our aggregation framework and a literature review on MSILP methodologies. Section 3 introduces our novel B&C framework integrated with the SDDP algorithm and describes the 2SLDR scheme with our proposed MC-based variant. Section 4 presents a case study of the proposed methodologies for HDR. Section 5 presents the numerical experiment results and we end with some concluding remarks in Section 6.

2. Aggregation Framework for Integer State Variables

This section describes the proposed aggregation framework to generate informative policies for a class of MSILP problems, with the main idea of aggregating integer state variables in the MSILP by leveraging the structure of the underlying stochastic process.

2.1. Problem Formulation for MSILP

We consider a class of MSILP models with continuous local variables and mixed-integer state variables. Although integer variables may also appear as local variables in a generic MSILP, the restriction here (local variables being continuous only) does not limit the model from being applicable to a broad range of problems that arise in real-world applications. For example, a typical multi-period stochastic optimization problem with local production and distribution constraints can be represented with this model, such as the one we study in Section 4, motivated by the application of disaster relief logistics planning. Other possible applications include hydro-power scheduling (Hjelmeland et al. 2018) and unit commitment (Zou et al. 2018) problems.

We assume that the underlying stochastic process has finite support and, as such, it can be represented by a scenario tree \mathcal{T} . This assumption is not restrictive since we can construct an approximate scenario tree by means of sampling if the stochastic process has continuous support, a common practice in the literature (Shapiro et al. 2021). Let T be the number of stages and let \mathcal{N} be the set of nodes associated with the scenario tree \mathcal{T} . The set of nodes in each stage $t \in [T] := \{1, \dots, T\}$ is given by \mathcal{N}_t . The root node is denoted by r and it is the only node in the first stage, i.e., $\mathcal{N}_1 = \{r\}$. Each node $n \in \mathcal{N}_t$ with $t > 1$ has a unique parent in \mathcal{N}_{t-1} , denoted by $a(n)$. Thus, there is a unique path from r to any node $n \in \mathcal{N}_t$ in stage $t > 1$, and we let $\mathcal{P}(n)$ represent the

set of nodes on this path (including r and n). For each non-leaf node n (i.e., $n \in \mathcal{N}_t$ for $t \in [T-1]$), $\mathcal{C}(n)$ is the set of children of n , that is, the nodes whose parent is n . The probability that node n occurs is p_n , and we have $p_n > 0$ for all $n \in \mathcal{N}$ and $\sum_{n \in \mathcal{N}_t} p_n = 1$ for all $t \in [T]$. The transition probability from node $n \in \mathcal{N}_t$ for $t \in [T-1]$ to $n' \in \mathcal{C}(n)$ can then be written as $\bar{p}_{nn'} := p_{n'}/p_n$.

Next, we present a generic problem formulation for the class of MSILP problems that we consider in (P), which is often referred to as the nested formulation. For each node $n \in \mathcal{N}$, let $x_n \in \mathbb{R}^k$ and $z_n \in \mathbb{Z}^\ell$ be the set of continuous and integer state variables, respectively, and $y_n \in \mathbb{R}^r$ be the continuous local variables. The objective coefficient vectors associated with x_n, y_n , and z_n are denoted by d_n, h_n , and c_n , respectively. The overall objective is to minimize the sum of the local cost and the expected future cost at the root node:

$$\begin{aligned} Q_r &= \min c_r^\top z_r + d_r^\top x_r + h_r^\top y_r + \sum_{n \in \mathcal{C}(r)} \bar{p}_{rn} Q_n(x_r, z_r), & (P) \\ \text{s.t. } & H_r z_r \geq g_r, J_r x_r \geq f_r, C_r x_r + D_r z_r + E_r y_r \geq b_r, & (1) \\ & x_r \in \mathbb{R}^k, y_r \in \mathbb{R}^r, z_r \in \mathbb{Z}^\ell, \end{aligned}$$

where the cost-to-go function $Q_n(\cdot, \cdot)$ associated with each non-root node $n \in \mathcal{N} \setminus \{r\}$ is:

$$\begin{aligned} Q_n(x_{a(n)}, z_{a(n)}) &= \min c_n^\top z_n + d_n^\top x_n + h_n^\top y_n + \sum_{n' \in \mathcal{C}(n)} \bar{p}_{nn'} Q_{n'}(x_n, z_n), \\ \text{s.t. } & H_n z_n \geq G_n z_{a(n)} + g_n, & (2a) \\ & J_n x_n \geq F_n x_{a(n)} + f_n, & (2b) \\ & C_n x_n + D_n z_n + E_n y_n \geq A_n x_{a(n)} + B_n z_{a(n)} + b_n, & (2c) \\ & x_n \in \mathbb{R}^k, y_n \in \mathbb{R}^r, z_n \in \mathbb{Z}^\ell. \end{aligned}$$

The cost-to-go function $Q_n(\cdot, \cdot)$ associated with a leaf node $n \in \mathcal{N}_T$ only involves the cost incurred in the terminal stage T since its set of children nodes is empty, i.e., $\mathcal{C}(n) = \emptyset$. Our proposed formulation considers, separately, one set of constraints (2a) for integer state variables, one set of constraints (2b) for continuous state variables, and one set of constraints (2c) that link state and local variables. Although alternative formulations are possible, this form accommodates the presentation of the proposed aggregation framework, which will be described in Section 3.

Problem (P) is notoriously challenging to solve for two reasons. The first challenge is the nature of multi-stage stochastic programs where the underlying stochastic process is modeled as a scenario tree of an exponentially large size. Thus, it is impractical to solve these large-scale problems directly and decomposition methods can be deemed necessary. The second challenge is the existence of integer decision variables, which makes the expected cost-to-go functions (i.e., the value functions that represent the expected future cost defined at each node of the scenario tree) to be nonconvex in general. The following section reviews existing methods in the literature to handle (P) and similar problems, focusing on how they attempt to overcome these challenges.

2.2. Literature Review

In this section, we discuss recent advances in the literature that address MSILPs like formulation (P). Nested Benders decomposition (Birge 1985) and its sampling variant under the assumption of stage-wise independence, SDDP (Pereira and Pinto 1991), are classical decomposition algorithms for multi-stage stochastic linear programs. Recently, Zou et al. (2019) develop the SDDiP algorithm for MSILPs with pure binary state variables, where the nonconvex expected cost-to-go functions are lower approximated by a piecewise linear convex envelope constructed by the so-called Lagrangian cuts. These cuts are guaranteed to be exact at points where the associated state variables are binary, leading to the exactness and finite convergence of the algorithm. The problem becomes more challenging when the set of state variables includes general mixed-integer decision variables, which is the case for formulation (P). One option is to perform a binarization procedure to reformulate general integer state variables and approximate continuous state variables via binary variables, respectively, and then apply the SDDiP algorithm (Zou et al. 2018). However, this approach may suffer from a large number of binary state variables as a result of the binarization.

Alternatively, when the expected cost-to-go functions are assumed to be Lipschitz continuous, Ahmed et al. (2020) propose nonlinear cuts to exactly approximate the nonconvex expected cost-to-go functions via the augmented Lagrangian method. More recently, Zhang and Sun (2022) propose and study a unified theoretical framework for general multistage nonconvex stochastic mixed integer nonlinear programming problems where the expected cost-to-go functions are not necessarily Lipschitz continuous. Specifically, they propose regularized expected cost-to-go functions that can be exactly approximated by the generalized conjugacy cuts, generalizing the results of Ahmed et al. (2020). However, despite these theoretical convergence properties, the nonlinear cut generation from augmented Lagrangian and the implementation of a variant of the nested Benders decomposition algorithm that incorporates these cuts can be computationally challenging.

Füllner and Rebennack (2022) extend the ideas of binarization procedure in Zou et al. (2019) and the regularization procedure in Zhang and Sun (2022) to develop the so-called nonconvex nested Benders decomposition algorithm for generic multi-stage mixed integer nonconvex nonlinear program. Despite its generality and successful implementation in *deterministic* multi-stage MINLP problems, it is unclear, to the best of our knowledge, whether or not applying this approach to MSILPs with a large-scale scenario tree is computationally feasible.

In sum, all the works mentioned above aim to directly address the nonconvex (and sometimes non-Lipschitz) expected cost-to-go functions defined at each node of the scenario tree by developing exact lower-bounding techniques. In contrast, we circumvent the challenge of approximating the nonconvex expected cost-to-go functions by relocating all integer state variables (defined in the node subproblems) to the first stage (defined in the root node), so that the resulting cost-to-go functions

defined at non-root nodes are convex and can be approximated (exactly) by a decomposition scheme (e.g., nested Benders or SDDP). To the best of our knowledge, our paper is one of the first that presents the performance of stage-wise decomposition based algorithms for MSILPs with both integer and continuous state variables. In addition, we showcase the performance of the proposed heuristic decision policies via decision rules, 2SLDRs, as well as the lower bounding technique constructed based on our exact algorithm. We next discuss the proposed reformulation in detail.

2.3. A Partially Extended Formulation for MSILP

We propose a reformulation for MSILP in the form of a partially extended formulation, in which all integer decisions can be considered as the first-stage variables. The advantage of this reformulation is two-fold. First, the resulting cost-to-go functions at each node $n \in \mathcal{N} \setminus \{r\}$ have only continuous variables (including local and continuous state variables), making it amenable to apply standard Benders-type cutting-plane approximation to these convex functions. Second, since integer variables only appear in the first-stage problem, we could apply, e.g., a B&C procedure by branching on these first-stage integer variables and approximating the expected cost-to-go functions using stage-wise decomposition algorithms such as the SDDP (see Section 3). Specifically, the proposed partially extended (nested) formulation is given by:

$$\begin{aligned} Q_r^R = \min & \sum_{n \in \mathcal{N}} p_n c_n^\top z_n + d_r^\top x_r + h_r^\top y_r + \sum_{n \in \mathcal{C}(r)} \bar{p}_{rn} Q_n^R(x_r, z), & (P^R) \\ \text{s.t.} & H_r z_r \geq g_r, J_r x_r \geq f_r, C_r x_r + D_r z_r + E_r y_r \geq b_r, \\ & H_n z_n \geq G_n z_{a(n)} + g_n, & \forall n \in \mathcal{N} \setminus \{r\}, \end{aligned}$$

where the cost-to-go function $Q_n^R(\cdot, \cdot)$ associated with $n \in \mathcal{N} \setminus \{r\}$ is defined in a nested fashion as:

$$Q_n^R(x_{a(n)}, z) = \min_{x_n \in \mathbb{R}^k, y_n \in \mathbb{R}^r} \left\{ d_n^\top x_n + h_n^\top y_n + \sum_{n' \in \mathcal{C}(n)} \bar{p}_{nn'} Q_{n'}^R(x_n, z) \mid (2b) - (2c) \right\}.$$

The key difference between this reformulation (P^R) and the original formulation (P) is that all the integer state variables $z = \{z_n\}_{n \in \mathcal{N}}$, along with their associated constraint sets (2a), are now “moved” to the problem at the root node r , i.e., they are treated as first-stage variables and constraints. As such, the cost-to-go function associated with each non-root node only includes the state and local continuous variables and their associated constraints. Since the continuous state variables remain to be defined in the individual node-based subproblems through the nested form, we call this reformulation a *partially extended formulation*. The validity of this reformulation is clear because we only relocate the (node-based) integer state variables in order to achieve the desired property of piece-wise linear convex expected cost-to-go functions, given its first-stage decisions, which accommodates the application of stage-wise decomposition methods such as SDDP.

2.4. An Aggregation Framework via MC-based Structural Policies

While the proposed partially extended formulation (P^R) has piece-wise linear convex expected cost-to-go functions, given its first-stage decisions, it may lead to exponentially many integer variables in the first stage. We propose an aggregation framework that imposes structures to the formulation by aggregating the integer state variables to alleviate this challenge while maintaining high-quality solutions. In that regard, the idea is to leverage the structure of the underlying stochastic process to impose a restriction to the problem and obtain an informative decision policy.

In this paper, we consider the case where the underlying stochastic process is modeled as a Markov chain (MC). We choose this specific structure for two reasons. First, many stochastic processes in real-world applications can be modeled or well-approximated using MCs (Löhndorf and Shapiro 2019), including disaster relief logistics planning problems that we focus on in our numerical experiments (see Section 4). Second, it facilitates the exposition of our aggregation framework and allows us to construct a wide range of structured policies. We discuss alternative structural policies beyond those based on MC models in Remark 2.

To start, let \mathcal{M} be the set of MC states where each $m \in \mathcal{M}$ is an s -dimensional vector (i.e., $m \in \mathbb{R}^s$). Let \mathcal{M}_t represent the set of MC states that are reachable in stage $t \in [T]$, starting from a given initial state. Notation $m^t = (m_1, \dots, m_t)$, where $m_{t'} \in \mathcal{M}_{t'}$ for $t' \in [t]$, represents a sequence of admissible MC states, that is, the probability to reach $m_{t'+1}$ given $m_{t'}$ is positive for all $t' \in [t-1]$.

We now describe the connection between the MC states and the scenario tree model used in the MSILP formulation (P). Each node $n \in \mathcal{N}_t$ ($t \geq 1$) of the scenario tree corresponds to an MC state, $m_t(n) \in \mathcal{M}_t$, and is uniquely determined by the trajectory of the stochastic process from the root node (initial MC state) to MC state $m_t(n)$, that is, the sequence of MC states $m^t(n) = (m_1(r), \dots, m_{t-1}(a(n)), m_t(n))$. To simplify the exposition, we write m^t instead of $m^t(n)$ and m_t instead of $m_t(n)$ when we refer to any node $n \in \mathcal{N}_t$.

As stated above, the MC representation of the stochastic process allows us to parameterize each node in the scenario tree by means of their corresponding stage and sequence of MC states, i.e., for each $n \in \mathcal{N}_t$, we have $n \equiv (t, m^t)$. We can then rewrite the integer state variables z_n associated with node n as $z_n \equiv z_{t, m^t}$. Based on this parameterization, our aggregation framework relies on a linear transformation that will compress the full MC history up to stage t , namely the vector m^t of size $s \cdot t$, into a vector $\Phi_t m^t$ of size $q_t \leq s \cdot t$, by means of a transformation matrix $\Phi_t \in \mathbb{Z}^{q_t \times s \cdot t}$. That is, it parametrizes the integer state variables of stage t by the limited information $\Phi_t m^t$:

$$z_n \equiv z_{t, m^t} \rightarrow z_{t, \Phi_t m^t}^A \quad \forall n \in \mathcal{N}_t, t \in [T]. \quad (3)$$

It is clear that this aggregation significantly reduces the number of integer state variables if $q_t \ll s \cdot t$.

| Name | Transformation matrix Φ_t | Aggr. var. | Condition for $z_n \equiv z_{n'}, \forall n, n' \in \mathcal{N}_t$ |
|------|----------------------------------------------------------------------------------------------------------------|-------------------------------|----------------------------------------------------------------------|
| HN | $\mathbf{0}^{1 \times s \cdot t}$ | $z_{t,0}^A$ | $t(n) = t(n')$ |
| MA | $[\mathbf{0}^{s \times s \cdot (t-1)} \mid \mathbb{I}^{s \times s}]$ | z_{t,m_t}^A | $m_t(n) = m_t(n')$ |
| MM | $[\mathbf{0}^{2s \times s \cdot (t-2)} \mid \mathbb{I}^{2s \times 2s}]$ | $z_{t,(m_{t-1},m_t)}^A$ | $m_t(n) = m_t(n') \ \& \ m_{t-1}(a(n)) = m_{t-1}(a(n'))$ |
| PM | $[\mathbf{0}^{(\bar{s}+s) \times (s \cdot t - \bar{s} - s)} \mid \mathbb{I}^{(\bar{s}+s) \times (\bar{s}+s)}]$ | $z_{t,(\bar{m}_{t-1},m_t)}^A$ | $m_t(n) = m_t(n') \ \& \ \bar{m}_{t-1}(a(n)) = \bar{m}_{t-1}(a(n'))$ |
| FH | $\mathbb{I}^{s \cdot t \times s \cdot t}$ | z_{t,m^t}^A | $m^t(n) = m^t(n')$ |

Table 1 MC-based transformations examples, the resulting aggregated variables (Aggr. var.), and the equivalence of the original variables. $\mathbf{0}$ and \mathbb{I} denote the zero and identity matrices, respectively.

Different transformations can impose different structural properties on the discrete decisions and, thus, lead to policies with different levels of aggregation. Table 1 summarizes a set of transformations, the resulting aggregated variables, and the conditions under which the original integer state variables associated with nodes n and n' in \mathcal{N}_t correspond to the same (aggregated) variables, i.e., z_n and $z_{n'}$ have the same value given by $z_{t,\Phi_t m^t}^A$. We explain each of these policies below.

The simplest transformation we consider is a *Here-and-Now* (HN) transformation, which aggregates integer state variables associated with nodes that are in the same stage: $z_n \equiv z_{t,m^t} \rightarrow z_{t,0}^A$ for all $n \in \mathcal{N}_t$. Thus, the corresponding integer state variables will be independent of the realization of the stochastic process, i.e., follow a static policy (see, e.g., [Basciftci et al. \(2019\)](#)).

An alternative is to leverage the structure of the MC to create a transformation that leads to more informative policies. For example, the *Markovian* (MA) transformation aggregates integer state variables associated with nodes that share the same stage and the same MC state: $z_n \equiv z_{t,m^t} \rightarrow z_{t,m_t}^A$ for all $n \in \mathcal{N}_t$. Another option is to retain information from previous stages at the price of a larger number of aggregated variables. One such example is what we refer to as the *Double Markovian* (MM) transformation, which goes one step beyond MA and aggregates integer state variables associated with nodes that share the same MC states in the current and previous stage. Following this pattern, the extreme case is to consider the entire history of the MC states starting from the initial MC state, referred to as the *Full History* (FH) transformation. Note that FH leads to the original partially extended formulation for MSILP (P^R), since no integer state variables are effectively aggregated.

Although all the aforementioned MC-based transformations are constructed based on the complete MC state information, we may also consider partial information from MC states. Consider a subset of attributes of size $\bar{s} < s$ associated with MC state m_t , and assume without loss of generality that these attributes of interest correspond to the last \bar{s} attributes in the MC state vector, i.e., we represent an MC state as $m_t = (\cdot, \bar{m}_t) \in \mathbb{R}^{s-\bar{s}} \times \mathbb{R}^{\bar{s}}$. Then, we can create a transformation based only on this partial vector \bar{m}_t associated with the MC state m_t . As an example, we introduce the *Partial Markovian* (PM) transformation that considers complete MC state information in the current stage but only partial information from the previous stage.




Figure 1 An illustrative example showing nodes with the same aggregated integer state variables for transformations HN, MA, and MM. Nodes with the same color and shape share the same set of integer state variables.

EXAMPLE 1. Consider an MC \mathcal{M} with two possible states (i.e., light \circ and dark \bullet) and positive transition probabilities for all transitions between states. The first drawing in Figure 1 depicts a scenario tree built for \mathcal{M} starting at a light state and representing all possible MC states for $T = 4$ stages, where the color of each node (light or dark) represents the corresponding MC state. The remaining drawings in Figure 1 illustrate the effect of aggregation on the nodes of the scenario tree for transformations HN, MA, and MM (we omit PM because in this example the MC state vector has only one attribute). The HN tree shows that all nodes in the same stage will share the same set of (aggregated) integer state variables z^A (i.e., they have the same color and shape). In contrast, MA aggregates integer state variables in nodes with the same stage and MC state. Thus, the MA tree shows nodes of two different colors in each stage, which correspond to two separate sets of aggregated integer state variables for each stage (e.g., the transformation for node $n \equiv (t, m^t) = (3, \circ\bullet\circ) \in \mathcal{N}_3$ is $\Phi_3 m^3(\circ\bullet\circ) = m_3(\circ) = \blacklozenge$, thus, $z_{3, \blacklozenge}^A$ represents its associated aggregated variables). Finally, nodes in each stage of the MM tree have up to four different colors, as a result of an aggregation based on the MC states in both the current stage and the previous stage (e.g., $z_{3, \blacklozenge}^A$ is associated to node $(3, \circ\bullet\bullet)$ because $\Phi_3 m^3(\circ\bullet\bullet) = (m_2(\bullet), m_3(\bullet)) = \blacklozenge$). Overall, HN contains four sets of z^A variables (i.e., one for each stage), MA seven, and MM eleven, while the FH variant contains 15 sets of z^A variables (i.e., one for each node of the scenario tree). \square

REMARK 1. These transformations are just a few examples of a wide range of aggregations that can be constructed within our framework. One key consideration in constructing these aggregations is that there is a clear trade-off between the quality of the underlying decision policy and the corresponding computational effort. Aggregations that include (preserve) more information would naturally lead to better decision policies but also make the resulting problem harder to solve due to the increased number of integer state variables in the first stage. Choosing the right transformation depends on the specific problem type and the available computational resources. We leave the problem of adaptively finding the best transformation as a future research direction.

REMARK 2. The main advantage of assuming that the stochastic process is defined by an MC is that most stage-dependent processes can be approximated by an MC using a discretization technique with high accuracy (Bally and Pagès 2003, Pagès et al. 2004). Also, the work by Löhndorf and Shapiro (2019) shows that such an approximation can lead to high-quality solutions compared to the ones obtained by approximating the stochastic process with autoregressive models in certain applications. We also note that our approach is applicable to stage-wise independent stochastic processes. In these cases, one can approximate the probability distribution of a random variable by partitioning its support into a set of clusters and creating aggregations based on these clusters. For example, considering an MSILP with stochastic demand parameters, we may define three levels of demand (low, medium, and high) and create aggregation schemes based on each demand level.

2.5. Aggregated MSILP

We now present the aggregated MSILP formulation after applying a transformation to the integer state variables z via matrices $\Phi_t, t \in [T]$. The new set of integer state variables is given by $z_t^A \in \mathbb{Z}^{\ell \cdot q_t}$ for each $t \in [T]$. We employ a mapping function $\phi_t: \mathcal{N}_t \rightarrow \{t\} \times [q_t]$ to relate a node $n \in \mathcal{N}_t$, which was originally used to index variables z , to the index corresponding to its aggregated variables, i.e., $\phi_t(n) = (t, \Phi_t m^t)$ for each node $n \in \mathcal{N}_t$. Thus, the aggregated MSILP can be written as follows:

$$Q_r^A = \min \sum_{n \in \mathcal{N}} p_n c_n^\top z_{\phi_t(n)}^A + d_r^\top x_r + h_r^\top y_r + \sum_{n \in \mathcal{C}(r)} \bar{p}_{rn} Q_n^A(x_r, z^A) \quad (P^A)$$

$$\text{s.t. } H_r z_r \geq g_r, J_r x_r \geq f_r, C_r x_r + D_r z_r + E_r y_r \geq b_r, \quad (4a)$$

$$H_n z_{\phi_t(n)}^A \geq G_n z_{\phi_{t-1}(a(n))}^A + g_n, \quad \forall n \in \mathcal{N} \setminus \{r\}, \quad (4b)$$

$$x_r \in \mathbb{R}^k, y_r \in \mathbb{R}^r, z^A \in \mathbb{Z}^{\ell \cdot \sum_{t \in [T]} q_t}, \quad (4c)$$

where the cost-to-go function $Q_n^A(\cdot, \cdot)$ associated with $n \in \mathcal{N} \setminus \{r\}$ is defined in a nested fashion as:

$$Q_n^A(x_{a(n)}, z^A) = \min d_n^\top x_n + h_n^\top y_n + \sum_{n' \in \mathcal{C}(n)} \bar{p}_{nn'} Q_{n'}^A(x_n, z^A)$$

$$\text{s.t. } J_n x_n \geq F_n x_{a(n)} + f_n,$$

$$C_n x_n + D_n z_{\phi_t(n)}^A + E_n y_n \geq A_n x_{a(n)} + B_n z_{\phi_{t-1}(a(n))}^A + b_n,$$

$$x_n \in \mathbb{R}^k, y_n \in \mathbb{R}^r.$$

We note that any feasible solution of (P^A) is also feasible to (P) with the same objective value, because the only difference between (P^A) and (P) is that (P^A) enforces some integer variables of (P) to take the same values. However, the reverse is not true in general, and hence an optimal solution of (P^A) might be sub-optimal for (P) .

Although the aggregated problem (P^A) potentially has much fewer integer variables than (P^R), depending on the transformation Φ_t , it may remain computationally challenging for an exact solution approach due to the existence of (aggregated) first-stage integer variables and the continuous state variables defined at each node of the scenario tree. We next introduce an exact solution approach based on a variant of the SDDP algorithm, and then present alternative approximation approaches to generate feasible decision policies and dual bounds that are computationally efficient.

3. Solution Methodology

We present two solution methods for the aggregated MSILP formulation (P^A). Both are decomposition methods implemented as B&C algorithms, taking advantage of the fact that all the integer variables in (P^A) are in the first stage. The first one is an exact method that employs a variant of the SDDP algorithm adapted to our setting as a subroutine. The second method employs 2SLDR to impose additional solution structures that lead to a two-stage approximation to the MSILP.

3.1. SDDP Integrated B&C Algorithm

We now present an exact algorithm for solving formulation (P^A) by integrating the SDDP algorithm within a B&C framework. The high-level idea is to decompose the problem into a master problem and a set of subproblems, where the master problem is concerned about the first-stage decisions including all the integer variables of (P^A), and the subproblems deal with the remaining stages, which involve continuous state and local variables.

The algorithm performs a branch-and-bound (B&B) procedure for the master problem and enters a cut generation procedure every time an integer relaxation solution is encountered (e.g., done through the so-called callback functions in modern IP solvers such as CPLEX). To generate a cut, following the basic idea of Benders decomposition, we need to solve a subproblem associated with each child node of r in the scenario tree, which deals with the remaining stages of (P^A) via a variant of the SDDP algorithm for a given candidate solution $(\hat{x}_r, \hat{z}^A, \hat{\theta})$ from the master problem. Note that we adapt the standard SDDP algorithm—which is applicable when we have stage-wise independence—to our aggregated MSILP, taking advantage of the MC structure in the stochastic process (representing stagewise-dependent uncertainty). In particular, we consider the so-called multi-cut version, where a variable θ_n is introduced for each $n \in \mathcal{C}(r)$ to represent the value of the outer approximation of the cost-to-go function at node n . Thus, the master problem is given by:

$$\bar{Q}_r^A = \min \sum_{n \in \mathcal{N}} p_n c_n^\top z_{\phi_t(n)}^A + d_r^\top x_r + h_r^\top y_r + \sum_{n \in \mathcal{C}(r)} \bar{p}_n \theta_n \quad (5a)$$

$$\text{s.t. (4a) – (4c),}$$

$$\theta_n \geq \alpha_n^\top x_r + \beta_n^\top z^A + \gamma_n, \quad \forall (\alpha_n, \beta_n, \gamma_n) \in \mathcal{B}_n, \quad n \in \mathcal{C}(r), \quad (5b)$$

where \mathcal{B}_n stores the coefficients associated with all the Benders cuts constructed to approximate $Q_n^A(x_r, z^A)$ during the solution procedure. We refer the reader to Appendix A for a detailed description of the overall B&C procedure, including a pseudo-code shown in Algorithm 2.

We next explain the implementation details of this B&C framework integrated with SDDP. We first describe the SDDP subproblem structure and then present the entire SDDP sub-routine.

3.1.1. SDDP Subproblems. A key feature of the SDDP algorithm is its capability to leverage the structure of the underlying stochastic process in defining the cost-to-go functions and the corresponding SDDP subproblems. For example, we only need to create one SDDP subproblem per stage in the case of stage-wise independence (Pereira and Pinto 1991), and one subproblem per stage and MC state in the case when the stage-wise dependence structure is modeled by an MC (Bonnans et al. 2012, Philpott and De Matos 2012). In what follows, we describe the SDDP subproblems for formulation (P^A) based on the MC structure and the employed transformation.

First, we observe that if the first-stage integer variables z^A do not exist, instead of defining one subproblem for each node $n \in \mathcal{N}$, we can have one subproblem defined for each stage $t \in [T]$ and MC state $m_t \in \mathcal{M}_t$, that is, all nodes $n \in \mathcal{N}_t$ with $m_t(n) = m_t$ share the same cost-to-go function (Philpott and De Matos 2012), thanks to the Markovian nature of the stochastic process. However, this might not be the case when we consider the aggregated variables z^A because cost-to-go functions with the same MC state might be associated with different aggregated variables depending on the chosen transformation. In such a case, the cost-to-go functions are defined over two different sets of variables. To see this, consider two nodes $n, n' \in \mathcal{N}_t$ with the same MC state (i.e., $m_t(n) = m_t(n') = m_t$) but with different sets of (aggregated) integer variables z^A (i.e., $\phi_t(n) \neq \phi_t(n')$), then we need to define two separate SDDP subproblems for n and n' . In other words, we can only define the same SDDP subproblem for two nodes $n, n' \in \mathcal{N}_t$ if they correspond to the same MC state (i.e., $m_t(n) = m_t(n')$) and the same aggregated variables (i.e., $\phi_t(n) = \phi_t(n')$). In the following example we use a policy graph—a concept introduced by Dowson (2020)—that represents the relationship between SDDP subproblems. Note that the policy graphs for (P^A) are constructed based on the MC transition graph and the chosen transformation.

EXAMPLE 2. Figure 2 illustrates different policy graphs for the HN, MA, and MM transformations based on the illustrative example shown in Figure 1. Recall that the underlying stochastic process is defined by an MC with two states (i.e., light \circ and dark \bullet). Since at a minimum we should have one SDDP subproblem defined for each MC state, we have at least two SDDP subproblems defined in each stage $t \geq 2$. In the case of MA, since nodes with the same MC state share the same set of aggregated integer state variables, we have exactly two SDDP subproblems per stage (i.e., one per MC state) and the policy graph only depends on the MC transitions (i.e., a Markovian




Figure 2 SDDP policy graphs and subproblems for various types of transformations, HN, MA, and MM.

policy graph). The same applies for HN, although we only have one set of aggregated integer state variables per stage. Since their corresponding policy graphs are identical, we use HN/MA to denote the first policy graph in Figure 2. On the other hand, the SDDP subproblem structure for MM is quite different: nodes with the same MC state can be associated with different aggregated integer state variables, because the MM transformation depends not only on the MC state in the current node but also its parent node. For example, nodes $(3, \bullet\bullet\bullet)$ and $(3, \circ\circ\bullet)$ are associated with the same stage and MC state ($m_3(\bullet\bullet\bullet) = m_3(\circ\circ\bullet) = \bullet$, see Figure 1), but since the MC states of their parent nodes differ, they have different sets of aggregated integer state variables (i.e., $z_{3,\blacklozenge}^A$ and $z_{3,\blacklozenge}^A$, respectively). Therefore, MM has one SDDP subproblem for each pair of MC states, corresponding to both the current node and the parent node for stages $t \geq 3$. The policy graph for MM can be seen as an aggregated version of its scenario tree (see Figure 1) with one copy for each type of nodes. \square

To formally define the SDDP subproblem for (P^A) , we represent the SDDP subproblem corresponding to node $n \in \mathcal{N}_t$ in the scenario tree by parameterizing it with mapping $\sigma(n) := (t, m_t(n), \phi_t(n))$. Consider \mathcal{S}_t to be the set of all SDDP subproblems defined for stage $t \geq 2$. The number of subproblems in stage $t=2$ is always the same (i.e., $|\mathcal{S}_2| = \mathcal{N}_2$) since each node in that stage is associated to a different MC state, whereas the number of subproblems in stages $t > 2$ depends on the transformation employed. For a given SDDP subproblem $\sigma \in \mathcal{S}_t$, we define $\mathcal{N}_t^S(\sigma)$ as the set of nodes associated with σ , i.e., all nodes $n \in \mathcal{N}_t$ such that $\sigma(n) = \sigma$.

To represent the relationship between SDDP subproblems illustrated by a policy graph, let $\mathcal{C}^S(\sigma)$ be the set of children subproblems for each subproblem σ in the policy graph. Specifically, σ' is a child of σ if $\sigma' \in \mathcal{S}_{t+1}$ and there exists a node $n' \in \mathcal{N}_{t+1}^S(\sigma')$ such that its parent is associated with σ (e.g., the children set of \blacksquare in the MM policy graph is $\mathcal{C}^S(\blacksquare) = \{\blacklozenge, \blacklozenge\}$, as shown Figure 2). The transition probability from $\sigma \in \mathcal{S}_t$ to one of its children $\sigma' \in \mathcal{C}^S(\sigma)$ is given by the MC transition probability, that is, $\bar{p}_{\sigma\sigma'} = p_{a(n')n'}$ for any node $n' \in \mathcal{N}_{t+1}^S(\sigma')$ with parent $a(n') \in \mathcal{N}_t^S(\sigma)$. Lastly, since there is at least one SDDP subproblem defined for every MC state, the random coefficients involved in the SDDP subproblems can be parameterized by the corresponding MC states, for example, the random cost vector in $\sigma \in \mathcal{S}_t$ is defined as $d_\sigma = d_n$ for any $n \in \mathcal{N}_t^S(\sigma)$.

To preserve the stage-wise dependency structure between SDDP subproblems in two consecutive stages, we introduce local copies of the first-stage integer state variables z^A in the subproblems.

In other words, we copy the first-stage integer solution from an SDDP subproblem to its children subproblems to ensure that each subproblem has the needed information given by z^A . We note that the use of local copies is common in other SDDP variants to maintain the stage-wise decomposition structure of the problem (see, e.g., [Zou et al. \(2019\)](#)). Then, the SDDP subproblem $\sigma := \sigma(n)$ for any node $n \in \mathcal{N}_t$ with $t > 1$ is given by:

$$\begin{aligned} \bar{Q}_\sigma^A(x_{a(\sigma)}, \zeta_{a(\sigma)}) = \min & d_\sigma^\top x_\sigma + h_\sigma^\top y_\sigma + \sum_{\sigma' \in \mathcal{C}^S(\sigma)} \bar{p}_{\sigma\sigma'} \theta_{\sigma'} \\ \text{s.t. } & J_\sigma x_\sigma \geq F_\sigma x_{a(\sigma)} + f_\sigma, \end{aligned} \quad (6a)$$

$$\zeta_\sigma = \zeta_{a(\sigma)}, \quad (6b)$$

$$C_\sigma x_\sigma + D_\sigma \zeta_{\sigma, \phi_t(\sigma)} + E_\sigma y_\sigma \geq A_\sigma x_{a(\sigma)} + B_\sigma \zeta_{a(\sigma), \phi_{t-1}(a(\sigma))} + b_\sigma, \quad (6c)$$

$$\theta_{\sigma'} \geq \alpha_{\sigma'}^\top x_\sigma + \beta_{\sigma'}^\top \zeta_\sigma + \gamma_{\sigma'}, \quad \forall (\alpha_{\sigma'}, \beta_{\sigma'}, \gamma_{\sigma'}) \in \mathcal{B}_{\sigma'}, \quad \sigma' \in \mathcal{C}^S(\sigma), \quad (6d)$$

where ζ_σ denotes the continuous local copy of the first-stage integer state variables z^A , which comes from its parent subproblem $\zeta_{a(\sigma)}$ via inequality (6b). We note that since \bar{Q}_r^A contains the original first-stage integer variables z^A , the SDDP subproblems for stage $t=2$ directly receive these variables instead of making local copies from the root (i.e., $\bar{Q}_\sigma^A(x_{a(\sigma)}, z^A)$ for each $\sigma \in \mathcal{C}^S(r)$, see [Figure 3](#)). A variable θ_σ is introduced for each SDDP subproblem σ to represent the outer approximation of the cost-to-go function, and a set \mathcal{B}_σ is introduced to store the coefficients associated with all the cutting planes constructed in this approximation during the solution procedure.

Finally, \bar{Q}_σ^A requires information from the parent of σ (i.e., $a(\sigma)$), however, an SDDP subproblem can have multiple parents in the policy graph, (e.g., both \blacklozenge and \blacktriangle are parents of \blacklozenge , see [Figure 2](#)). This ambiguity is not an issue for our SDDP implementation as we only consider one scenario at a time in the SDDP forward pass. So each time \bar{Q}_σ^A is solved, a specific node $n \in \mathcal{N}_t^S(\sigma)$ is considered as its parent node, and, thus, the parent of the SDDP subproblem would be $a(\sigma) = \sigma(a(n))$.

3.1.2. SDDP Algorithm. [Algorithm 1](#) describes the SDDP sub-routine implemented in our B&C algorithm. The procedure takes a candidate solution $(\hat{x}_r, \hat{z}^A, \hat{\theta})$ from the master problem (5) and one child node of the root $n' \in \mathcal{C}(r)$, and runs the SDDP algorithm until we find a valid inequality that cuts off $(\hat{x}_r, \hat{z}^A, \hat{\theta})$ or show that no violated cuts can be found. We use the multi-cut version of SDDP (i.e., one variable to approximate the objective value of each SDDP subproblem ([Philpott and De Matos 2012](#))) and the single-scenario iteration scheme (i.e., sample one scenario at a time from the scenario tree during the SDDP forward pass ([Philpott and Guan 2008](#))). Lastly, our algorithm considers the quick pass version of SDDP (i.e., we evaluate all the subproblems in a sample path before generating cuts in the backward pass ([Morton 1996](#))). We refer the reader to [Füllner and Rebennack \(2021\)](#) for further details on this and other SDDP variants.

Algorithm 1 SDDP procedure

```

1: procedure SDDP( $(\hat{x}_r, \hat{z}^A, \hat{\theta}), n', \epsilon, K, \mathbf{exact}$ )
2:   repeat
3:     Sample  $K$  paths from the sub-tree rooted at  $n'$  (i.e.,  $\mathcal{T}(n')$ ) and store indices in  $\mathcal{K}$ .
4:     cutAdded:= false.
5:     for  $k \in \mathcal{K}$  do
6:       for  $t \in \{2, \dots, T\}$  do                                     %% Forward pass sub-routines
7:         Select node of scenario  $k$  and stage  $t$ ,  $n := n_{t,k}$ , and its subproblem  $\sigma := \sigma(n)$ .
8:         Solve subproblem  $\sigma$ , let  $\hat{Q}_\sigma^A := \overline{Q}_\sigma^A(\hat{x}_{a(\sigma)}, \hat{\zeta}_{a(\sigma)})$ . Save primal and dual solutions.
9:       for  $t \in \{T, \dots, 2\}$  do                                     %% Backward pass sub-routines
10:        Select node of scenario  $k$  and stage  $t$ ,  $n := n_{t,k}$ , and its subproblem  $\sigma := \sigma(n)$ .
11:        if  $|\hat{\theta}_\sigma - \hat{Q}_\sigma^A| \geq \epsilon |\hat{Q}_\sigma^A|$  then
12:          Add inequality (6d) to all subproblems with variables  $\theta_\sigma$ , including  $\overline{Q}_{a(\sigma)}^A$ .
13:          cutAdded:= true.
14:          if  $t = 2$  then return  $(\alpha_{\sigma(n')}, \beta_{\sigma(n')}, \gamma_{\sigma(n')})$ .           %% Add cut to master problem
15:   until SDDPConverge( $K, \mathbf{cutAdded}, \mathbf{exact}$ ).                               %% Termination criteria

```

We now provide a detailed explanation of the main components of the algorithm SDDP. We first sample (without replacement) K scenarios (sample paths) from the sub-tree rooted at node $n' \in \mathcal{N}_2$ to generate cuts for its associated variable $\theta_{n'}$ in the master problem. For each scenario, we then perform the forward and backward pass of the SDDP algorithm (lines 5-14). The forward pass iterates over all stages from $t = 2$ to $t = T$ on a randomly sampled scenario (sample path), selects the associated node at each stage, and solves the corresponding SDDP subproblem (lines 6-8). The backwards pass iterates over all stages from $t = T$ to $t = 2$ (lines 9-14) and checks the relative difference between the approximated value $\hat{\theta}_\sigma$ and \hat{Q}_σ^A (line 11). If this difference is larger than a given tolerance ϵ , we generate a Benders cut as represented in (6d). In particular, the $(\alpha_{\sigma'}, \beta_{\sigma'}, \gamma_{\sigma'})$ coefficients for any $\sigma' \in \mathcal{C}^S(\sigma)$ are given by:

$$\alpha_{\sigma'} = \pi^{(2b)\top} F_{\sigma'} + \pi^{(6c)\top} A_{\sigma'}, \quad \beta_{\sigma'} = \pi^{(6c)\top} B_{\sigma'}, \quad \gamma_{\sigma'} = \hat{Q}_{\sigma'}^A - \left(\alpha_{\sigma'}^\top \hat{x}_\sigma + \beta_{\sigma'}^\top \hat{\zeta}_{\phi_{t-1}(\sigma)} \right).$$

where $\pi^{(2b)}$ and $\pi^{(6c)}$ are the dual solutions associated with inequalities (2b) and (6c) of subproblem σ' , respectively. This cut is added to the parent subproblem (i.e., $\overline{Q}_{\sigma(a(n))}^A$) but it is also valid for any subproblem in stage $t - 1$ involving variable $\theta_{\sigma(a(n))}$ (Philpott and De Matos 2012). We note that while the forward pass randomly samples scenarios rooted at one particular child node of the root node, the backward pass can add cuts to subproblems in subtrees rooted at other children nodes of the root node. Also, we consider feasibility cuts during the forward pass if one of the subproblems is infeasible. Lastly, we terminate the current SDDP iterations for scenarios rooted at n' if a cut is added to the first-stage (master) problem (i.e., $t = 2$), otherwise, we continue iterating over the remaining scenarios until the termination criterion is met.

Procedure `SDDPConverge` ensures that we terminate the SDDP algorithms when we can guarantee that no cut can be added to the master problem. To do so, we increase the number of sampled scenarios K to be equal to the total number of scenarios (i.e., $K = |\mathcal{N}_T|$) when no cut was added using the original sample size. Once $K = |\mathcal{N}_T|$, we continue iterating until no violated cut can be added to any subproblem (i.e., `cutAdded=false`). We note that this termination criterion can lead to several iterations over the full set of scenarios of the scenario tree, which is computationally expensive. However, this scenario increment is necessary to ensure an optimal solution to (P^A) .

EXAMPLE 3. Figure 3 illustrates the main components of our SDDP integrated B&C algorithm for the MA transformation. Note that in this case an SDDP subproblem is defined for each MC state in each stage $t > 1$. The two illustrations show the decomposition structure over the policy graph: the root node \circ corresponds to the master problem while other nodes correspond to the SDDP subproblems. The drawing on the left depicts the forward pass of Algorithm 1 for scenario $\circ \rightarrow \bullet \rightarrow \bullet \rightarrow \bullet$. Note that the solution of the SDDP master problem, $(\hat{z}^A, \hat{x}_\circ)$, is used as part of the input parameters to \blacksquare and, similarly, $(\hat{\zeta}_\blacksquare, \hat{x}_\blacksquare)$ is used as part of the input to \blacklozenge . Recall that the ζ variables are just the copies of z^A , so $\hat{z}^A = \hat{\zeta}_\blacksquare = \hat{\zeta}_\blacklozenge$. The drawing on the right illustrates the backward pass where the dual optimal solution of an SDDP subproblem is used to generate cuts to improve the outer-approximation of the cost-to-go function defined at its parent subproblems. For example, the optimal dual solution of \blacklozenge is used to generate cuts for subproblems \blacksquare and \blacksquare . \square




Figure 3 An illustrative example of the SDDP integrated B&C algorithm. It depicts decomposition structure and the two main SDDP subroutines (i.e., forward and backward pass).

3.1.3. Lower Bound Computation. Algorithm 2 guarantees an optimal solution to (P^A) , but the procedure can be computationally expensive because the number of SDDP subproblems depends on the underlying MC structure and the chosen transformation. In addition, we may need to call the SDDP sub-routine multiple times to correctly evaluate each candidate integer relaxation solution encountered during the B&C procedure. With this in mind, we propose an alternative to the aforementioned procedure to considerably reduce its computation time and return valuable information about the problem in the form of a lower bound. This lower bound can be used, for example, to provide an overestimate of the optimality gap associated with a known feasible

solution. We utilize this lower bound to evaluate the quality of solutions obtained by our proposed 2SLDR approximations (see Section 3.2 for details).

The overall idea for this lower-bounding technique is to apply a limited number of iterations of the SDDP sub-routine. First, we increase the cut violation tolerance parameter ϵ only to add cuts that have a significant relative difference to its true value (e.g., using $\epsilon = 0.1$ instead of a conservative value $\epsilon = 0.0001$). Also, we limit the number of rounds in the `SDDPConverge` procedure in Algorithm 1. Specifically, when Boolean variable `exact = false`, we limit the number of rounds to only three and avoid increasing the number of sampled scenarios K . This is in contrast with the `exact = true` version (see the previous subsection), where we iterate until no more cuts are added to any subproblem and the number of sampled scenarios increases to cover all possibilities. These modifications to the termination criteria considerably reduce the computational time of the SDDP procedure. However, we only obtain lower-approximations to the cost-to-go functions as we might be missing cuts in the SDDP master and subproblems, resulting in a lower bound for (P^A) .

3.2. Two-stage Linear Decision Rule Approximation

In this section, we describe an approximation scheme to create feasible solutions for (P^A) by imposing additional structures to the decision policies in the form of 2SLDR (Bodur and Luedtke 2022). Imposing additional structures in such a way transforms the multi-stage stochastic program into a two-stage stochastic program, which requires significantly less computational effort than the B&C framework integrated with the SDDP algorithm that we presented in the previous section. Moreover, the 2SLDR can be customized to leverage the structure of the underlying stochastic process, and we provide new schemes for the MC case.

Let ξ_t be the vector of random variables at stage $t \in [T]$, and $\xi_{t,n}$ be the realization of these random variables at node $n \in \mathcal{N}_t$, i.e., $\xi_{t,n} = \{A_n, B_n, C_n, D_n, E_n, F_n, J_n, b_n, c_n, d_n, f_n, h_n\}$. Note that ξ_r , the data associated with the root node of the scenario tree, is assumed to be a deterministic vector. In addition, $\xi^t = \{\xi_1, \dots, \xi_t\}$ and $\xi_n^t = \{\xi_{1,r}, \dots, \xi_{t-1,a(n)}, \xi_{t,n}\}$ represent the trajectory of the random variables up to stage t and their realizations for a particular node n in the scenario tree, respectively. We consider ξ_t and ξ^t as vectors of random variables with dimensions l_t and $l^t = \sum_{i=1}^t l_i$ for all $t \in [T]$, respectively, and the same applies to the realizations of these random variables.

The general idea behind 2SLDRs is to replace the original continuous state variables x_n with a new set of variables for each node $n \in \mathcal{N}$ by means of a linear relation of node-specific data. By doing so, we impose additional (linear relation) structure over the state variables, and in return transform the multi-stage stochastic program into a two-stage stochastic program (Bodur and Luedtke 2022). Specifically, 2SLDRs employ basis functions $\Psi_t(\cdot)$ to map ξ_n^t into values that are then used to create a linear representation for the state variables x_n , i.e., $x_n = \mu_n^\top \Psi_t(\xi_n^t) \in \mathbb{R}^k$,

where μ_t is a vector of the so-called LDR variables at stage t . Although the basis functions $\Psi_t(\cdot)$ can be any form of functions that map the node-specific data ξ_n^t to certain values, for ease of exposition, we consider simple basis functions that return the actual random variable realizations, that is, $\Psi_t(\xi_n^t) = \xi_n^t$. In this case, the LDR is given by:

$$x_n = \mu_t^\top \xi_n^t, \quad \forall n \in \mathcal{N}_t, t \in [T]. \quad (7)$$

We observe that these simple basis functions return high-quality decision policies for our case study (see Section 5), although more complex basis functions can potentially lead to better decision policies (Chen et al. 2008, Bampou and Kuhn 2011).

3.2.1. Aggregation Framework for 2SLDR. Similar to the aggregation framework introduced in Section 2, we create different aggregation schemes for the LDR variables μ , and as a result, obtain approximate 2SLDR-based decision policies. Table 2 summarizes a few LDR variable structures that we consider. The first two LDR variants (LDR-TH and LDR-T) are common in the literature (Shapiro and Nemirovski 2005, Kuhn et al. 2011), while LDR-M is a new LDR variant based on the MC structure of the stochastic process.

| Name | LDR variables for $t \in [T]$ | Standard basis function | Resulting LDR |
|--------|------------------------------------------------------------------------------|-------------------------------|-------------------------------------|
| LDR-TH | $\mu_t \in \mathbb{R}^{k \times l^t}$ | $\Psi_t(\xi_n^t) = \xi_n^t$ | $x_n = \mu_t^\top \xi_n^t$ |
| LDR-T | $\mu_t \in \mathbb{R}^{k \times l_t}$ | $\Psi_t(\xi_n^t) = \xi_{t,n}$ | $x_n = \mu_t^\top \xi_{t,n}$ |
| LDR-M | $\mu_{t,m} \in \mathbb{R}^{k \times l_t}, \quad \forall m \in \mathcal{M}_t$ | $\Psi_t(\xi_n^t) = \xi_{t,n}$ | $x_n = \mu_{t,m(n)}^\top \xi_{t,n}$ |

Table 2 Examples of LDR variable structures and basis functions for each $n \in \mathcal{N}_t$ and $t \in [T]$.

The first LDR variant, LDR-TH, which we refer to as a stage-history LDR, constructs the linear relation based on the trajectory of realizations of the random variables up to the current stage. The LDR-TH variant can lead to high-quality decision policies (thanks to the amount of information incorporated) but may result in a large number of LDR variables, especially for problems with a large number of stages. An alternative is LDR-T which constructs the linear relation only based on the realization of the random vector at the current stage, which leads to a smaller number of LDR variables at the expense of lower solution quality (Bodur and Luedtke 2022).

In addition to these LDR variants, we propose LDR-M which takes advantage of the MC structure of the underlying stochastic process. The idea is to have one copy of LDR variables $\mu_{t,m}$ for each MC state t and stage m . It is clear that LDR-M will lead to better policies compared to LDR-T because both LDR variants use the same information (realizations of random vectors), but the latter has a single copy of the LDR variables for all nodes in the same stage. We note that, however, LDR-M and LDR-TH are incomparable in general, as shown in our empirical results.

3.2.2. 2SLDR Models. We now present the resulting model after applying the LDR transformation (7) to (P^A). For ease of exposition, we only present the 2SLDR variant LDR-TH and the other two alternatives are similar. The overall idea is to replace the state variables x_n with the LDR variables using linear function (7). The resulting model for the root node is as follows:

$$Q_r^L = \min \sum_{n \in \mathcal{N}} p_n (c_n^\top z_{\phi_t(n)}^A + d_n^\top (\mu_{t(n)}^\top \xi_n^{t(n)})) + h_r^\top y_r + \sum_{n \in \mathcal{N} \setminus \{r\}} p_n Q_n^L(\mu, z^A) \quad (P^L)$$

$$\text{s.t. } H_r z_r \geq g_r, J_r \mu_1^\top \xi_r^1 \geq f_r, C_r \mu_1^\top \xi_r^1 + D_r z_r + E_r y_r \geq b_r, \quad (8a)$$

$$H_n z_{\phi_t(n)}^A \geq G_n z_{\phi_{t-1}(a(n))}^A + g_n, \quad \forall n \in \mathcal{N} \setminus \{r\}, \quad (8b)$$

$$J_n \mu_{t(n)}^\top \xi_n^{t(n)} \geq F_n \mu_{t(a(n))}^\top \xi_{a(n)}^{t(a(n))}, \quad \forall n \in \mathcal{N} \setminus \{r\}, \quad (8c)$$

$$(y_r, \mu, z^A) \in \mathbb{R}^r \times \mathbb{R}^{k \cdot \sum_{t \in [T]} l^t} \times \mathbb{Z}^{\ell \cdot \sum_{t \in [T]} q_t}.$$

Model (P^L) contains all the root node variables and constraints of (P^A) and includes all the LDR variables and its associated constraints (8c) from all nodes in the scenario tree. We note that there is no need to define LDR variables for the root node since the x_r variables are already first-stage variables, but we include them in the model for ease of exposition. Lastly, the cost-to-go function for node $n \in \mathcal{N} \setminus \{r\}$ only depends on the first-stage variables μ and z^A and is given by:

$$Q_n^L(\mu, z^A) = \min_{y_n \in \mathbb{R}^r} \left\{ h_n^\top y_n \mid C_n \mu_{t(n)}^\top \xi_n^{t(n)} + D_n z_{\phi_t(n)}^A + E_n y_n \geq A_n \mu_{t(n)-1}^\top \xi_{a(n)}^{t(n)-1} + B_n z_{\phi_{t-1}(a(n))}^A + b_n \right\}.$$

Note that (P^L) is indeed a two-stage stochastic program. Applying (7) we remove the dependency between parent and child nodes, thus, all cost-to-go functions become functions of only the first-stage decisions. This two-stage stochastic program is easier to solve than its multi-stage counterpart since it eliminates the nested dependency between the cost-to-go functions in different stages.

Formulation (P^L) considers one cost-to-go function for each node in the scenario tree. However, the number of cost-to-go functions can be considerably smaller depending on the LDR variant and transformation Φ_t . For example, LDR-TH requires one cost-to-go function for each node in the scenario tree because it considers the entire trajectory of realizations of random variables. Therefore, the cost-to-go functions for LDR-TH are unique for each node even if they share the same stage and MC state. On the contrary, LDR-T and LDR-M only employ local information, so the cost-to-go functions for two nodes $n, n' \in \mathcal{N}_t$ are the same if $\sigma(n) = \sigma(n')$, that is, if they share the same stage t , MC state $m(n) = m(n')$, and aggregated variables $\phi_t(n) = \phi_t(n')$. Therefore, we have the same number of distinct cost-to-go functions as SDDP subproblems, which can be significantly smaller than the size of the scenario tree. In conclusion, LDR-T and LDR-M have a significant computational advantage over LDR-TH due to the fewer LDRs variables and cost-to-go functions employed (along with all the variables and constraints necessary to define them).

We solve the resulting two-stage stochastic programming model using Benders decomposition, a standard approach in the stochastic programming literature (see, e.g., Zverovich et al. (2012)). Our specific implementation details can be found in Appendix B.

4. Case Study: A Hurricane Disaster Relief Planning Problem

In this section, we present a case study to exemplify the aggregated framework and evaluate the proposed solution methodologies. We consider a class of multi-period hurricane disaster relief logistics planning (HDR) problem, where a number of contingency modality options can be activated during the planning horizon to increase the capacities at the distribution centers (DCs).

The HDR problems are well-studied in the operations research literature because of their practical impacts on reducing economic loss and human suffering (Graumann et al. 2006, Seraphin 2019). The operational challenges in HDR problems arise from hurricanes' stochastic nature (Sabbaghtorkan et al. 2020). To address these challenges, most of the HDR literature focuses on two-stage stochastic programming models, where the first-stage model makes strategic facility location and resource pre-positioning decisions in terms of shelter preparation and resource allocation, while the second stage deals with disaster relief logistics decisions once the damage and demand information are revealed after the hurricane's landfall (see, e.g., Duran et al. (2011), Lodree Jr et al. (2012), Davis et al. (2013), Alem et al. (2016), and a survey paper by Sabbaghtorkan et al. (2020)).

Recent works have also considered the multi-stage HDR variants where sequential logistics decisions are made in a rolling horizon (RH) fashion (Pacheco and Batta 2016, Siddig and Song 2022, Yang et al. 2022). For example, Pacheco and Batta (2016) consider a relief supply preposition and reposition problem starting from the time when the hurricane is first detected, with updated forecast information on the hurricane's attributes every six hours. The authors proposed a forecast-driven dynamic model for the problem, and a solution approach combining scenario analysis and RH. Similarly, Siddig and Song (2022) consider a multi-period preposition problem where the demand is realized at landfall (i.e., last period). The authors proposed an MSLP formulation and showed the value of multi-stage stochastic programming solutions compared to solutions given by an RH approach and a two-stage approximation.

Despite their wide range of problem settings, the HDR problems have several assumptions in common. For example, the evolution of the hurricane and intensity are commonly modeled using an MC (Taskin and Lodree Jr 2010, Pacheco and Batta 2016) and the estimated demand at each shelter can be modeled using a deterministic mapping from the current hurricane state (i.e., hurricane location and intensity) (Siddig and Song 2022). Similarly, the HDR problem that we study follows these assumptions as well. One unique feature of our HDR problem is that we consider a set of contingency modality options that can be activated during the planning horizon to increase the inventory capacities of the DCs. We introduce more details about the problem next.

4.1. Problem Description and Model Formulation

Our multi-stage HDR problem is concerned with the logistics decisions of producing and distributing relief commodities to shelters prior to the landfall of a hurricane. The relief commodities are

used to fulfill the demand for civilians that are evacuated to shelters. As in the related problems, we consider that demand can be estimated according to the location and intensity of the hurricane, updated every six hours, which is the frequency of the hurricane forecast update by the National Hurricane Center (Regnier and MacKenzie 2019). We assume that shelters, DCs, and possible positions of the hurricane are located in a two-dimensional grid with a discrete set of x - and y -coordinates. The goal is to produce and transport relief commodities from the DCs to the shelters at a minimum (expected) cost, which consists of the penalty cost for unsatisfied demand and operational costs incurred in all time periods.

In addition, we consider a set of contingency modalities that can be activated to increase the DCs' capacities. These modalities make the disaster relief logistics system more resilient to extreme hurricane events that may lead to unexpectedly high demand at the price of higher logistics costs. We make a practically relevant assumption that these modalities can be activated at most once during the planning horizon, and once activated, the modality will stay the same until the end of the planning horizon. An active modality represents an incremental capacity increase at every stage on the planning horizon. These considerations are consistent with the practical constraints associated with these large-scale critical logistics operations (e.g., the phased activation of national guards for disaster relief efforts), which have this "all or nothing" feature and have to be implemented in small increments. Lastly, we consider a one-leg delay for these modalities (i.e., the capacity starts to increase from the stage immediately after the stage when a modality activation decision is made), which partially captures the logistical challenges of modality changes during disaster relief.

We assume that the hurricane's evolution can be characterized by an MC \mathcal{M} . Each MC state corresponds to two attributes of the hurricane: location and intensity, that is, $m = (m^x, m^y, m^i) \in \mathcal{M}$ where m^x and m^y are the x - and y -coordinates of the hurricane's location, respectively, and m^i represents the hurricane's intensity level. We consider independent probability transition matrices for each hurricane attribute. For the hurricane intensity, We use the transition probability matrix described in Pacheco and Batta (2016) in our test instances. For the hurricane movement, we assume that the hurricane originates from the bottom row of the grid (i.e., $m^y = 0$ for the initial MC state) and that m^y increases by one in each period (i.e., the hurricane advances upwards by one step in each period) and will reach land in exactly T periods (see Appendix D for additional details). We note that our approach can be easily extended to the case when the number of stages until landfall is random, as long as the underlying stochastic process is modeled by an MC.

Given the stochastic behavior of the hurricane, we consider an MSILP model with a stage defined for each time period when new information is realized (i.e., the hurricane moves to a new MC state). We construct a scenario tree \mathcal{T} by considering an initial MC state and all possible MC transitions for up to T stages (i.e., until landfall). We consider a set of contingency modalities \mathcal{L} , a

set of shelters \mathcal{I} , and a set of DCs \mathcal{J} , where each DC $j \in \mathcal{J}$ has an initial capacity C_j and inventory I_j . At each node $n \in \mathcal{N}$ of the scenario tree \mathcal{T} , one needs to decide the amount of relief commodities $v_j \geq 0$ to produce for each DC $j \in \mathcal{J}$, the amount of relief commodities $y_{ij} \geq 0$ to transport from each DC $j \in \mathcal{J}$ to each shelter $i \in \mathcal{I}$, and the amount of unsatisfied demand $w_i \geq 0$ at each shelter $i \in \mathcal{I}$. We introduce two continuous state variables $x_n = (x_n^I, x_n^C)$ to represent the inventory and capacity for each DC $j \in \mathcal{J}$ at node $n \in \mathcal{N}$, respectively. Lastly, we introduce a binary state variable $z_{n\ell} \in \{0, 1\}$ to represent whether or not a contingency modality $\ell \in \mathcal{L}$ is active at node $n \in \mathcal{N}$.

Each node $n \in \mathcal{N}$ is associated with realizations of random variables according to the corresponding MC state of the hurricane. Specifically, for a given node $n \in \mathcal{N}_t$, $t \in [T]$, the demand of a shelter $i \in \mathcal{I}$, d_{ni} , the production cost at DC $j \in \mathcal{J}$, q_{nj} , and the transportation cost from DC j to shelter i , f_{nij} , depend solely on the MC state $m_t(n)$ associated to node n . Thus, two nodes $n, n' \in \mathcal{N}_t$ with the same MC state at stage t , $m_t(n) = m_t(n')$, share the same realizations of random variables. The remaining model parameters are assumed to be deterministic, which include the unit inventory cost g_j at $j \in \mathcal{J}$, the contingency modality cost c_ℓ for $\ell \in \mathcal{L}$, the penalty cost for each unit of unsatisfied demand b_i for $i \in \mathcal{I}$, and the capacity increase $K_{j\ell}$ at DC j in each stage when modality ℓ is active. Further details on the problem instance generation can be found in Appendix D. The MSILP model for the HDR problem is given by:

$$\begin{aligned}
\min \quad & \sum_{j \in \mathcal{J}} \left(g_j x_{rj}^I + q_{rj} v_j + \sum_{i \in \mathcal{I}} f_{rij} y_{ij} \right) + \sum_{i \in \mathcal{I}} b_i w_i + \sum_{\ell \in \mathcal{L}} z_{r\ell} c_\ell + \sum_{n \in \mathcal{C}(r)} \bar{p}_{rn} Q_n(x_r, z_r) & (HDR) \\
\text{s.t.} \quad & \sum_{j \in \mathcal{J}} y_{ij} + w_i \geq d_{ri}, & \forall i \in \mathcal{I}, \quad (9a) \\
& x_{rj}^I = I_j - \sum_{i \in \mathcal{I}} y_{ij} + v_j, & \forall j \in \mathcal{J}, \quad (9b) \\
& v_j \leq x_{rj}^C, & \forall j \in \mathcal{J}, \quad (9c) \\
& x_{rj}^C = C_j, & \forall j \in \mathcal{J}, \quad (9d) \\
& \sum_{\ell \in \mathcal{L}} z_{r\ell} \leq 1, & (9e) \\
& x_{rj}^I, x_{rj}^C, v_j, w_i, y_{ij} \geq 0, z_{r\ell} \in \{0, 1\}, & \forall j \in \mathcal{J}, i \in \mathcal{I}, \ell \in \mathcal{L}.
\end{aligned}$$

Constraints (9a) and (9b) represent the demand and inventory constraints, respectively. Constraints (9c) enforce the production capacity and constraints (9d) represent the initial capacity of each DC. Constraint (9e) enforces that at most one modality can be activated at the root node of the scenario tree. Lastly, the cost-to-go function for a node $n \in \mathcal{N} \setminus \{r\}$ is given by:

$$Q_n(x_{a(n)}, z_{a(n)}) = \min \sum_{j \in \mathcal{J}} \left(g_j x_{nj}^I + q_{nj} v_j + \sum_{i \in \mathcal{I}} f_{nij} y_{ij} \right) + \sum_{i \in \mathcal{I}} b_i w_i + \sum_{\ell \in \mathcal{L}} z_{n\ell} c_\ell + \sum_{n' \in \mathcal{C}(n)} \bar{p}_{nn'} Q_{n'}(x_n, z_n)$$

$$\begin{aligned}
\text{s.t. } \quad & \sum_{j \in \mathcal{J}} y_{ij} + w_i \geq d_{ni}, & \forall i \in \mathcal{I}, \quad (10a) \\
& x_{nj}^I = x_{a(n)j}^I - \sum_{i \in \mathcal{I}} y_{ij} + v_j, & \forall j \in \mathcal{J}, \quad (10b) \\
& v_j \leq x_{nj}^C, & \forall j \in \mathcal{J}, \quad (10c) \\
& x_{nj}^C = x_{a(n)j}^C + \sum_{\ell \in \mathcal{L}} K_{j\ell} z_{a(n)\ell}, & \forall j \in \mathcal{J}, \quad (10d) \\
& \sum_{\ell \in \mathcal{L}} z_{n\ell} \leq 1, & (10e) \\
& z_{a(n)\ell} \leq z_{n\ell}, & \forall \ell \in \mathcal{L}, \quad (10f) \\
& x_{nj}^I, x_{nj}^C, v_j, w_i, y_{ij} \geq 0, z_{n\ell} \in \{0, 1\}, & \forall j \in \mathcal{J}, i \in \mathcal{I}, \ell \in \mathcal{L}.
\end{aligned}$$

The meanings of constraints (10a)-(10c) and (10e) are identical to their counterparts in (*HDR*). Constraints (10d) represent the capacity increase at a DC depending on whether or not a contingency modality is active in the previous stage. Constraints (10f) enforce that an active contingency modality in the previous stage will remain active in the current stage.

4.2. Aggregation Framework

We consider the four MC-based transformations introduced in Table 1 of Section 2, i.e., *HN*, *MA*, *MM*, and *PM*, for the aggregated model. The mathematical formulation of the aggregated model (*HDR^A*) can be found in Appendix C, which is a straightforward application of our proposed framework.

In the context of our *HDR* problem, aggregation is done for the contingency modality decisions. *HN* considers the same decisions for all nodes in a given stage; *MA* aggregates variables based on the hurricane state information in a given stage, which corresponds to the location and intensity of the hurricane; and *MM* incorporates the hurricane states at the current node and its parent node. The transformation *PM* can be customized depending on how one defines the partial MC state information. In our experiments, we consider the full hurricane state information (i.e., both location and intensity) at the current node and only the hurricane intensity information at its parent node, that is, the difference between *MM* and *PM* for the *HDR* problem is that the former incorporates the hurricane location information at the parent node into the aggregation, while the latter does not.

4.3. Two-stage LDR approximation

As explained in Section 3.2, we can apply a further restriction to the aggregated MSILP model (*HDR^A*) to obtain a two-stage model by applying LDR to the continuous state variables (i.e., inventory variables x_n^I). We present three LDRs variants similar to the ones introduced in Table 2.

First, we introduce notation $\xi_{t,n} = (d_n, q_n, f_n)$ to represent the vector of realizations of random variables for each node $n \in \mathcal{N}_t$ in the scenario tree, including the demand d_n , production cost q_n , and transportation cost f_n . Since the demand is one of the most important parameters for the

HDR problem, in our experiment we choose to use basis functions (for constructing LDRs) that are only based on either the demand realization at the current node $\Psi_t(\xi_n^t) = d_n$ or the entire history of demand realizations up to the current node $\Psi_t(\xi_n^t) = (d_r, \dots, d_{a(n)}, d_n)$. Then, the LDR variants for the inventory variables for a given node $n \in \mathcal{N}_t$ and DC $j \in \mathcal{J}$ are defined as:

$$\text{LDR-TH: } x_{nj}^I = \sum_{n' \in \mathcal{P}(n)} \sum_{i \in \mathcal{I}} \mu_{t(n')tji} d_{n'i} \quad \text{LDR-T: } x_{nj}^I = \sum_{i \in \mathcal{I}} \mu_{t(n)ji} d_{ni} \quad \text{LDR-M: } x_{nj}^I = \sum_{i \in \mathcal{I}} \mu_{tm_t(n)ji} d_{ni}$$

where $\mu_{t't} \in \mathbb{R}^{|\mathcal{J}| \times |\mathcal{I}|}$ for $t' \in [t]$, $\mu_t \in \mathbb{R}^{|\mathcal{J}| \times |\mathcal{I}|}$, and $\mu_{tm} \in \mathbb{R}^{|\mathcal{J}| \times |\mathcal{I}|}$ for $m \in \mathcal{M}_t$ are the LDR variables associated to each variant, respectively. LDR-TH considers the full demand history, while LDR-T and LDR-M only consider the demand realization at the current node. LDR-M has one set of variables for each hurricane state, while LDR-T and LDR-TH employ one set of variables per stage. We discuss the implications of these differences in Section 3.2, also in our numerical results in Section 5.2.

These LDR variants lead to three different two-stage approximations of (HDR^A), which are solved via Benders decomposition as discussed in Section 3.2 and Appendix B. Note that while the MSILP problem has the relatively complete recourse property, these two-stage approximations do not necessarily have this property. For instance, LDR variables μ can take negative values, which may lead to an infeasible solution with negative inventory values. We could force $\mu \geq 0$ to void this issue, but this further limits the approximation and can lead to a worse solution, as observed during preliminary experimentation. Thus, we incorporate Benders feasibility cuts in this regard.

5. Numerical Experiment Results

We now present the numerical results for the proposed aggregation framework and solution approaches. We test the models and approaches over various instances of the described HDR problem. In what follows, we first provide details on the test instances and the experimental setup. We then evaluate the computational performances of the proposed methodologies, focusing on the approximation methods (i.e., 2SLDR solutions with SDDP-based lower bounds) which perform best in practice. Lastly, we analyze the obtained decision policies and provide managerial insights.

5.1. Experimental Setup

We generate instances for our MSILP HDR problem for two grid sizes, 4×5 and 5×6 , which correspond to small-size and large-size instances, respectively (see Appendix E for more details on the sizes of these instances). We experiment with different levels of initial capacities for the DCs (i.e., C_j for $j \in \mathcal{J}$) to analyze the impact of decision policies derived from the aggregation framework to systems with different capacities to address demand fluctuations. Specifically, we consider instances where the initial capacities across all DCs are 20%, 25%, or 30% of the maximum demand across all shelters (the maximum demand is a predefined parameter for instance generation, see Appendix

D for details). We also consider two types of contingency modality options: conservative (Type-1) and aggressive (Type-2). For Type-1 modality, DC capacity expansion options are 10%, 20%, 30%, or 40% per stage, while for Type-2 modality, the options are 15%, 30%, 45%, or 60% per stage. We randomly generate a number of instances for each instance configuration and eliminate the ones where the optimal decision policies under transformations HN and FH are identical to make a meaningful comparison between different transformations. Overall, we obtain ten random instances for each of the six different instance configurations on two grid sizes.

The code for the decomposition algorithms was implemented in C++ using solver IBM ILOG CPLEX 20.1 with callback functions. All experiments were run using a single core/thread and a six-hour time limit. The experiments were run in the University of Toronto SciNet server Niagara, using cores with 16GB RAM each¹. We will make our code and instances available upon publication.

5.2. Computational Performance of the Proposed Methodologies

We now present the empirical performance of the proposed methodologies to solve the aggregated MSLP models. We focus on the best performing procedures which could be more valuable for practitioners, that is, the 2SLDR variants described in Section 4 (i.e., LDR-TH, LDR-T, and LDR-M) and the SDDP-based methods to obtain a lower bound (S-LB) and an upper bound (S-UB). We briefly discuss the results of the exact methods at the end of this section, that is, the B&C SDDP approach (S) and the extensive model (Ex), and refer to Appendix F for further details. For brevity, we only show results for the PM transformation since it achieves the best trade-off between policy quality and computational effort. See Appendix G for the results of the other transformations.

We first present results over small-size instances in Table 3. Columns “Average Time (sec)” report each approach’s time to obtain their respective optimal solution (all techniques obtained optimal solutions within the time limit), averaged over instances with the same configuration. Columns “Relative Difference (%)” compare the solution found by each approach and the (true) optimal solution obtained from Ex, that is, $(|Obj_{\text{Ex}} - Obj_i|)/Obj_{\text{Ex}}$ for all $i \in \{\text{LDR-TH, LDR-T, LDR-M, S-UB, S-LB}\}$. We observe that LDR-T and LDR-M take the least time on average. Moreover, they yield the smallest relative difference ($< 0.3\%$ on average in all configurations), indicating the high quality of their corresponding solutions. In particular, the newly proposed LDR variant, LDR-M, provides the best approximation and takes the least time in most configurations. This shows the value of developing decision policies that leverage the structure of the underlying stochastic process.

Table 3 also shows the performance of the SDDP lower (S-LB) and upper (S-UB) bound techniques. Both provide high-quality bounds with small relative differences (i.e., less than 1.5%). Also,

¹ See https://docs.scinet.utoronto.ca/index.php/Niagara_Quickstart for further server specifications.

Table 3 Solution time and quality of 2SLDR and SDDP bounds. Results for PM over small-size instances.

| | | Average Time (sec) | | | | | Relative Difference (%) | | | | |
|----------|------|--------------------|--------------|--------------|-------|--------|-------------------------|-------|-------------|-------------|------|
| Modality | Cap. | LDR-TH | LDR-T | LDR-M | S-LB | Ex | LDR-TH | LDR-T | LDR-M | S-UB | S-LB |
| Type-1 | 20% | 813.0 | 164.4 | 157.9 | 549.1 | 961.3 | 0.12 | 0.26 | 0.11 | 0.01 | 0.36 |
| | 25% | 380.1 | 84.3 | 77.7 | 347.0 | 487.2 | 0.08 | 0.11 | 0.03 | 0.26 | 0.71 |
| | 30% | 408.2 | 79.0 | 85.9 | 400.7 | 222.0 | 0.12 | 0.25 | 0.00 | 0.55 | 1.24 |
| Type-2 | 20% | 1406.0 | 247.2 | 265.2 | 940.5 | 2047.3 | 0.13 | 0.26 | 0.13 | 0.00 | 0.28 |
| | 25% | 670.5 | 108.2 | 98.1 | 534.4 | 651.0 | 0.09 | 0.13 | 0.04 | 0.23 | 0.49 |
| | 30% | 532.6 | 90.0 | 89.5 | 434.8 | 270.6 | 0.12 | 0.26 | 0.01 | 0.98 | 1.29 |
| Average | | 701.7 | 128.9 | 129.1 | 534.4 | 773.2 | 0.11 | 0.21 | 0.05 | 0.34 | 0.73 |

Table 4 Solution time and quality of 2SLDR and SDDP bounds. Results for PM over large-size instances.

| | | Average Time (sec) | | | | Opt. Gap (%) | | |
|-------------|------|--------------------|--------------|--------|-------|--------------|-------------|-------------|
| Modality | Cap. | LDR-T | LDR-M | S-LB | (opt) | LDR-T | LDR-M | S-UB |
| Type-1 | 20% | 5,353 | 3,732 | 13,564 | (9) | 0.62 | 0.39 | 0.24 |
| | 25% | 4,235 | 4,716 | 13,557 | (6) | 2.56 | 2.06 | 1.91 |
| | 30% | 1,555 | 1,675 | 2,067 | (10) | 0.80 | 0.63 | 0.27 |
| Type-2 | 20% | 7,489 | 4,855 | 18,323 | (5) | 9.00 | 7.33 | 9.01 |
| | 25% | 6,664 | 5,560 | 14,790 | (4) | 4.98 | 4.38 | 4.28 |
| | 30% | 1,717 | 1,640 | 6,971 | (10) | 0.81 | 0.63 | 0.30 |
| Av. (Total) | | 4,502 | 3,696 | 11,545 | (44) | 3.13 | 2.57 | 2.67 |

S-LB is computationally more efficient compared to Ex in most cases. Nonetheless, S-UB (evaluating the incumbent solution from S-LB) can require significant computational effort due to the additional SDDP calls: around 1-2 hours for small-size instances and up to 12 hours for large ones.

Table 4 presents results obtained by these approximation methods over large-size instances. The LDR variant LDR-TH is not included here as the corresponding formulation cannot be loaded into CPLEX due to the large number of first-stage LDR variables and constraints. Also, due to the lack of exact optimal solutions for large-size instances, the optimality gaps are conservative estimations computed using the best lower bound found by S-LB, that is, we replace Obj_{Ex} with $Obj_{\text{S-LB}}$.

As in the small-size instances, the 2SLDR variants find high-quality solutions with relatively low computational effort. In fact, LDR-T and LDR-M find their respective optimal solutions within the time limit for all the instances, while S-LB proves optimality in 44 of the 60 tested instances (as seen in the “(opt)” column). Nonetheless, S-LB finds high-quality lower bounds and S-UB yields marginally higher-quality solutions in most cases. The results suggest that 2SLDR, specifically LDR-M, is the best for finding high-quality solutions with limited computational effort, while the SDDP-based methods can yield high-quality lower and upper bounds given sufficient time.

REMARK 3. Appendix F analyzes the results of using exact methodologies (i.e., Ex and S) to solve the aggregated models. To summarize, Ex is faster and solves more instances than S in small-size instances. However, Ex cannot be solved for large-size problems due to memory requirements (the models use $\sim 16\text{GB}$ RAM after the pre-solve phase of CPLEX), while S can solve several instances to optimality and find feasible solutions with small optimality gaps. The performance

of \mathbf{S} is mostly explained by the large branching factor of the MC process and, consequently, the large number of SDDP subproblems (up to 593 for large-size instances). We believe that \mathbf{S} may be competitive in applications where the MC has fewer reachable states per stage (e.g., ten or fewer).

5.3. Decision Policies and Managerial Insights

We now analyze the quality and structure of the decision policies (in activating contingency modalities) associated with transformations HN, MA, PM and MM, and the original non-aggregated model (FH). The analysis considers small-size instances (i.e., 4×5 grids) since its purpose is to provide insights from the provably optimal solutions, and the non-aggregated model FH can only be solved by Ex, which cannot handle the large-size instances due to computational limitations.

Table 5 summarizes the performances of the considered models in terms of the resulting objective values. Columns “% Gap closed” present the percentage of the gap between HN and FH that is closed by each transformation averaged over all the instances, where the gap closed for each instance is calculated by $(Obj_{HN} - Obj_i)/(Obj_{HN} - Obj_{FH}), \forall i \in \{MA, PM, MM\}$. Thus, 0% corresponds to the same objective value given by HN and 100% corresponds to the objective value given by FH.

Table 5 Policy quality comparison for all transformations in our aggregation framework.

| | | Average Objective Value | | | | | % Gap closed | | |
|----------|------|-------------------------|---------|--------|--------|--------|--------------|------|------|
| Modality | Cap. | HN | MA | PM | MM | FH | MA | PM | MM |
| Type-1 | 20% | 104,162 | 102,572 | 92,924 | 92,924 | 82,193 | 7.5 | 51.9 | 51.9 |
| | 25% | 73,879 | 73,442 | 66,650 | 66,238 | 62,972 | 3.2 | 69.4 | 73.1 |
| | 30% | 48,970 | 48,951 | 47,425 | 47,314 | 47,117 | 0.3 | 89.5 | 96.7 |
| Type-2 | 20% | 104,135 | 102,625 | 92,612 | 92,612 | 81,182 | 6.7 | 50.9 | 50.9 |
| | 25% | 73,919 | 73,487 | 66,573 | 65,958 | 63,654 | 3.3 | 75.6 | 81.2 |
| | 30% | 48,970 | 48,968 | 47,540 | 47,349 | 47,253 | 0.0 | 80.3 | 97.0 |

We first observe that the difference between the objective values given by HN and FH is smaller when the initial DC capacities are higher—about 4% when the initial capacity is 30%, and about 28% when the initial capacity is 20%. This suggests that even the most restrictive transformation HN can lead to high-quality policies when the DCs are capable of addressing random fluctuating demand without much adaptability. Conversely, less restrictive transformations, which yield more adaptive contingency modality activation plans, become valuable when the initial capacity is small.

Table 5 also shows that the solution quality of MA and HN are similar (with $< 10\%$ difference) in all instances, indicating that information in the current stage is insufficient to obtain high-quality policies for our test instances. This effect is mitigated by in PM and MM since both can lead to significant performance improvements—closing more than 50% of the gap between FH and HN in all cases. In particular, the objective values of PM and MM are very close to those obtained by FH when the initial capacity is 30%, but this is not the case when the initial capacity is 20%. PM and

Table 6 Solution structures for different policies and initial capacities for Type-1 instances.

| Cap. | Nodes (%) | | | | | # of Contingencies | | | | | Aggressiveness (%) | | | | | Intensity | | | | |
|------|-----------|----|----|----|----|--------------------|-----|-----|-----|-----|--------------------|----|----|----|----|-----------|-----|-----|-----|-----|
| | HN | MA | PM | MM | FH | HN | MA | PM | MM | FH | HN | MA | PM | MM | FH | HN | MA | PM | MM | FH |
| 20% | 100 | 93 | 72 | 72 | 53 | 1.0 | 1.0 | 1.0 | 1.0 | 3.5 | 16 | 38 | 30 | 30 | 35 | 2.3 | 2.5 | 2.9 | 2.9 | 2.9 |
| 25% | 10 | 50 | 22 | 22 | 47 | 1.0 | 1.0 | 1.5 | 2.1 | 4.4 | 10 | 15 | 35 | 36 | 21 | 2.3 | 2.9 | 4.0 | 4.0 | 3.0 |
| 30% | 0 | 7 | 16 | 15 | 19 | - | 1.0 | 2.0 | 4.3 | 4.6 | - | 9 | 17 | 18 | 15 | - | 3.0 | 4.1 | 4.1 | 4.0 |

MM perform similarly, suggesting that the hurricane intensity state captures most of the valuable information from the previous stage. In Section 5.4 we further explore the reasons behind these.

We now take a closer look at the underlying contingency modality activation policies associated with different transformations. Specifically, we characterize the contingency modality activation by defining some key metrics and report them in Table 6 for Type-1 instances (see Appendix H for similar results on Type-2 instances). Columns “Nodes (%)” corresponds to the average percentage of nodes in the scenario tree with an active contingency modality, and columns “# of Contingencies” provide the average number of contingency modalities used by a policy among all instances. Columns “Aggressiveness (%)” refer to the percentage of capacity increase (between 10% and 40%) resulting from the contingency modality activation, and “Intensity” corresponds to the hurricane’s intensity when a contingency modality is first activated. We note that the average values in the second and third columns are computed only over nodes with active modalities. For example, PM only activates one modality for the instance with 20% initial capacity, which is an aggressive modality (i.e., 30%) activated when the hurricane reaches intensity three or higher.

From Table 6, we see that a lower initial capacity typically leads to more contingency modality activation and modality choices that are more aggressive. Comparing HN with other transformations, we see that it has more active nodes when the initial capacity is lower, and has active fewer nodes when the initial capacity is higher. This is attributed to the structure of the policy under HN which resembles an “all-or-nothing” feature due to the lack of adaptability. The behaviors of HN and MA are quite similar, although the activation timing of MA is later, resulting in a different number of nodes with an active modality. We attribute this similarity to the “propagation effect” in MC-based aggregation policies, which we discussed in Section 5.4. In contrast, PM and MM result in more adaptive policies that resemble a “wait-and-activate” type of behavior, similar to FH. In these more adaptive policies, contingency modalities tend to be activated when the hurricane reaches a higher intensity state. In addition, the more adaptive policies lead to a more diverse set of modality choices and a better balance between aggressiveness and the number of nodes actives depending on the initial capacity. Comparing the MC-based transformations PM and MM with FH, which has full adaptability, we see the value of additional adaptability in activating less nodes throughout the scenario tree with more diverse contingency modalities with appropriate aggressiveness.




Figure 4 Activation solutions over an abstract scenario tree for different MC-based transformations.

5.4. Implications of the MC-based Aggregation in HDR

We now analyze impacts of the MC-based aggregation on the resulting HDR solution policies. In particular, we study why the number of possible contingency modalities is low for some of the transformations (e.g., at most one for MA) in our test instances. First, recall that binary variables $z_{n\ell}$ represent whether contingency modality $\ell \in \mathcal{L}$ is active in node $n \in \mathcal{T}$. In addition, the model imposes that only one modality is active in each state (10e) and every active modality remains active in the following stages (10f). Thus, these constraints restrict the number of different modalities that can be active in any solution (including FH), especially when an aggregation is applied.

To understand the implications of an MC-based aggregation over nodes in the scenario tree with active modalities, Figure 4 illustrates solutions of different transformations in an abstract scenario tree with two states (from Example 1). Specifically, \blacksquare are nodes where a modality is *initially activated* (as opposed to passively activated due to propagation, as explained below), \blacklozenge are nodes that have active modalities because of prior activation at their parent nodes, and $\color{orange}\blacklozenge$ are nodes where the modality is *passively activated* (i.e., the activation is a consequence of the propagation caused by the transformation that forces certain nodes to have the same z values). For example, we can see that the MA policy initially activates a modality in $t = 2$ at the dark node. However, this causes the two $\color{orange}\blacklozenge$ nodes to be passively activated because nodes with the same MC state must have the same z values according to transformation MA. This propagation, as a result of the MC-based aggregation, may cause an initial activation to have “unintended consequences” of passively activating nodes that may not necessarily benefit from the modality activation. This explains the similarity between HN and MA that we observe in Section 5.3. We can also observe such propagation behaviors with PM and MM but they occur less often since the aggregation depends on the MC state of the current stage and the previous one. In contrast, the activation when using FH does not have any such propagation effect since no aggregation is imposed (i.e., no $\color{orange}\blacklozenge$ can be observed from FH).

Lastly, Figure 5 illustrates the resulting policies for a single instance in a series of images corresponding to the grid of the problem where the hurricane originates at $y = 0$ (i.e., bottom row)




Figure 5 Graphical representation of the solutions of an instance of Type-1 with 20% initial capacity.

and lands at $y = 5$ (i.e., top row). The top grids show the percentage of nodes in the scenario tree with an active contingency modality for each cell. The bottom grids are heat-maps showing the proportion of nodes in a cell where a contingency modality is the first activated (either initially activated or passively activated), and the color scheme corresponds to the hurricane intensity level (i.e., darker colors for higher intensity). Each row is associated with a symbol \blacksquare or \bullet , which distinguishes if the nodes are initially or passively activated, respectively. We observe that MA, PM, and MM lead to a wait-and-activate type behavior by activating modalities in high-intensity nodes at later stages to reduce the number of active modalities in the last stage, a similar strategy to FH. In contrast, HN activates a modality in the first stage even though it may not be necessary for many nodes in later stages. We also observe that the intensity levels of \bullet nodes tend to be much smaller than the ones associated with \blacksquare nodes. This reflects the “unintended consequence” of forcing nodes of low intensity to be passively activated due to the restrictions imposed by the transformations.

6. Conclusions

We present an aggregation framework to address MSILP problems with mixed-integer state variables and continuous local variables. The main idea of the framework is to reformulate the original MSILP such that all integer state variables are first-stage and then apply a suitable transformation to aggregate these variables to reduce the problem size while not sacrificing the solution quality. We describe a number of transformations that rely on the underlying structure of the stochastic process, which we assume is given by an MC. We demonstrate an exact solution approach based on a B&C framework integrated with the SDDP algorithm to solve the resulting aggregated MSILP. Also, we propose a computationally more tractable alternative based on the 2SLDR approximations

and introduce new MC-based decision rules. These approaches are tested on a hurricane disaster relief logistics application with contingency modality decisions. Empirical results show the trade-off between different transformations regarding the solution quality and computational effort, and the best trade-off is yielded by transformations that consider information from current and previous MC states (i.e., PM). The SDDP lower and upper bounding procedures have the best performance in terms of solution quality for large-scale instances, while the MC-based 2SLDR approach returns high-quality solutions with significantly less computational effort.

As shown by our empirical results, the performance of our aggregation framework largely depends on the chosen transformation. With this in mind, one future research direction is to investigate how to choose an appropriate transformation for a given problem instance. An alternative research direction is to develop a dynamic extension of the current approach to iteratively refine the aggregation over time, e.g., in a rolling horizon procedure.

References

- Ahmed S, Cabral FG, Freitas Paulo da Costa B (2020) Stochastic Lipschitz dynamic programming. *Mathematical Programming* 1–39.
- 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.
- Bally V, Pagès G (2003) A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli* 9(6):1003–1049.
- Bampou D, Kuhn D (2011) Scenario-free stochastic programming with polynomial decision rules. *2011 50th IEEE Conference on Decision and Control and European Control Conference*, 7806–7812 (IEEE).
- Basciftci B, Ahmed S, Gebrael N (2019) Adaptive two-stage stochastic programming with an application to capacity expansion planning. *arXiv preprint arXiv:1906.03513* .
- Birge JR (1985) Decomposition and partitioning methods for multistage stochastic linear programs. *Operations Research* 33(5):989–1007.
- Birge JR, Louveaux FV (1988) A multicut algorithm for two-stage stochastic linear programs. *European Journal of Operational Research* 34(3):384–392.
- Bodur M, Dash S, Günlük O, Luedtke J (2017) Strengthened Benders cuts for stochastic integer programs with continuous recourse. *INFORMS Journal on Computing* 29(1):77–91.
- Bodur M, Luedtke JR (2022) Two-stage linear decision rules for multi-stage stochastic programming. *Mathematical Programming* 191:347—380.
- Bonnans JF, Cen Z, Christel T (2012) Energy contracts management by stochastic programming techniques. *Annals of Operations Research* 200(1):199–222.

- Chen X, Sim M, Sun P, Zhang J (2008) A linear decision-based approximation approach to stochastic programming. *Operations Research* 56(2):344–357.
- Davis LB, Samanlioglu F, Qu X, Root S (2013) Inventory planning and coordination in disaster relief efforts. *International Journal of Production Economics* 141(2):561–573.
- Dowson O (2020) The policy graph decomposition of multistage stochastic programming problems. *Networks* 76(1):3–23.
- Duran S, Gutierrez MA, Keskinocak P (2011) Pre-positioning of emergency items for care international. *Interfaces* 41(3):223–237.
- Füllner C, Rebennack S (2021) Stochastic dual dynamic programming and its variants, preprint, Karlsruhe Institute of Technology.
- Füllner C, Rebennack S (2022) Non-convex nested Benders decomposition. *Mathematical Programming* 1–38.
- Graumann A, Houston TG, Lawrimore JH, Levinson DH, Lott N, McCown S, Stephens S, Wuertz DB (2006) Hurricane Katrina: A climatological perspective: Preliminary report. Technical Report 2005-01, NOAA’s National Climatic Data Center, Washington, DC.
- Hjelmeland MN, Zou J, Helseth A, Ahmed S (2018) Nonconvex medium-term hydropower scheduling by stochastic dual dynamic integer programming. *IEEE Transactions on Sustainable Energy* 10(1):481–490.
- Kuhn D, Wiesemann W, Georghiou A (2011) Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming* 130(1):177–209.
- Lodree Jr EJ, Ballard KN, Song CH (2012) Pre-positioning hurricane supplies in a commercial supply chain. *Socio-Economic Planning Sciences* 46(4):291–305.
- Löhndorf N, Shapiro A (2019) Modeling time-dependent randomness in stochastic dual dynamic programming. *European Journal of Operational Research* 273(2):650–661.
- Morton DP (1996) An enhanced decomposition algorithm for multistage stochastic hydroelectric scheduling. *Annals of Operations Research* 64(1):211–235.
- Pacheco GG, 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.
- Pagès G, Pham H, Printems J (2004) An optimal markovian quantization algorithm for multi-dimensional stochastic control problems. *Stochastics and Dynamics* 4(04):501–545.
- Pereira MV, Pinto LM (1991) Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52(1):359–375.
- Philpott AB, De Matos VL (2012) Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research* 218(2):470–483.

-
- Philpott AB, Guan Z (2008) On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters* 36(4):450–455.
- Regnier ED, MacKenzie CA (2019) The hurricane decision simulator: A tool for marine forces in new orleans to practice operations management in advance of a hurricane. *Manufacturing & Service Operations Management* 21(1):103–120.
- 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.
- Seraphin H (2019) Natural disaster and destination management: The case of the caribbean and hurricane irma. *Current Issues in Tourism* 22(1):21–28.
- Shapiro A, Dentcheva D, Ruszczyński A (2021) *Lectures on stochastic programming: modeling and theory* (SIAM).
- Shapiro A, Nemirovski A (2005) On complexity of stochastic programming problems. *Continuous optimization*, 111–146 (Springer).
- Siddig M, Song Y (2022) Multi-stage stochastic programming methods for adaptive disaster relief logistics planning. *arXiv preprint arXiv:2201.10678* .
- Taskin S, Lodree Jr EJ (2010) Inventory decisions for emergency supplies based on hurricane count predictions. *International Journal of Production Economics* 126(1):66–75.
- Van Slyke RM, Wets R (1969) L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics* 17(4):638–663.
- Yang H, Duque D, Morton DP (2022) Optimizing diesel fuel supply chain operations to mitigate power outages for hurricane relief. *IIEE Transactions* 1–14.
- Zhang S, Sun XA (2022) Stochastic dual dynamic programming for multistage stochastic mixed-integer nonlinear optimization. *Mathematical Programming* 1–51.
- Zou J, Ahmed S, Sun XA (2018) Multistage stochastic unit commitment using stochastic dual dynamic integer programming. *IEEE transactions on Power Systems* 34(3):1814–1823.
- Zou J, Ahmed S, Sun XA (2019) Stochastic dual dynamic integer programming. *Mathematical Programming* 175(1):461–502.
- Zverovich V, Fábíán CI, Ellison EF, Mitra G (2012) A computational study of a solver system for processing two-stage stochastic LPs with enhanced Benders decomposition. *Mathematical Programming Computation* 4(3):211–238.

Appendix A: B&C Integrated with the SDDP algorithm

Algorithm 2 describes the B&C procedure. The procedure receives as input a few parameters used by an SDDP sub-routine to be explained in detail later. Note that lines 4-7 correspond to the standard B&B algorithm of commercial solvers.

Algorithm 2 A B&C framework integrated with the SDDP algorithm

```

1: procedure SDDP-B&C(exact,  $K$ ,  $\epsilon$ )
2:   Initialize B&B search for  $\overline{Q}_r^A$  (master problem) and the SDDP environment.
3:   Initialize  $\mathcal{B}_n := \emptyset$  for all  $n \in \mathcal{C}(r)$  and cutAdded := false.
4:   repeat
5:     Choose a B&B node from its list of open nodes. Solve the associated node relaxation problem.
6:     if Current B&B node is infeasible or the relaxation bound is worse than the bound given by an
       incumbent solution then
7:       Prune this B&B node.
8:       if Values  $\widehat{z}^A$  of the current solution  $(\widehat{x}_r, \widehat{z}^A, \widehat{\theta})$  are integer then
9:         repeat                                     %% Cutting plane subroutine
10:          cutAdded := false.
11:          for  $n' \in \mathcal{C}(r)$  do                       %% Find cuts parameters for each child subproblem
12:            Get cut coefficients  $(\alpha_{n'}, \beta_{n'}, \gamma_{n'}) := \text{SDDP}((\widehat{x}_r, \widehat{z}^A, \widehat{\theta}), n', \epsilon, K, \text{exact})$ .
13:            if Cut with coefficients  $(\alpha_{n'}, \beta_{n'}, \gamma_{n'})$  is violated by  $(\widehat{x}_r, \widehat{z}^A, \widehat{\theta})$  then
14:               $\mathcal{B}_{n'} := \mathcal{B}_{n'} \cup (\alpha_{n'}, \beta_{n'}, \gamma_{n'})$ , cutAdded := true.
15:            until cutAdded = false.
16:          else
17:            Branch on the current B&B node and update the list of open nodes.
18:          until The list of open B&B nodes is empty.

```

There are a few important points to mention about this procedure. First, we call the SDDP sub-routine for each child of the root node r , in which we perform the SDDP forward pass only considering sample paths that contain that specific child. Also, we look for cuts for all children nodes of the root node r before re-solving the current B&B node relaxation. An alternative would be to re-solve the B&B node relaxation whenever we find a violated cut, but this option showed worse performance in our preliminary experiments. Lastly, all cuts added to enhance the outer approximation of the cost-to-go functions associated with the SDDP subproblems in the previous iterations are automatically carried over to the current iteration.

Appendix B: Benders Decomposition for the 2SLDR Model

In this section we present a Benders decomposition approach to solve (P^L) . We consider a Benders master problem that handles the first-stage decisions and decompose the second-stage problem into one Benders subproblem per node. Because the number of subproblems can be exponentially many (e.g., when LDR-TH is used), we consider an aggregated approach where we have one variable $\theta_{t,m}^L$ that approximates all the cost-to-go functions associated with nodes in stage t and with MC state $m \in \mathcal{M}_t$. This variant can be seen as a hybrid between the single-cut version (i.e., using a single θ variable) and the multi-cut version (i.e., using one $\theta_{t,n}^L$ variable for each subproblem) of Benders decomposition (Van Slyke and Wets 1969, Birge and

Louveaux 1988). Our preliminary experiments show that our Benders variant performs the best on our test instances, however, other variants of Benders decomposition and computational enhancements can also be considered (see, e.g., Zverovich et al. (2012) and Bodur et al. (2017)).

In what follows we present the details of our Benders decomposition for LDR-TH (the other two cases, LDR-T and LDR-M are similar). The Benders master problem is given by

$$\begin{aligned} \bar{Q}_r^L = \min & \sum_{n \in \mathcal{N}} p_n (c_n^\top z_{\phi_t(n)}^A + d_n^\top (\mu_{t(n)}^\top \xi_n^{t(n)})) + h_r^\top y_r + \sum_{t=2}^T \sum_{m \in \mathcal{M}_t} \theta_{t,m}^L \\ \text{s.t.} & \text{(8a) - (8c),} \\ & \theta_{t,m}^L \geq \alpha_{t,m}^L \mu + \beta_{t,m}^L z^A + \gamma_{t,m}^L, \quad \forall (\alpha_{t,m}^L, \beta_{t,m}^L, \gamma_{t,m}^L) \in \mathcal{B}_{t,m}^L, \quad m \in \mathcal{M}_t, \quad t \in \{2, \dots, T\}, \\ & (y_r, \mu, z^A) \in \mathbb{R}^r \times \mathbb{R}^{k \times \sum_{t \in [T]} t^t} \times \mathbb{Z}^{\ell \cdot \sum_{t \in [T]} q_t}, \end{aligned} \quad (11)$$

where $\theta^L \in \mathbb{R}^{\sum_{t=2}^T |\mathcal{M}_t|}$ is introduced to represent the cost-to-go function approximation for each Markovian state in all remaining stages (except for the first stage) via constraints (11). Here $\mathcal{B}_{t,m}^L$ stores the coefficients associated with the Benders cuts for each MC state $m \in \mathcal{M}_t$.

Algorithm 3 Benders Decomposition

```

1: procedure BENDERS( $\epsilon$ )
2:   Initialize B&B search for  $\bar{Q}_r^L$  (master problem)
3:   repeat
4:     Choose a node B&B from its node list. Solve the associated node relaxation problem.
5:     if Current B&B node is infeasible or the relaxation bound is worse than the bound given by an
       incumbent solution then
6:       Prune this B&B node.
7:       if Current solution  $(\hat{\mu}, \hat{z}^A, \hat{\theta}^L)$  is integer then
8:         repeat %% Cutting plane subroutine
9:           for  $m \in \mathcal{M}_t$  and  $t \in \{2, \dots, T\}$  do
10:            for  $n \in \mathcal{N}_t(m)$  do
11:              Solve  $\hat{Q}_n^L := Q_n^L(\hat{\mu}, \hat{z}^A)$  and save dual solutions  $\pi^n$ .
12:              Save value  $\hat{Q}_{t,m}^L := \sum_{n \in \mathcal{N}_t(m)} p_n \hat{Q}_n^L$ .
13:              if  $|\hat{Q}_{t,m}^L - \hat{\theta}_{t,m}^L| \geq \epsilon |\hat{Q}_{t,m}^L|$  then %% Add cut if needed
14:                Add inequality (11) and exit the for loop.
15:            until No new inequality is added to the master problem  $\bar{Q}_r^L$ .
16:         else
17:           Branch on the current B&B node and update list of nodes.
18:         until No more nodes in the B&B node list.

```

Algorithm 3 describes our Benders decomposition for (P^L) . The procedure is a B&C algorithm similar to the one described in Algorithm 2. Once an integer solution is found during B&B search, we enter the cutting plane subroutine. The algorithm iterates over each stage and MC state and its corresponding nodes, where $\mathcal{N}_t(m)$ is the set of all nodes in stage t associated with MC state $m \in \mathcal{M}_t$. It then evaluates the Benders subproblems for all the nodes associated with a particular m and t and checks if the current approximation

$\widehat{\theta}_{t,m}^L$ is close enough to the actual value $\widehat{Q}_{t,m}^L$ for a given tolerance ϵ . If the approximation is not good enough, then the algorithm returns a cut and stops evaluating other subproblems. Similar to our SDDP variant, we add cuts to each candidate integer solution until no more cuts are found. The complete procedure ends when there are no more nodes to explore in the B&B search tree, i.e., we have found an optimal solution or proven infeasibility.

The cuts added to the master problem (11) are Benders cuts over all the nodes associated to a given stage and MC state. Specifically, the cut coefficients for $t \in \{2, \dots, T\}$ and $m \in \mathcal{M}_t$ are

$$\begin{aligned}\alpha_m^L &= \sum_{n \in \mathcal{N}_t(m)} p_n \pi^{n \top} A_n, & \beta_m &= \sum_{n \in \mathcal{N}_t(m)} p_n \pi^{n \top} B_n, \\ \gamma_m &= \sum_{n \in \mathcal{N}_t(m)} p_n \left(\widehat{Q}_n^L - (\pi^{n \top} A_n)^\top \widehat{\mu}_{t(n)-1} - (\pi^{n \top} B_n)^\top \widehat{z}_{\phi_t(a(n))}^A \right),\end{aligned}$$

where π^n are the dual solutions associated with the cost-to-go function $Q^L(n)$. Note that each coefficient here corresponds to a weighted sum of the associated coefficients in the subproblems. Also, the algorithm generates Benders feasibility cuts when necessary, which is common in 2SLDR approximations (Bodur and Luedtke 2022).

Appendix C: Aggregated Model for HDR Applicaton

We now present how to transform the MSILP model (*HDR*) with integer state variables at each node of the scenario tree to an MSILP model with integer variables only at the root node, following the ideas presented in Section 2 and Section 3. The general idea is to consider all integer state variables as first-stage variables and apply a suitable transformation Φ_t to aggregate the number of variables to a manageable size. In what follows, we present the resulting aggregated model (*HDR^A*) for a generic transformation and then show four alternatives given the specific structure of (*HDR*).

Following the notation introduced in Section 2, we consider z^A to be the set of aggregated integer state variables for each node. There are two aspects to consider when applying this transformation. First, constraints (10e) and (10f) are now first-stage constraints. Second, we can omit state variables x_n^c because these variables are only used to model the capacity dependency from one stage to the next, which is fully dictated by the first-stage variables z^A . The resulting aggregated model is:

$$\min \sum_{j \in \mathcal{J}} \left(g_j x_{rj}^I + q_{rj} v_j + \sum_{i \in \mathcal{I}} f_{rij} y_{ij} \right) + \sum_{i \in \mathcal{I}} b_i w_i + \sum_{n \in \mathcal{N}} \sum_{\ell \in \mathcal{L}} z_{\phi_t(n)\ell}^A c_\ell + \sum_{n \in \mathcal{C}(r)} \bar{p}_{rn} Q_n^A(x_r^I, z^A) \quad (\text{HDR}^A)$$

s.t. (9a) – (9b)

$$v_j \leq C_j, \quad \forall j \in \mathcal{J}, \quad (12a)$$

$$\sum_{\ell \in \mathcal{L}} z_{\phi_t(n)\ell}^A \leq 1, \quad \forall n \in \mathcal{N}, \quad (12b)$$

$$z_{\phi_{t-1}(a(n))\ell}^A \leq z_{\phi_t(n)\ell}^A, \quad \forall \ell \in \mathcal{L}, n \in \mathcal{N} \setminus \{r\}, \quad (12c)$$

$$x_{rj}^I, v_j, w_i, y_{ij} \geq 0, z_{\phi_t(n)\ell}^A \in \{0, 1\}, \quad \forall j \in \mathcal{J}, i \in \mathcal{I}, \ell \in \mathcal{L}, n \in \mathcal{N} \setminus \{r\},$$

where the cost-to-go function for a node $n \in \mathcal{N} \setminus \{r\}$ is given by:




Figure 6 A 6×4 grid example that shows the initial location of the hurricane $(m^x, m^y) = (5, 0)$ and all possible locations in the following stages (i.e., shaded region).

$$\begin{aligned}
Q_n^A(x_{a(n)}^I, z^A) &= \min \sum_{j \in \mathcal{J}} \left(g_j x_{nj}^I + q_{nj} v_j + \sum_{i \in \mathcal{I}} f_{nij} y_{ij} \right) + \sum_{i \in \mathcal{I}} b_i w_i + \sum_{n' \in \mathcal{C}(n)} \bar{p}_{nn'} Q_{n'}^A(x_n^I, z^A) \\
&\text{s.t. (10a) - (10b)} \\
v_j &\leq C_j + \sum_{n' \in \mathcal{P}(n)} \sum_{\ell \in \mathcal{L}} K_{j\ell} z_{\phi_t(n')\ell}^A, & \forall j \in \mathcal{J}, & \quad (13a) \\
x_{nj}^I, v_j, w_i, y_{ij} &\geq 0, & \forall j \in \mathcal{J}, i \in \mathcal{I}. &
\end{aligned}$$

Note that the right-hand side of constraint (13a) represents the current capacity of a DC by considering the modality activation and capacity increases at previous stages. Also, although constraints (12b) and (12c) are imposed for each node, the fact is that some of them might be identical because they share the same problem parameters and aggregated variables. For example, transformation HN only needs one set of constraints (12b) per stage because of the stage-based variable aggregation.

Appendix D: Problem Description and Instance Generation for HDR

We now describe the main components of our HDR problem and present details of the instance generation procedure. We focus on the grid representation for the problem, the MC, and three main parameters (i.e., demand, modalities, and capacity). All other details about instances generation (e.g., cost parameters) can be found in our instances generation code. We will make the instance generation code and the set of instances publicly available upon publication.

D.1. Grid Representation and Locations of Shelters and DCs

Similar to previous works in the literature, we use a grid network to represent the potential locations of shelters and DCs, as well as possible locations of the hurricane (see Figure 6). The top row on the grid represents the land, where each cell contains a subset of shelters and DCs. The remaining cells in the grid correspond to the sea and are used to represent possible locations of the hurricane at any stage during the planning horizon. For simplicity, we assume that any shelter can be supplied by any DC in the network.

Land and sea cells have pre-defined dimensions: land cells have a size of 50×100 , while sea cells have a size of 20×100 . We randomly generate 3 to 7 shelters and 2 to 4 DCs in each land cell. Shelters and DCs are positioned uniformly at random inside their corresponding cells. The coordinates of the hurricane located in a cell are assumed to be the coordinates of the center of the cell.

D.2. MC description

As described in the main text, each MC state corresponds to two attributes of the hurricane: location and intensity, that is, $m = (m^x, m^y, m^i) \in \mathcal{M}$ where m^x and m^y are the x - and y -coordinates of the hurricane's location, respectively, and m^i represents the hurricane's intensity level. The MC is given by two independent transition probability matrices: one for the hurricane intensity level and one for its location. The intensity level transition matrix is identical to the one used by Pacheco and Batta (2016), which considers six levels of intensity $\{0, 1, \dots, 5\}$ where level 5 corresponds to the maximum hurricane intensity and level 0 corresponds to the case when the hurricane dissipates. The following matrix shows the transition probabilities where P_{ij} represents the probability of transitioning from intensity level i to j .

$$P = \begin{matrix} & \begin{matrix} 0 & 1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0.11 & 0.83 & 0.06 & 0 & 0.6 & 0 \\ 0 & 0.15 & 0.6 & 0.25 & 0 & 0 \\ 0 & 0 & 0.04 & 0.68 & 0.28 & 0 \\ 0 & 1 & 0 & 0.18 & 0.79 & 0.03 \\ 0 & 0 & 0 & 0 & 0.5 & 0.5 \end{pmatrix} \end{matrix}$$

For the hurricane movement, we assume that the hurricane originates at the bottom row of the grid (i.e., $m^y = 0$ for the initial MC state) and that m^y increases by one in each period (i.e., the hurricane advances upwards by one step in each period) and will reach land in exactly T periods. The hurricane can also move on the x -axis by either staying in the same x -coordinate, moving one cell to the left or one cell to the right. To determine the transition probabilities of each movement, we assign random weights to each movement and normalize them. In particular, we assign a uniform weight between 30 and 40 for staying in the same x -coordinate and a weight between 20 and 40 for moving to the right and left cells. If the current x -coordinate is at one of the grid borders (e.g., $m^x = 0$), then the probability of stepping out of the grid is zero. As an example, consider a 4×5 grid, current hurricane location $(m^x, m^y) = (1, 1)$, and weights 25, 36, and 22, for moving to locations $(0, 2)$, $(1, 2)$, $(2, 2)$, respectively. Then, the transition probability of moving from location $(1, 1)$ to $(0, 2)$ is $\frac{25}{83} \approx 0.3$.

The hurricane movement transition probability matrix is generated at random for each instance. The initial x -coordinate is chosen uniformly at random considering all possible options. The initial intensity level is also generated with a uniform random distribution considering hurricane intensity levels of two or higher.

D.3. Demand Generation

The demand for each shelter depends on the current MC state of the hurricane (i.e., location and intensity). We consider a maximum demand parameter d^{\max} and a maximum distance parameter δ^{\max} to model this dependency. The maximum demand d^{\max} is chosen at random between values 1000 and 1500 for each instance and represents the maximum demand for all shelters in a cell. The maximum demand at a cell is randomly distributed across all shelters in that cell. Therefore, each shelter $j \in \mathcal{I}$ has a maximum demand d_j^{\max} and the total maximum demand of all shelters in the same cell is equal to d^{\max} .

The maximum distance $\delta^{\max} = 150$ represents the distance of influence of the hurricane, i.e., all shelters which are farther than δ^{\max} from the hurricane have zero demand. If the distance between a shelter $j \in \mathcal{I}$ and the hurricane is less than δ^{\max} , then the shelters demand for MC state $m = (m^x, m^y, m^i) \in \mathcal{M}$ is given by:

$$d_{j,m} = d_j^{\max} \left(1 - \frac{\delta_{m,j}}{\delta^{\max}}\right) \cdot \left(\frac{m^i}{5}\right)^2,$$

where $\delta_{m,j}$ is the Euclidean distance between the current location of the hurricane and the shelter. Thus, the demand generation is linear with respect to the hurricane location and quadratic with respect to the hurricane intensity. Note that the demand is zero if the intensity level is 0 and that the maximum demand is achieved when the hurricane is on top of the shelter and has the highest intensity level (i.e., $\delta_{m,j} = 0$ and $m^i = 5$).

D.4. Initial Capacity, Contingency Modality Options, and Capacity Expansion

The initial capacity of a DC depends on both d^{\max} and the initial capacity percentage defined in Section 5. The total initial capacity over all DCs in a cell is given by $C^{\text{ini}} = d^{\max} \cdot p$, where $p \in \{0.2, 0.25, 0.3\}$ is the initial capacity percentage. The total initial capacity of a cell C^{ini} is then distributed uniformly at random across all DCs in that cell. Thus, the initial capacity C_j of a particular DC $j \in \mathcal{J}$ is given by a portion of C^{ini} .

We generate different contingency modality options depending on the grid size (i.e., number of land cells) and the modality Type (i.e., either Type-1 or Type-2, see Section 5). Our modality set \mathcal{L} considers options that increase DCs capacity in adjacent cells and options that increase the capacity for all DCs. For example, in a 4×5 grid we consider modality options that increase capacity for DCs in cells 0 and 1, 1 and 2, 2 and 3, and options that increase the capacity in all DCs, that is, a total of four different contingency modality options based on location.

In addition, we consider capacity increments of Type-1 or Type-2, each one with four different options: Type-1 has options $\{10\%, 20\%, 30\%, 40\%\}$ and Type-2 has options $\{15\%, 30\%, 45\%, 60\%\}$. Then each modality option in \mathcal{L} corresponds to a combination of the affected DCs based on location and the capacity increase. Considering a 4×5 grid example with Type-1 modalities, a possible modality option is to increase the capacity of DCs in cells 2 and 3 by 20% increments. In this example, we have a total of 16 modality options (i.e., four location options and four increments options).

The capacity increments for each modality are based on the initial capacity of the DCs. For example, consider the 4×5 grid with Type-1 modalities and active modality $\ell \in \mathcal{L}$ affecting cells 2 and 3 by 20% increments. Then, a DC $j \in \mathcal{J}$ in cell 0 has no capacity increments given the current active modality (i.e., $K_{j\ell} = 0$), but a DC $j' \in \mathcal{J}$ in cell 2 with initial capacity $C_{j'} = 100$ has capacity increments $K_{j'\ell} = 20$.

Appendix E: Model Size for HDR Instances

Tables 7 and 8 show the average number of variables and constraints for both small-size (4×5 grid) and large-size (5 grid) instances.

Table 7 Average number of variables and constraints for different aggregations over small-size instances.

| | HN | MA | PM | MM | FH |
|----------------------|---------|---------|---------|---------|---------|
| Continuous Variables | 566,370 | 566,370 | 566,370 | 566,370 | 566,370 |
| Integer Variables | 80 | 1,120 | 6,720 | 16,528 | 41,760 |
| Total Variables | 566,450 | 567,490 | 573,090 | 582,898 | 608,130 |
| Constraints | 120,129 | 124,834 | 129,373 | 137,171 | 164,414 |

Table 8 Average number of variables and constraints for different aggregations over large-size instances.

| | HN | MA | PM | MM | FH |
|----------------------|------------|------------|------------|------------|------------|
| Continuous Variables | 12,281,770 | 12,281,770 | 12,281,770 | 12,281,770 | 12,281,770 |
| Integer Variables | 120 | 2,300 | 13,800 | 46,800 | 733,240 |
| Total Variables | 12,281,890 | 12,284,070 | 12,295,570 | 12,328,570 | 13,015,010 |
| Constraints | 2,163,164 | 2,175,033 | 2,188,428 | 2,224,132 | 2,932,940 |

Table 9 Performance comparison between S and Ex for small-size instances (4×5 grid size).

| | | Average Time (sec) | | | | | | | | | | Gap (%) | | |
|----------|------|--------------------|-------|------------|-------|--------------|--------|--------------|--------|-------|------|---------|-----|----|
| | | HN | | MA | | PM | | MM | | FH | PM | | MM | FH |
| Modality | Cap. | Ex | S | Ex | S | Ex | S | Ex | S | Ex | S | S | Ex | |
| Type-1 | 20% | 260 | 1,221 | 712 | 5,969 | 961 | 4,553 | 1,853 | - | 1,402 | - | 14.9 | - | |
| | 25% | 132 | 195 | 361 | 1,313 | 487 | 2,760 | 626 | - | 3,974 | - | 3.5 | 0.8 | |
| | 30% | 73 | 102 | 109 | 272 | 222 | 3,380 | 367 | - | - | - | 2.1 | 0.4 | |
| Type-2 | 20% | 253 | 1,138 | 919 | 5,880 | 2,047 | 13,224 | 2,951 | 20,699 | 2,546 | 19.9 | 17.8 | 0.9 | |
| | 25% | 167 | 140 | 393 | 1,324 | 651 | 3,635 | 817 | - | - | - | 6.2 | 0.8 | |
| | 30% | 82 | 102 | 119 | 205 | 271 | 4,347 | 556 | - | - | - | 2.3 | 0.6 | |
| Average | | 161 | 483 | 436 | 2,494 | 773 | 5,316 | 1,195 | 20,699 | 2,641 | 19.9 | 7.8 | 0.7 | |

Appendix F: Computational Effort for Exact Methods

We now compare two solution methods for the aggregated models with four types of transformations, that is, solving the extensive form model directly with CPLEX (**Ex**) versus the B&C procedure integrated with the SDDP algorithm proposed in Section 3 (**S**). Table 9 shows the average solution time (for instances where optimal solutions are obtained within the time limit) and the average optimality gaps (for unsolved instances) for the small-size instances. We also include the results of solving **FH** with **Ex** as a point of reference. We note that **Ex** solves to optimality all the instances for all transformations except for 33 when solving **FH**, while **S** fails to solve one instance for **PM** and solves only one instance to optimality for **MM**.

From Table 9 we see a clear advantage of **Ex** over **S** in both computational time and number of instances solved within the time limit across all transformations and instance configurations. These results are mostly explained by the fact that approximately 90% of the computational time is spent inside the SDDP subroutine. One main factor is the number of SDDP subproblems for each transformation: **HN** and **MA** have on average 69 SDDP subproblems, **PM** has 144, and **MM** has 294, which explains the relatively poor performance of **MM**.

The previous results may suggest that **Ex** requires less computational effort for small-size instances than **S**. However, large-size instances cannot even be loaded into the solver due to the large number of variables and constraints in the formulation, that is, approximately 16GB of memory after the presolved phase (see Appendix E for further model size information), so **S** is our only resort for these instances. Table 10 shows the number of instances solved to optimality, the number of instances where the algorithm found a feasible

Table 10 Performance of the SDDP integrated B&C algorithm in large-size instances (5×6 grid size).

| | | # Optimal | | | | # Feasible | | | | Opt. Gaps (%) | | | |
|-----------|------|-----------|----|----|----|------------|----|----|----|---------------|------|------|------|
| Modality | Cap. | HN | MA | PM | MM | HN | MA | PM | MM | HN | MA | PM | MM |
| Type-1 | 20% | 2 | 0 | 0 | 0 | 8 | 10 | 7 | 1 | 53.4 | 51.8 | 56.1 | 82.0 |
| | 25% | 5 | 2 | 0 | 0 | 5 | 8 | 6 | 0 | 35.7 | 26.2 | 27.9 | - |
| | 30% | 10 | 8 | 0 | 0 | 0 | 1 | 4 | 0 | - | 15.4 | 4.3 | - |
| Type-2 | 20% | 3 | 0 | 0 | 0 | 7 | 10 | 7 | 1 | 43.5 | 52.9 | 61.4 | 86.0 |
| | 25% | 6 | 3 | 0 | 0 | 4 | 7 | 6 | 0 | 37.2 | 24.9 | 32.8 | - |
| | 30% | 9 | 8 | 0 | 0 | 1 | 1 | 6 | 0 | 12.0 | 13.3 | 6.4 | - |
| Total/Av. | | 35 | 21 | 0 | 0 | 25 | 37 | 36 | 2 | 36.4 | 30.7 | 31.5 | 84.0 |

solution (but could not prove optimality), and the average optimality gap over instances with a feasible solution for \mathbf{S} over large-size instances. We see that \mathbf{S} performs quite well for transformations HN and MA solving several instances to optimality; however, it fails to even find any feasible solution for MM. As before, these results are explained mainly by a large number of SDDP subproblems, and for these large-size instances, HN and MA have on average 114 subproblems, PM has 249, and MM has 593.

Appendix G: Additional Results for Approximation Methods

Tables 11 to 16 show the performance of the approximated methods based on 2SLDR and SDDP for small-size and large-size instances, and transformations HN, PM, and MM.

Table 11 Solution time and quality of 2SLDR and SDDP bounds. Results for HN over small-size instances.

| | | Average Time (sec) | | | | Relative Difference (%) | | | | |
|----------|------|--------------------|-------------|-------------|------|-------------------------|-------|-------------|-------------|------|
| Modality | Cap. | LDR-TH | LDR-T | LDR-M | S-LB | LDR-TH | LDR-T | LDR-M | S-UB | S-LB |
| Type-1 | 20% | 193.7 | 103.6 | 71.7 | 76.5 | 0.04 | 0.05 | 0.00 | 0.06 | 0.44 |
| | 25% | 132.2 | 31.3 | 33.1 | 15.2 | 0.05 | 0.07 | 0.00 | 0.00 | 0.63 |
| | 30% | 82.5 | 20.0 | 19.0 | 9.9 | 0.10 | 0.21 | 0.00 | 0.00 | 1.36 |
| Type-2 | 20% | 262.8 | 64.3 | 76.0 | 61.6 | 0.03 | 0.05 | 0.00 | 0.03 | 0.46 |
| | 25% | 131.7 | 29.3 | 35.8 | 12.4 | 0.05 | 0.07 | 0.00 | 0.01 | 0.62 |
| | 30% | 80.4 | 19.7 | 19.0 | 10.0 | 0.10 | 0.21 | 0.00 | 0.00 | 1.36 |
| Average | | 147.2 | 44.7 | 42.4 | 30.9 | 0.06 | 0.11 | 0.00 | 0.02 | 0.81 |

Table 12 Solution time and quality of 2SLDR and SDDP bounds. Results for HN over large-size instances.

| | | Average Time (sec) | | | | Opt. Gap (%) | | |
|-------------|------|--------------------|---------|--------------|-------|--------------|-------|-------------|
| Modality | Cap. | LDR-T | LDR-M | S-LB | (opt) | LDR-T | LDR-M | S-UB |
| Type-1 | 20% | 2,100.9 | 1,961.8 | 449.3 | (10) | 0.51 | 0.33 | 0.13 |
| | 25% | 1,679.4 | 1,628.6 | 288.0 | (10) | 0.78 | 0.39 | 0.15 |
| | 30% | 965.1 | 864.2 | 56.6 | (10) | 0.61 | 0.48 | 0.18 |
| Type-2 | 20% | 2,501.3 | 1,794.3 | 730.9 | (10) | 0.56 | 0.31 | 0.12 |
| | 25% | 1,660.5 | 1,764.6 | 203.7 | (10) | 0.78 | 0.40 | 0.13 |
| | 30% | 950.1 | 1,035.3 | 57.9 | (10) | 0.62 | 0.49 | 0.18 |
| Av. (Total) | | 1,642.9 | 1,508.1 | 297.7 | (60) | 0.64 | 0.40 | 0.15 |

Appendix H: Additional Results for Policy Structures and Managerial Insights

Table 17 presents the same metrics shown in Table 6 but for Type-2 small-size instances.

Table 13 Solution time and quality of 2SLDR and SDDP bounds. Results for MA over small-size instances.

| | | Average Time (sec) | | | | Relative Difference (%) | | | | |
|----------|------|--------------------|-------------|--------------|-------|-------------------------|-------|-------------|-------------|------|
| Modality | Cap. | LDR-TH | LDR-T | LDR-M | S-LB | LDR-TH | LDR-T | LDR-M | S-UB | S-LB |
| Type-1 | 20% | 431.1 | 104.2 | 95.2 | 367.3 | 0.26 | 0.29 | 0.23 | 0.08 | 0.19 |
| | 25% | 368.3 | 58.4 | 49.2 | 167.4 | 0.14 | 0.18 | 0.11 | 0.22 | 1.11 |
| | 30% | 119.3 | 25.7 | 25.4 | 17.1 | 0.11 | 0.22 | 0.02 | 0.00 | 1.56 |
| Type-2 | 20% | 630.6 | 119.4 | 117.1 | 439.4 | 0.12 | 0.14 | 0.10 | 0.09 | 0.26 |
| | 25% | 464.3 | 49.3 | 46.9 | 111.2 | 0.15 | 0.18 | 0.11 | 0.18 | 1.15 |
| | 30% | 113.8 | 22.7 | 23.1 | 16.4 | 0.10 | 0.21 | 0.01 | 0.00 | 1.58 |
| Average | | 354.5 | 63.3 | 59.5 | 186.5 | 0.15 | 0.21 | 0.09 | 0.10 | 0.97 |

Table 14 Solution time and quality of 2SLDR and SDDP bounds. Results for MA over large-size instances.

| | | Average Time (sec) | | | | Opt. Gap (%) | | |
|-------------|------|--------------------|----------------|----------------|-------|--------------|-------|-------------|
| Modality | Cap. | LDR-T | LDR-M | S-LB | (opt) | LDR-T | LDR-M | S-UB |
| Type-1 | 20% | 2,266.5 | 2,855.8 | 1,913.0 | (10) | 0.50 | 0.32 | 0.13 |
| | 25% | 3,067.1 | 2,458.6 | 1,295.8 | (10) | 0.84 | 0.46 | 0.19 |
| | 30% | 1,051.1 | 1,028.4 | 133.1 | (10) | 0.62 | 0.49 | 0.19 |
| Type-2 | 20% | 4,065.5 | 3,488.9 | 3,500.0 | (10) | 0.59 | 0.34 | 0.14 |
| | 25% | 4,295.9 | 4,107.4 | 1,960.6 | (10) | 0.89 | 0.40 | 0.18 |
| | 30% | 1,272.7 | 1,252.9 | 150.3 | (10) | 0.65 | 0.52 | 0.20 |
| Av. (Total) | | 2,669.8 | 2,532.0 | 1,492.1 | (60) | 0.68 | 0.42 | 0.17 |

Table 15 Solution time and quality of 2SLDR and SDDP bounds. Results for MM over small-size instances.

| | | Average Time (sec) | | | | Relative Difference (%) | | | | |
|----------|------|--------------------|--------------|--------------|---------|-------------------------|-------|-------------|-------------|------|
| Modality | Cap. | LDR-TH | LDR-T | LDR-M | S-LB | LDR-TH | LDR-T | LDR-M | S-UB | S-LB |
| Type-1 | 20% | 2766.2 | 262.2 | 277.0 | 3521.5 | 0.12 | 0.26 | 0.11 | 0.00 | 0.39 |
| | 25% | 1272.7 | 176.6 | 154.9 | 12447.8 | 0.09 | 0.13 | 0.04 | 0.25 | 0.98 |
| | 30% | 4049.5 | 235.3 | 310.3 | 14746.8 | 0.13 | 0.25 | 0.00 | 1.37 | 1.77 |
| Type-2 | 20% | 3916.3 | 504.6 | 396.0 | 11703.3 | 0.13 | 0.26 | 0.13 | 0.02 | 0.22 |
| | 25% | 1486.7 | 226.3 | 164.9 | 21226.9 | 0.09 | 0.12 | 0.01 | 0.38 | 2.30 |
| | 30% | 8042.7 | 431.0 | 720.4 | 15323.7 | 0.13 | 0.24 | 0.01 | 0.95 | 1.80 |
| Average | | 3589.0 | 306.0 | 337.2 | 13161.7 | 0.11 | 0.21 | 0.05 | 0.50 | 1.24 |

Table 16 Solution time and quality of 2SLDR and SDDP bounds. Results for MM over large-size instances.

| | | Average Time (sec) | | | | Opt. Gap (%) | | |
|-------------|------|--------------------|-----------------|----------|-------|--------------|--------------|-------|
| Modality | Cap. | LDR-T | LDR-M | S-LB | (opt) | LDR-T | LDR-M | S-UB |
| Type-1 | 20% | 8,288.1 | 4,776.9 | 20,125.0 | (1) | 39.71 | 33.00 | 39.39 |
| | 25% | 14,398.1 | 12,203.6 | 21,601.1 | (0) | 15.73 | 15.08 | 17.53 |
| | 30% | 1,911.4 | 2,076.6 | 21,601.6 | (0) | 2.41 | 0.67 | 2.09 |
| Type-2 | 20% | 14,668.7 | 13,135.3 | 21,600.7 | (0) | 39.55 | 37.42 | 43.68 |
| | 25% | 16,280.7 | 16,051.0 | 21,601.6 | (0) | 23.98 | 23.22 | 26.70 |
| | 30% | 2,391.2 | 1,925.4 | 21,601.4 | (0) | 3.63 | 1.51 | 3.50 |
| Av. (Total) | | 9,656.4 | 8,361.5 | 21,355.2 | (1) | 20.84 | 18.48 | 22.15 |

Table 17 Solution structures for different policies and initial capacities for Type-2 instances.

| Cap. | Nodes (%) | | | | | # of Contingencies | | | | | Aggressiveness (%) | | | | | Intensity | | | | |
|------|-----------|----|----|----|----|--------------------|-----|-----|-----|-----|--------------------|----|----|----|----|-----------|-----|-----|-----|-----|
| | HN | MA | PM | MM | FH | HN | MA | PM | MM | FH | HN | MA | PM | MM | FH | HN | MA | PM | MM | FH |
| 20% | 100 | 92 | 72 | 72 | 53 | 1.0 | 1.0 | 1.0 | 1.0 | 4.1 | 14 | 39 | 52 | 52 | 46 | 2.3 | 2.5 | 2.9 | 2.9 | 2.9 |
| 25% | 10 | 38 | 22 | 16 | 31 | 1.0 | 1.0 | 1.9 | 3.8 | 5.3 | 15 | 20 | 51 | 55 | 33 | 2.3 | 2.9 | 4.0 | 4.1 | 3.5 |
| 30% | 0 | 7 | 15 | 14 | 17 | - | 1.0 | 1.9 | 4.8 | 5.7 | - | 14 | 29 | 27 | 23 | - | 3.0 | 4.2 | 4.1 | 4.0 |